

Scanning Micro-Photoluminescence Studies of GaAs Photonic Crystals and Perovskite Structures



Luke P. Nuttall
St. Anne's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Hilary Term 2019

Abstract

The first part of this thesis focuses on a novel design of photonic crystal microcavity coupled to InGaAs quantum dots. Such coupled dot-cavity systems can be used as enhanced single photon sources for quantum information applications and more complicated arrangements could even be used as optical switches in a quantum computer. A photolithography process is used to fabricate these cavities, allowing them to overcome many of the difficulties involved in achieving reliable dot-cavity coupling in traditional e-beam defined cavities.

Theoretical FDTD simulations are used to predict the Q factor and mode volume ($1.44(\lambda_0/n)^3$) of this cavity design. The fabrication process is given in detail, and micro-photoluminescence measurements are used to verify successful cavity fabrication. A success rate of 85% is achieved with Q factors as high as 7.4×10^3 at a wavelength of around $1.25 \mu\text{m}$. These cavities are shown to have comparable performance to existing designs such as L3 and Notomi cavities fabricated using e-beam lithography.

The second part covers studies of four different polycrystalline perovskite films with compositions of the form $\text{FA}_{0.83}\text{Cs}_{0.17}\text{Pb}(\text{Br}_x\text{I}_{1-x})_3$ and varying bromine fraction $x \in \{0.1, 0.2, 0.3, 0.4\}$. These perovskites are promising candidates for commercially scalable photovoltaic applications and have received a great deal of scientific interest over the past decade. This particular composition has been shown to have improved stability and optoelectronic properties compared to other perovskites.

Micro-photoluminescence mapping is used to study the temperature dependence and structure of these samples. The diffusion lengths are found to be in the range from $2 \mu\text{m}$ to $5 \mu\text{m}$, and evidence of photon recycling over longer distances is identified. Time-resolved photoluminescence measurements are carried out at cryogenic temperatures to study the carrier decay dynamics. A theoretical model of the decay process is developed and fitted to the data. Both excitons and free carriers are found to contribute to the emission, with the 10% bromine sample having the highest exciton binding energy.

To my parents, Richard and Joan.

Acknowledgements

First I wish to express my sincere gratitude to my supervisor, Robert Taylor, for his invaluable guidance, support, encouragement, and patience throughout my time in Oxford. Despite his many responsibilities he has always found the time to discuss problems, look over results, provide assistance in the lab, and share his unparalleled wisdom in all areas of physics, wine, and the Welsh language. Robert's easygoing but dedicated mentorship has nurtured a warm and collaborative atmosphere in the group, and without him my DPhil would not have been nearly as fruitful nor enjoyable.

I am extremely grateful to all of the students, post-docs, and visitors in the Taylor group whom I have been fortunate enough to get to know. Ben Reid, Chris Chan, and Dave Coles all gave me a very warm welcome into the group and taught me many of the essential skills needed to carry out experiments. Stephen Lennon worked along side me throughout much of the SU-8 project, providing enormous help with both theoretical and practical aspects. Many of the later figures in chapter 5 are based on data we took together, and Stephen's FDTD simulations using Lumerical were very useful for verifying the accuracy of my own simulation results. Guanhua Ying worked with me on parts of the perovskite studies, assisting with both data collection and data analysis. In addition to those already mentioned, Youngsin Park, Tong Wang, Tim Puchtler, Claudius Kocher, Mo Li, and Vitaly Osokin have all provided many stimulating conversations and fun memories. I couldn't have asked for a better group of friends and colleagues to share my time in Oxford with.

I owe my gratitude to Fred Brossard and Jonathan Griffiths from HCL for carrying out all of the e-beam lithography and etching work on the PhC samples. Fred also provided invaluable insights into the intricacies of FDTD simulations and took the time to thoroughly proofread my paper. Jiang Wu (UCL) and Ed Clarke (University of Sheffield) deserve thanks for growing all of the GaAs samples used in this work. I'm grateful to Jacob Wang and Zhiping Wang from Henry Snaith's group in Oxford who produced the perovskite samples, and to Robin Nicholas for many interesting and helpful discussion of the results from those samples.

Additionally I wish to thank the numerous technical and support staff who keep the department clean and functional. In particular, Paul Pattinson deserves special mention for putting in an enormous amount of effort running the DWB

cleanroom, as do Robert Storey and Tony Hickman for always providing us with copious liquid helium, sometimes at very short notice. I also want to thank Ghassan Yassin for generously allowing us to use his very nice 5000 $\times$ optical microscope.

This research was only made possible by the funding generously provided by the Engineering and Physical Sciences Research Council (EPSRC) and Hitachi Cambridge Laboratory (HCL). I also wish to thank St. Anne's College for their financial assistance with travel expenses to China and for providing a friendly and welcoming home over the course of my studies here.

Another huge contribution to my DPhil came from the numerous contributors to free and open source software. Without the generous contributions made to software projects such as GNU+Linux, Meep, Python, NumPy, matplotlib, Inkscape, Blender, Gwyddion, and L^AT_EX, performing data analysis and writing this thesis would have been significantly more troublesome. Special mention should also go to the various individuals and organisations who form the broader Open Access movement in academic publishing.

I also want to thank the friends who've enriched my life outside of my studies. In particular I'm grateful to my flatmates Erhui Cai, Brittany Hause, and Ramy Aboushelbaya for all of the excellent conversations in the kitchen at RSH, and to the members of the Oxford University Role Playing Game Society (OURPGSoc) for helping me unwind and enjoy my time in Oxford.

Last but certainly not least I am eternally indebted to my parents for their limitless compassion, unyielding support, and for the countless hours they spent teaching me throughout my childhood. I have my dad to thank for instilling within me a love of science and problem solving, and I have my mum to thank for giving me the emotional stability required to complete this thesis. They have stuck by me through all the highs and lows of my DPhil.

Contents

List of Publications	viii
1 Introduction	1
2 Numerical Simulations of Photonic Crystals	3
2.1 Theory of Photonic Crystals	4
2.2 Finite Difference Time Domain Methods	9
2.3 Calculating the Band Structure of a 2D PhC	15
2.4 Modes of a PhC Waveguide	25
2.5 Simulating an SU-8 Cavity	31
2.6 Conclusions	44
3 Experimental Techniques	45
3.1 Introduction	45
3.2 Excitation Sources	46
3.3 Time-Integrated Micro-Photoluminescence	48
3.4 Micro-Photoluminescence Mapping	53
3.5 Time-Resolved Micro-Photoluminescence	59
3.6 Microscope Imaging	62
3.7 White Light Interferometry	62
3.8 Automation Techniques	65
4 Materials	68
4.1 Introduction	68
4.2 GaAs and other III-As Semiconductors	69
4.3 SU-8 Photoresist	82
4.4 Hybrid Organic-Inorganic Metal-Halide Perovskites	86
5 Lithographically Defined Photonic Crystal Microcavities	92
5.1 Introduction	93
5.2 Motivation and Background Theory	93
5.3 Device Design	105
5.4 Optimising SU-8 Photolithography	106

5.5	Cavity Fabrication and Optical Characterisation	125
5.6	Conclusions	134
6	Photoluminescence from Hybrid Perovskites	136
6.1	Introduction	136
6.2	Micro-Photoluminescence Studies	137
6.3	Time-Resolved μ PL Studies	157
6.4	Conclusions	179
7	Conclusions	182
7.1	PhC microcavities	182
7.2	Hybrid Perovskites	184
7.3	Closing Thoughts	186
Appendices		
A	Experiment Control Software	188
A.1	Spectrometer Control Software	188
A.2	Position Optimisation Pseudocode	191
B	FDTD Simulation Code	192
B.1	Band Structure of a hexagonal PhC	193
B.2	SU-8 Cavity Q Factor	200
C	Perovskite TRPL Fitting Code	211
C.1	Carrier Decay Simulation Code	212
C.2	Parameter Fitting Code	216
Bibliography		220

List of Publications

- **L. P. Nuttall**, F. S. F. Brossard, S. A. Lennon, B. P. L. Reid, J. Wu, J. Griffiths, and R. A. Taylor, “Optical fabrication and characterisation of SU-8 disk photonic waveguide heterostructure cavities”, *Optics Express*, 25, 24615-24622 (2017)
- S. A. Lennon, F. S. F. Brossard, **L. P. Nuttall**, J. Wu, J. Griffiths, and R. A. Taylor, “Photonic molecules defined by SU-8 photoresist strips on a photonic crystal waveguide”, *Optics Express*, 26, 32332-32345 (2018)
- Y. Park, C. C. S. Chan, B. P. L. Reid, **L. P. Nuttall**, R. A. Taylor, N.-S. Lee, and Y. M. Lee, “Carrier trapping and confinement in Ge nanocrystals surrounded by Ge_3N_4 ”, *Scientific Reports*, 6, 25449 (2016)
- T. Wang, T. J. Puchtler, T. Zhu, J. C. Jarman, **L. P. Nuttall**, R. A. Oliver, and R. A. Taylor, “Polarisation-controlled single photon emission at high temperatures from InGaN quantum dots”, *Nanoscale*, 9, 9421-9427 (2017)
- C. C. Kocher, T. J. Puchtler, J. C. Jarman, T. Zhu, T. Wang, **L. P. Nuttall**, R. A. Oliver, and R. A. Taylor, “Highly polarized electrically driven single-photon emission from a non-polar InGaN quantum dot”, *Applied Physics Letters*, 111, 251108 (2017)

1

Introduction

Ever since lenses were first used to light fires in ancient Greece, mastery of light has been one of the key pillars of human scientific advancement. Our understanding of optics has been essential for discoveries in every field from astronomy to biology. Without it we wouldn't have any of our modern technology, and lately the use of sunlight as a power source has become one of our best hopes for limiting the devastation of anthropogenic climate change.

In this thesis we present optical studies of two major types of semiconductor structures: photonic crystal microcavities coupled to quantum dots in GaAs and polycrystalline mixed-cation lead-halide perovskite films. The former have applications as single photon sources for quantum information technologies, while the latter are a promising candidate for next generation photovoltaics.

The following chapter covers the background theory of photonic crystals in addition to numerical simulations of various structures. The finite-difference time-domain simulation techniques are explained in detail, starting with calculations of the optical band structure for a simple 2D infinite lattice, then moving on to modelling waveguides and heterostructure cavities. Predictions are made for the frequency, quality factor, and mode volume of the cavities which are later fabricated in chapter 5.

Chapter 3 is an overview of the laboratory setup and methodology used in the experimental chapters. These mainly consist of micro-photoluminescence spectroscopy measurements, both time-integrated and time-resolved, carried out at cryogenic temperatures. The results of some characterisation measurements are also presented, quantifying the capabilities and limitations of the apparatus.

In the fourth chapter we discuss the different materials which the samples are composed from. Some physical, optical and electronic properties of both hybrid metal-halide perovskites and GaAs are presented, as well as an overview of the quantum confinement effects in low-dimensional semiconductors such as quantum wells and dots. There is also a section covering the chemistry and processing of SU-8 photoresist which forms part of the photonic crystal cavities.

Chapter 5 gives details of the experimental fabrication and characterisation of the photonic crystal microcavities. These optical cavities are fabricated using photolithographically patterned SU-8 on top of GaAs photonic crystal waveguides containing embedded InGaAs quantum dots. Micro-photoluminescence measurements at a temperature of 4.2 K are used to verify the successful creation of optical cavities and measure their quality factors. There is also a theoretical discussion of single photon sources and dot-cavity coupling in this chapter.

In the final experimental chapter the results of micro-photoluminescence studies of hybrid metal-halide perovskites are presented. The first half of the chapter deals with time-integrated measurements over a range of temperatures from 4.2 K to 300 K. Hyperspectral maps (2 spatial dimensions and 1 wavelength dimension) with variable excitation-collection separations are analysed to study the diffusion of photogenerated carriers. The second half of the chapter covers time-resolved micro-photoluminescence measurements at cryogenic temperatures. A sophisticated theoretical model of the carrier decay is developed and fitted to the experimental data.

2

Numerical Simulations of Photonic Crystals

Contents

2.1	Theory of Photonic Crystals	4
2.1.1	Photonic Crystal Defects	9
2.2	Finite Difference Time Domain Methods	9
2.2.1	Overview	9
2.2.2	The Yee Algorithm	10
2.2.3	Boundary Conditions and PML	13
2.2.4	Symmetry	14
2.3	Calculating the Band Structure of a 2D PhC	15
2.3.1	Overview	15
2.3.2	Selecting a k Vector	16
2.3.3	Initial Simulation	17
2.3.4	The Origin of Folded Bands	18
2.3.5	Approach 1: Phased Sources	20
2.3.6	Approach 2: Exploiting Symmetry	21
2.3.7	Approach 3: Phased Monitors	23
2.3.8	Final 3D Band Structure	24
2.4	Modes of a PhC Waveguide	25
2.4.1	Waveguide Band Structure	26
2.4.2	Fabry-Pérot Modes	30
2.5	Simulating an SU-8 Cavity	31
2.5.1	Cavity Structure	31
2.5.2	Finding the Cavity Frequency	32
2.5.3	Subpixel Averaging, Stability and Numerical Accuracy	37
2.5.4	Determining the Q Factor and Mode Volume	39
2.6	Conclusions	44




Figure 2.1: A schematic of a very simple 1D photonic crystal.

2.1 Theory of Photonic Crystals

A photonic crystal (PhC) is, in the most general terms, a structure with spatially varying optical properties (usually refractive index) which exhibits discrete translational symmetry along one or more lattice vectors. As the name suggests, photonic crystals are the optical analog to semiconductor crystals. Just as the wavefunctions of electrons are affected by the periodic lattice of varying electric potential created by atoms within a semiconductor crystal, the electromagnetic waves of light are affected by the periodic lattice of varying refractive index. Given the revolutionary versatility of semiconductors, it is therefore not surprising that there are a great variety of useful optical properties achievable using carefully designed PhCs.

The simplest example of a photonic crystal is an infinite one dimensional (1D) lattice with alternating layers of high and low refractive index material. This structure is depicted in Fig. 2.1. If we consider a plane wave travelling through the structure along the z axis, some part of the wave would be reflected at each

interface where the refractive index changes. Multiple interfaces produce multiple overlapping reflections which are then able to interfere with one another. This effect can be observed in a thin film of oil floating on water. Light is reflected from both the top and bottom surfaces of the film. At some wavelengths the reflections interfere constructively and are enhanced, while at other wavelengths the reflections are suppressed by destructive interference. This gives rise to the multicoloured appearance of thin films of oil and other transparent materials.

It is possible to produce useful interference effects with just two dielectric interfaces. For instance, many cheap anti-reflection coatings are made using a single layer of dielectric with a thickness equal to a quarter of the wavelength in the material. However, increasing the number of layers can greatly enhance such effects. For the theoretical case of an infinite photonic crystal, the reflectivity at a given wavelength and angle of propagation can be tuned from 0% to 100% simply by changing the thicknesses of the layers. These kind of 1D PhCs are used in high quality anti-reflection coatings, in dielectric (Bragg) mirrors, and in certain optical filters.

This simple consideration of optical path lengths and interference may be sufficient for calculating the properties of finite 1D PhCs, but for infinite or higher dimensional structures a more sophisticated mathematical treatment is necessary.

We begin with the macroscopic form of Maxwell's equations. We assume the material is linear, isotropic, and lossless, and we also assume there are no free charges or currents. Given these assumptions, Maxwell's equations become:

$$\begin{aligned} \nabla \cdot [\mu(\mathbf{r})\mathbf{H}(\mathbf{r}, t)] &= 0 & \nabla \times \mathbf{E}(\mathbf{r}, t) &= -\mu_0\mu(\mathbf{r})\frac{\partial\mathbf{H}(\mathbf{r}, t)}{\partial t} \\ \nabla \cdot [\epsilon(\mathbf{r})\mathbf{E}(\mathbf{r}, t)] &= 0 & \nabla \times \mathbf{H}(\mathbf{r}, t) &= \epsilon_0\epsilon(\mathbf{r})\frac{\partial\mathbf{E}(\mathbf{r}, t)}{\partial t} \end{aligned} \quad (2.1)$$

where ϵ_0 and μ_0 are the vacuum permittivity and permeability respectively, and $\epsilon(\mathbf{r})$ and $\mu(\mathbf{r})$ are the relative permittivity and permeability. For most dielectric materials $\mu(\mathbf{r})$ is very close to one, so we can make the maths slightly simpler by setting it equal to one.

Since any function can be written as a (possibly infinite) Fourier series, we do not lose any physical accuracy by restricting our analysis to a set of electromagnetic modes which vary sinusoidally in time. A given harmonic mode of angular frequency ω can thus be written as:

$$\begin{aligned}\mathbf{E}(\mathbf{r}, t) &= \mathbf{E}(\mathbf{r}) \exp(-i\omega t) \\ \mathbf{H}(\mathbf{r}, t) &= \mathbf{H}(\mathbf{r}) \exp(-i\omega t)\end{aligned}\tag{2.2}$$

Any two harmonic modes of differing ω are orthogonal, so substituting (2.2) into (2.1) lets us separate out the time dependence and write a new set of equations describing only the spatial variation of a given harmonic mode.

$$\begin{aligned}\nabla \cdot \mathbf{H}(\mathbf{r}) &= 0 & \nabla \times \mathbf{E}(\mathbf{r}) &= i\omega\mu_0\mathbf{H}(\mathbf{r}) \\ \nabla \cdot [\epsilon(\mathbf{r})\mathbf{E}(\mathbf{r})] &= 0 & \nabla \times \mathbf{H}(\mathbf{r}) &= -i\omega\epsilon_0\epsilon(\mathbf{r})\mathbf{E}(\mathbf{r})\end{aligned}\tag{2.3}$$

We can then combine the two curl equations to produce the “master equation” which depends only on $\mathbf{H}$.

$$\nabla \times \left[\frac{1}{\epsilon(\mathbf{r})} \nabla \times \mathbf{H}(\mathbf{r}) \right] = \left(\frac{\omega}{c} \right)^2 \mathbf{H}(\mathbf{r})\tag{2.4}$$

This equation is essentially just an eigenvalue problem where the left side is a Hermitian operator operating on the eigenvector, $\mathbf{H}(\mathbf{r})$, while the right side is $\mathbf{H}(\mathbf{r})$ multiplied by an eigenvalue, $(\omega/c)^2$. It is directly analogous to the eigenvalue problem of electron wavefunctions within a crystal lattice, except in that case the operator is the Hamiltonian. It is convenient to define this operator as:

$$\hat{\Theta} = \nabla \times \left[\frac{1}{\epsilon(\mathbf{r})} \nabla \times \right]\tag{2.5}$$

We now introduce a dielectric function with discrete translational symmetry, $\epsilon(\mathbf{r}) = \epsilon(\mathbf{r} + \mathbf{R})$, where $\mathbf{R}$ is a primitive lattice vector. It can be shown that each mode of such a system is simply a plane wave multiplied by a function,

$u_{\mathbf{k}}(\mathbf{r}) = u_{\mathbf{k}}(\mathbf{r} + \mathbf{R})$, which has the same periodicity as the lattice. This result is known as Bloch's theorem [1]. It was originally developed to model the wavefunctions of electrons in atomic lattices, but is equally applicable to photonic crystals.

$$\mathbf{H}_{\mathbf{k}}(\mathbf{r}) = u_{\mathbf{k}}(\mathbf{r})e^{i\mathbf{k}\cdot\mathbf{r}} \quad (2.6)$$

In free space any two plane waves with different wave vectors are orthogonal and cannot interact. The spatial Fourier transform of a wave of wave vector $\mathbf{k}$ is a single delta function at the point $\mathbf{k}$ in reciprocal space. If the wave vectors are different, there is no overlap. Inside a PhC, however, the $u_{\mathbf{k}}(\mathbf{r})$ factor modifies the spatial profile of the mode and hence introduces additional Fourier components. In general, the spatial Fourier transform of a PhC mode with wave vector $\mathbf{k}$ consists of a delta function at the point $\mathbf{k}$ in reciprocal space, plus some (possibly infinite) number of delta functions at points $\mathbf{k} + 2n\pi/\mathbf{R}$. This means that any two modes with wave vectors which differ only by an integer multiple of $2\pi/\mathbf{R}$ are likely to overlap in Fourier space and hence not be orthogonal. This naturally leads to the concept of the first Brillouin zone which is the region in $\mathbf{k}$ -space closest to the origin containing all the $\mathbf{k}$ vectors which are unique modulo $2\pi/\mathbf{R}$ [2]. Fig. 2.2 shows the shape of the first Brillouin zone for a simple square lattice. By exploiting additional symmetries of the lattice it is possible to reduce the first Brillouin zone down further to what is known as the irreducible Brillouin zone without any loss of information.

An alternative way to understand the first Brillouin zone is in terms of the phase shift of the EM wave from one primitive lattice cell to the next. Two modes are orthogonal if and only if they have different phase shifts. Adding $2\pi/\mathbf{R}$ to the $\mathbf{k}$ vector does not change this phase shift. Two modes with the same phase shift and the same frequency are functionally equivalent, so it does not make sense to plot them in different locations on a band diagram.

Another important concept to understand when dealing with 2D PhCs is the separability of solutions into transverse electric (TE) and transverse magnetic (TM) modes. This is a consequence of a mirror symmetry within a PhC structure. We




Figure 2.2: (a) A simple square lattice with the unit cell indicated in green. (b) The corresponding reciprocal lattice with the first Brillouin zone shown in green and the irreducible Brillouin zone (IBZ) shown in blue.

consider a mirror operator, $\hat{M}_z$, which reflects a vector field about the xy plane. If the dielectric function has the same mirror symmetry, then $\hat{\Theta}$ must commute with $\hat{M}_z$. It follows that if $\mathbf{H}(\mathbf{r})$ is an eigenvector of $\hat{\Theta}$, then so is $\hat{M}_z \mathbf{H}(\mathbf{r})$. Assuming no degeneracy, $\hat{M}_z \mathbf{H}(\mathbf{r}) = \alpha \mathbf{H}(\mathbf{r})$, where α is a constant. This is an eigenvalue problem, and because $\hat{M}_z \hat{M}_z \mathbf{H}(\mathbf{r}) = \mathbf{H}(\mathbf{r})$, we know that $\alpha = \pm 1$. A similar equation can also be derived for the electric field. Therefore the solutions to (2.4) can be separated into those which are symmetric under the mirror operation and those which are antisymmetric. $\mathbf{E}$ transforms as a vector while $\mathbf{H}$ transforms as a pseudovector, so the only non-zero components of the symmetric (TE) modes must be E_x , E_y , and H_z , while for the anti-symmetric (TM) modes they must be H_x , H_y , and E_z . For any PhC possessing mirror symmetry, these two polarizations are orthogonal and are generally considered independently. As we will see later, if the mirror symmetry is violated, this assumption breaks down and the two sets of modes can exchange energy.

We now have all of the theoretical basis needed to understand the band structure of a an infinite 2D photonic crystal. However, to take full advantage of PhCs when designing devices, we also need to understand the effect of introducing defects into the structure.

2.1.1 Photonic Crystal Defects

When considering 2D PhCs there are two main classes of defects: point-like defects (e.g. a missing hole or small group of holes) and line-like defects (e.g. an entire row of missing holes). Strictly speaking these defects break the symmetry of the entire PhC so the band structure model is no longer an exact model anywhere, however in practice the perturbation can be treated as localised around the defects. In effect they alter the local density of states either at a point or along a line respectively.

Introducing a point-like defect into a PhC with a photonic band gap can, depending on the exact parameters, create a peak in the local density of states at a frequency inside the band gap. In this case light of that frequency is allowed to exist at the site of the defect but is still prohibited from propagating into the surrounding PhC. Effectively this creates a localised optical mode which decays exponentially away from the defect. In other words, point-like defects in PhCs allow for the creation of optical cavities. These defects have become a mainstay of photonic structure engineering and many different designs exist. Two of the more common ones are H1 and L3 cavities, consisting respectively of one and three missing holes in a hexagonal 2D PhC [3, 4].

The same reasoning can be applied to line-like defects, only the dimensionality of the confinement is different. Rather than fully confining the light to create a cavity, line-like defects can confine the light to a 1D path and hence act as waveguides. One significant difference is that line-like defects still preserve translational symmetry along one dimension and hence their modes can be described by a band structure.

2.2 Finite Difference Time Domain Methods

2.2.1 Overview

Finite difference time domain (FDTD) algorithms involve dividing time and space up into discrete intervals and then iteratively solving the time evolution of a system by calculating the next state based on the current state and the length of the time step. The differential equations describing the physical system are converted into a




Figure 2.3: A diagram of a single unit cell of the Yee lattice which stores electric and magnetic field data in an FDTD simulation. The components of $\mathbf{E}_{[i,j,k]}$ and $\mathbf{H}_{[i,j,k]}$ are labelled.

set of finite difference equations describing the value at each point as a function of the values at neighbouring points in time or space. Finite difference methods have been used to solve problems in fluid dynamics since the 1920s [5], but the approach used in modern FDTD schemes for solving Maxwell's equations was first pioneered by Kane Yee [6] and further developed by Allen Taflove [7].

There are numerous different pieces of software which can perform electromagnetic FDTD simulations. The one used in this work is Meep, a GNU GPL licensed free software package developed by Oskooi *et al.* [8].

The following subsections will explain the key components of an electromagnetic FDTD simulation. Most of the information about the Yee algorithm, except where otherwise noted, is taken from Taflove's book on the subject [9].

2.2.2 The Yee Algorithm

One of the main advances of Yee's FDTD method is the use of a staggered grid of points for tracking both the electric and magnetic fields, as shown in Fig. 2.3.

It would be possible to solve for only a single field (either electric or magnetic) using a wave equation, but the Yee algorithm solves for both the $\mathbf{E}$ and $\mathbf{H}$ fields using the coupled Maxwell curl equations. This makes the simulation more robust and allows for a wider range of material properties and physical phenomena to be simulated. Each field component is tracked at a different location within a grid cell of the Yee lattice, and these locations are chosen such that each $\mathbf{E}$ field component is surrounded by a perpendicular loop of four $\mathbf{H}$ field components and vice versa. This arrangement makes it computationally faster to calculate the curl of each field and allows for second-order accuracy by making the finite difference equations take the form of central differences. Additionally, the Yee lattice has the advantage of implicitly enforcing the two divergence equations when the state is updated with the curl equations.

When performing the time stepping process, the FDTD algorithm follows a similar staggered approach as in the spatial domain. The magnetic field is calculated using the electric field data from half a time step ago, and vice versa for the electric field. These staggered update calculations are fully explicit so they can be computed quickly, but the derivatives used in the curl equations are central differences which gives second-order accuracy in the time domain. This time stepping scheme is often referred to as a leapfrog approach.

The six discrete finite difference equations for updating the $\mathbf{E}$ and $\mathbf{H}$ fields are as follows:

$$\begin{aligned}
H_x^{n+1/2}[i,j+1/2,k+1/2] &= H_x^{n-1/2}[i,j+1/2,k+1/2] + \frac{\Delta t}{\mu[i,j+1/2,k+1/2]\mu_0} \times \\
&\quad \left(\frac{E_y^n[i,j+1/2,k+1] - E_y^n[i,j+1/2,k]}{\Delta z} + \frac{E_z^n[i,j,k+1/2] - E_z^n[i,j+1,k+1/2]}{\Delta y} \right) \\
H_y^{n+1/2}[i+1/2,j,k+1/2] &= H_y^{n-1/2}[i+1/2,j,k+1/2] + \frac{\Delta t}{\mu[i+1/2,j,k+1/2]\mu_0} \times \\
&\quad \left(\frac{E_z^n[i+1,j,k+1/2] - E_z^n[i,j,k+1/2]}{\Delta x} + \frac{E_x^n[i+1/2,j,k] - E_x^n[i+1/2,j,k+1]}{\Delta z} \right) \\
H_z^{n+1/2}[i+1/2,j+1/2,k] &= H_z^{n-1/2}[i+1/2,j+1/2,k] + \frac{\Delta t}{\mu[i+1/2,j+1/2,k]\mu_0} \times \\
&\quad \left(\frac{E_x^n[i+1/2,j+1,k] - E_x^n[i+1/2,j,k]}{\Delta y} + \frac{E_y^n[i,j+1/2,k] - E_y^n[i+1,j+1/2,k]}{\Delta x} \right) \\
E_x^{n+1}[i+1/2,j,k] &= E_x^n[i+1/2,j,k] + \frac{\Delta t}{\epsilon[i+1/2,j,k]\epsilon_0} \times \\
&\quad \left(\frac{H_z^{n+1/2}[i+1/2,j+1/2,k] - H_z^{n+1/2}[i+1/2,j-1/2,k]}{\Delta y} + \frac{H_y^{n+1/2}[i+1/2,j,k-1/2] - H_y^{n+1/2}[i+1/2,j,k+1/2]}{\Delta z} \right) \\
E_y^{n+1}[i,j+1/2,k] &= E_y^n[i,j+1/2,k] + \frac{\Delta t}{\epsilon[i,j+1/2,k]\epsilon_0} \times \\
&\quad \left(\frac{H_x^{n+1/2}[i,j+1/2,k+1/2] - H_x^{n+1/2}[i,j+1/2,k-1/2]}{\Delta z} + \frac{H_z^{n+1/2}[i-1/2,j+1/2,k] - H_z^{n+1/2}[i+1/2,j+1/2,k]}{\Delta x} \right) \\
E_z^{n+1}[i,j,k+1/2] &= E_z^n[i,j,k+1/2] + \frac{\Delta t}{\epsilon[i,j,k+1/2]\epsilon_0} \times \\
&\quad \left(\frac{H_y^{n+1/2}[i+1/2,j,k+1/2] - H_y^{n+1/2}[i-1/2,j,k+1/2]}{\Delta x} + \frac{H_x^{n+1/2}[i,j-1/2,k+1/2] - H_x^{n+1/2}[i,j+1/2,k+1/2]}{\Delta y} \right)
\end{aligned}$$

The particular choice of time step, Δt , and grid spacing, Δx , are critically important for the numerical stability of the simulation. The relationship between these two quantities is given by the Courant factor which is defined as $C_f = \Delta t / \Delta x$. For a d dimensional system with minimum refractive index $n_{\min}$, the numerical stability condition is given by the relation:

$$C_f < \frac{n_{\min}}{\sqrt{d}}$$

This restriction can be understood by considering the speed at which information can propagate through the simulation volume [10]. The finite difference equations

which update the field at a particular grid point only depend on the field at the immediately adjacent grid points in the previous time step. This places an upper bound on the speed of information propagation of:

$$v_{\max} = \frac{\Delta x}{\Delta t \sqrt{d}}$$

This speed must be greater than or equal to the maximum speed of light in the simulation which is determined by the minimum refractive index. The dependence on the number of spatial dimensions arises due to the fact that light can propagate diagonally but information in the Yee lattice can only travel orthogonally. For physical accuracy it is also necessary to have at least about 10 grid points per wavelength, thus placing an upper bound on Δx .

2.2.3 Boundary Conditions and PML

The Yee lattice, like all good things, must come to an end at some point. The simplest way to terminate the simulation volume is simply to set all fields outside of it to zero. This is equivalent to surrounding it with walls made of a perfectly conductive metal with zero skin depth. Incident electromagnetic waves are perfectly reflected across all frequencies. Unfortunately, while simple to implement, this isn't usually what we want in a simulation. In most cases, including all of the ones considered in this thesis, we wish to simulate structures surrounded by a transparent medium of effectively infinite volume.

Another relatively simple type of boundary condition supported by most FDTD software is the periodic boundary condition, often referred to as a Bloch boundary. Essentially this makes each side of the simulation volume 'wrap around' to the opposite side. This type of boundary is particularly useful for simulating infinite periodic structures. More details of the use of periodic boundary conditions for calculating a PhC band structure are given in section 2.3.

The most generally useful kind of boundary conditions are not technically boundary conditions at all. First proposed in 1944 by Berenger [11], perfectly matched layers (PML) are actually regions of a non-physical material with arbitrarily

high absorption which must be placed inside the simulation volume. Despite this, PMLs are often thought of as boundary conditions because they absorb all incident light and effectively render the choice of actual boundary conditions irrelevant. While in the analytical case a PML is a perfect absorber, the discretisation of space in an FDTD simulation introduces small numerical errors which lead to small reflections. The best way to minimise these numerical errors is to make the PMLs thick and gradually ‘turn on’ their absorption over a number of grid points. When using PMLs it is important that the refractive index within the PML region is constant along any line perpendicular to the boundary, otherwise the changing refractive index can interfere with the proper functioning of the PML and introduce erroneous reflection or transmission.

2.2.4 Symmetry

It is possible to take advantage of the symmetry of the structure in order to speed up the simulation. The main symmetries of interest to us for simulations of 2D photonic crystals are mirror symmetries, although MEEP also supports certain types of rotational symmetry. If the structure possesses mirror symmetry about the $x = 0$, $y = 0$, or $z = 0$ planes, it is possible to reduce the simulation volume by a factor of 2 by imposing the same symmetry on the fields. This allows up to an $8\times$ increase in simulation speed if mirror symmetries are applied to all three axes.

When applying the mirror symmetry to the fields there are in fact two possible solutions corresponding to even and odd modes of the system, one of which must be selected. In exactly the same way that the presence of a mirror symmetry in a PhC gives rise to separate TE and TM modes, as detailed in section 2.1, the choice of even or odd mirror symmetry within the simulation only permits certain components of the electric and magnetic fields to be non-zero on the mirror plane.

2.3 Calculating the Band Structure of a 2D PhC

2.3.1 Overview

One of the most important numerical calculations which is of interest to us is the calculation of the band structure for a given PhC structure. All of the interesting properties of PhCs arise from the arrangement of the different bands and band-gaps, so this is often the first thing which is calculated for any new structure. In particular, when designing PhC waveguides and cavities it is essential to know the frequency range of the band-gap relative to the frequencies of the guided or confined modes.

We will begin by calculating the band structure for a simple hexagonal lattice of air holes in GaAs, as shown in Fig. 2.4. A hexagonal PhC structure was chosen for this since in the 2D case a hexagonal structure tends to give the greatest omnidirectional band gap [12]. The PhCs simulated in this chapter all have a lattice constant, a , of 340 nm and a hole radius, r , of $0.27a$.

In all these simulations the refractive index of GaAs was set to 3.33. The final devices were expected to operate at a temperature of 4.2 K and a vacuum wavelength of around 1.3 μm , but this value was extrapolated from data taken at $T = 2$ K and wavelengths between 820 nm and 1.1 μm [13] using the equations given in that paper. Given the small difference in temperature and the close fit between the provided model and the data, the uncertainty is almost certainly less than 0.02.

For the SU-8, used in later parts of this chapter, low temperature refractive index data was hard to come by. A value of 1.57 was used based on measurements at room temperature [14]. However, the error is likely to be quite large, not only because of the difference in temperature, but also because of differences in composition caused by different fabrication procedures (e.g. spin speed, soft bake temperature, exposure power).

The basic process of calculating a photonic band structure using FDTD algorithms involves running a separate simulation for each $\mathbf{k}$ vector of interest, then performing a temporal Fourier transform to extract the frequency of each band at each $\mathbf{k}$ vector. Typically this is done for a set of $\mathbf{k}$ points around the




Figure 2.4: (a) Diagram of the 2D PhC structure to be simulated. White is air or vacuum. Grey is GaAs. The green tinted hexagon is one possible choice of unit cell for this lattice.

(b) Diagram of the first Brillouin zone corresponding to the structure on the left. The irreducible Brillouin zone is shown in blue, and the high symmetry points are labelled.

edge of the irreducible Brillouin zone. As we will see, in addition to the usual considerations such as choosing the correct sources to excite the modes of interest and estimating numerical errors, there are a number of subtle difficulties unique to the problem of band structure calculations.

2.3.2 Selecting a $\mathbf{k}$ Vector

As explained in section 2.1, $\mathbf{k}$ vectors in a lattice can be understood as corresponding to particular phase differences between adjacent unit cells. The Brillouin zone contains the full range of possible phase differences between 0 and 2π in all directions, hence why it isn't meaningful to consider any $\mathbf{k}$ vectors outside of the Brillouin zone. This immediately suggests a possible way of restricting the simulation to a particular $\mathbf{k}$ vector.

For standard bandstructure calculations we wish to simulate an infinite PhC lattice, so periodic boundary conditions are the natural choice. The negative x boundary meets the positive x boundary, the negative y boundary meets the positive y boundary, and in the case of a 3D simulation the negative z boundary meets the positive z boundary. In this way the simulation can approximate an infinite lattice of repeated cells.

Enforcing the phase offset between adjacent unit cells is simply a matter of adding or subtracting the appropriate phase factor at the boundaries. This is only possible if the FDTD algorithm is designed to work with complex fields, which Meep fortunately is. When the field at one boundary is being updated, the field values it reads from the opposite boundary are multiplied by a phase factor $\exp(i\mathbf{k} \cdot \mathbf{r})$ where $\mathbf{r}$ is the vector from one boundary to the other.

When a $\mathbf{k}$ -point is specified, Meep applies these boundary conditions to all sides of the simulation volume, however it is still possible to add PMLs to the boundaries for which periodicity is not desired. For example, these are used at the top and bottom surfaces when running 3D simulations of PhC slabs to approximate free space. It is also important to note that mirror symmetry cannot be used along a particular axis if the component of the $\mathbf{k}$ vector along that axis is non-zero because the two sides of the simulation volume have different phases.

2.3.3 Initial Simulation

To begin with, the results of a naïve, minimal simulation will be shown in order to demonstrate one of the key problems with calculating the band structure for a hexagonal crystal structure. Because the FDTD software used here requires the simulation volume to be a rectangle or cuboid, we could not simulate just one single unit cell of the lattice. Instead we had to simulate a super-cell with at least twice the volume of the unit cell. The super-cell we chose is shown in Fig. 2.5. In order to minimise the simulation time and reduce the complexity of the band structure diagrams, 2D simulations were used in these initial tests, however we'll switch to a full 3D simulation at the end of the section when we're putting everything together.

The FDTD simulation was performed for 61 different $\mathbf{k}$ -vectors tracing a path around the perimeter of the irreducible Brillouin zone. The spatial resolution was set at 20 grid points per PhC lattice spacing. For each $\mathbf{k}$ vector, the fields were excited by 5 sources exciting the H_z field component, each with the same Gaussian frequency profile but a random phase, placed in random locations within the xy plane. The simulation was allowed to run for 1000 time steps after those sources




Figure 2.5: Diagram of the simulation volume. The dark and light regions correspond to GaAs and air respectively. The red rectangle is the simulation volume which corresponds to a super-cell of the lattice.

stopped emitting, and at each of those 1000 time steps the H_z fields were recorded at 5 randomly chosen locations for later analysis. Multiple source and monitor locations were chosen in order to ensure that all the modes of the system would be observed. After the simulation completed, software was used to perform harmonic inversion [15] in order to decompose the time-series field data into a list of exponentially decaying sinusoidal modes. The resulting band diagram is plotted in Fig. 2.6. The FDTD simulation does not provide any way to identify which particular band a given point belongs to, so the data points are not connected by lines.

The most obvious problem with this band structure is the band which approaches zero frequency at the M point. By considering the way in which bands form, it follows that the only place where modes should tend towards zero frequency is the Γ point. In total, approximately half the bands can be identified as non-physical when comparing these results to the literature [16]. To explain why these erroneous modes appear, it is necessary to take a more in-depth look at what is going on in reciprocal space when the Bloch boundary conditions are applied to the simulation volume.

2.3.4 The Origin of Folded Bands

The first thing to note is that the super-cell shown in Fig. 2.5 actually contains a total of two unit cells of the lattice. The Bloch boundary conditions restrict




Figure 2.6: Plot of the TE band structure of a PhC lattice calculated using a naive 2D FDTD simulation. A number of erroneous non-physical bands appear in the data.

the phase relationship between adjacent super-cells, but they do not place any restrictions on the phase relationship between the two unit cells within the simulation volume. In effect, this means that rather than solving for the band structure of a hexagonal lattice, we were solving for the band structure of a rectangular lattice with twice as many air holes per unit cell. The unit cell in real space has twice the area, so it follows that the first Brillouin zone in reciprocal space will have half the inverse area. This is shown in Fig. 2.7 (a).

It is a well known fact of semiconductor physics that any $\mathbf{k}$ vectors which lie outside the BZ get “folded” and are physically equivalent to $\mathbf{k}$ vectors inside the first BZ. When the simulation selects a particular $\mathbf{k}$ vector, what it is really doing is setting the phase relationship corresponding to the appropriate $\mathbf{k}$ vector inside the first BZ. $\mathbf{k}$ vectors outside the first BZ correspond to phase offsets outside the interval $[0, 2\pi)$, which are obviously equivalent to phase offsets within that range. Fig. 2.7 (b) shows the result of applying this folding to the $\mathbf{k}$ vectors used in Fig. 2.6.

This provides a very elegant explanation of why the M point in the band diagram looks like the Gamma point. The M point of the hexagonal lattice is equivalent to the Gamma point of the super-cell lattice. In fact, every band of




Figure 2.7: (a) Diagram of the first Brillouin zones of the hexagonal lattice (green) and the rectangular super-cell lattice (red).

(b) Diagram showing the band folding effect produced by the smaller BZ of the super-cell lattice. The grey dotted triangle shows where one choice of irreducible BZ of the hexagonal lattice would normally be. The blue area shows where those same $\mathbf{k}$ vectors get mapped to on the smaller BZ of the super-cell lattice. The Gamma and M points end up in the same place.

the hexagonal lattice appears twice in our calculated band structure: once as the real band and once as a folded copy.

There are a number of ways to resolve this problem which will be covered in the following sections. All of them fundamentally involve restricting the phase offset between the two unit cells within the simulation volume.

2.3.5 Approach 1: Phased Sources

One way of selecting for modes with the correct phase relationship is to preferentially excite them by embedding that phase relationship into the sources. Instead of exciting the modes with sources of completely arbitrary position and phase, it is possible to place pairs of sources in symmetrically equivalent locations in the two unit cells, then set the phase relationship between them according to the $\mathbf{k}$ -point being simulated. It is important that the phase relationships between different pairs of sources remain random to avoid preferentially selecting for any particular mode.

The same simulation was repeated with phase-locked sources, and the calculated band structure is plotted in Fig. 2.8. It can be seen that this approach is very effective at low frequencies, however at high frequencies there is much less difference in amplitude between the real bands and the folded bands. The




Figure 2.8: Plot of the TE band structure of a 2D PhC lattice calculated using FDTD with excitation provided by 5 pairs of phase-locked sources. The amplitude of the folded bands has been significantly reduced at low frequencies, however at higher frequencies the discrimination is much weaker. The blue shaded area indicates the light cone, although this is a 2D simulation so it does not introduce losses.

advantage of this method is that it requires very little extra computation. The extra harmonic inversion operations are negligible compared to the duration of the FDTD simulation, especially in 3D. However the results are still not perfect. The wide variation in amplitudes across both real and folded bands means that thresholding is not sufficient to eliminate the folded bands.

2.3.6 Approach 2: Exploiting Symmetry

Another way of enforcing the phase relationship is to take advantage of the fact that the BZ of the primitive cell lattice has a different symmetry to the BZ of the supercell lattice. The key observation is that the real band structure remains unchanged under a rotation by $\pi/3$ while the same is not true of the folded bands. Hence, if you calculate the band structure for two points in $\mathbf{k}$ -space which differ only by a rotation of $\pi/3$, the real bands will be the same in both, but the folded bands will be different. This technique is illustrated in Fig. 2.9.




Figure 2.9: Diagram of two choices of irreducible Brillouin zone for a hexagonal lattice. (a) shows the two equivalent irreducible BZs inside the BZ of the primitive cell lattice. (b) shows the folded versions of those same irreducible BZs when a supercell lattice is used. Note that the two choices of irreducible BZ are no longer equivalent after applying band folding.

The simulation was again repeated using 5 sources of random phase and 5 monitor points, but this time each simulation run was performed for two symmetrically equivalent points in the BZ. After converting the time-domain data to frequency-domain spectra, each mode was only accepted if it appeared in the spectra for both equivalent $\mathbf{k}$ vectors with a frequency difference no greater than $1.9 \times 10^{-2}c/a$. In theory the frequencies should be exactly equal, but there is always some small variation due to the finite resolution of the Yee lattice. The resulting band diagram is plotted in Fig. 2.10.

It's clear that most of the folded bands were removed by this procedure. At higher frequencies the performance is somewhat worse because there are a greater number of bands and hence more opportunities for coincidental overlap. However, the main limitation of this approach is visible at the Γ and K points where there are a large number of isolated data points. This can be explained by looking back at Fig. 2.9 and observing that the Γ and K points are the only two points which are always mapped to equivalent points by the band folding. For those particular points in $\mathbf{k}$ -space, there is no difference between the symmetry of the primitive cell BZ and that of the supercell BZ. This approach also has the disadvantage of requiring approximately twice the computation time, since each




Figure 2.10: Plot of the TE band structure of a 2D PhC lattice calculated using FDTD. Modes are only shown if they appear in simulations of two symmetrically equivalent choices of the irreducible Brillouin zone. Folded bands are suppressed due to the different symmetry of the BZ of the supercell lattice. The blue shaded region indicates the light cone, although this is a 2D simulation so it does not introduce losses.

point in $\mathbf{k}$ -space must be simulated twice.

2.3.7 Approach 3: Phased Monitors

The final way of restricting the phase within the simulation volume is to look at the phase difference between the modes detected by monitors at equivalent points in the two unit cells. This is very similar to the first approach, except here we fix the phase of the monitors instead of the sources. After the harmonic inversion has been performed on the time-series data from all the monitor points, modes are discarded if the phase relationship deviates from the predicted value by more than a certain threshold.

Once again, the simulation was run using 10 random sources and 10 monitor points, however this time the monitor points were grouped into symmetrical pairs. Each pair was used to extract a list of bands with the correct phase relationship, then the data from all 5 monitor points were combined, taking the maximum amplitude at each frequency and $\mathbf{k}$ vector. The resulting band diagram is plotted in Fig. 2.11.




Figure 2.11: Plot of the TE band structure of a 2D PhC lattice calculated using FDTD. Modes were filtered by comparing the phase at two symmetrically equivalent points within the unit cell.

This approach worked extremely well. The folded bands were neatly separated from the real bands by a phase error of almost exactly π , so they could be removed by a simple threshold operation. Only a very small number of erroneous data points slipped through this filter, and most of those were due to numerical errors where bands were almost parallel and overlapping. Additionally there was only a small amount of processing overhead due to having to perform twice as many harmonic inversion operations.

2.3.8 Final 3D Band Structure

From the previous results it can be seen that the technique of comparing the phase of each mode at symmetrically paired monitor points is by far the most effective way to eliminate erroneous folded bands from a simulation. Even considering each data point individually, in almost every case the phased monitor method performed equally well or better as a classifier than the other two approaches. Therefore, rather than combining multiple approaches, I decided simply to use the phased




Figure 2.12: Diagram of the 3D simulation volume used to calculate the band structure of an infinite PhC lattice via the FDTD method. The in-plane x and y faces have periodic Bragg boundary conditions, while the top and bottom faces have thick PML absorbers.

monitor method on its own when running the final 3D band structure calculation. The MEEP code used for this simulation is given in appendix B.

The 3D structure being simulated is shown in Fig. 2.12. It was necessary to add at least half a wavelength of free space above and below the PhC membrane to prevent the PML from interacting with the evanescent field in that region. Without that free space, the PML would absorb energy from the evanescent field and cause all modes to decay unnaturally quickly. Obviously an even larger gap would be slightly better, but the returns diminish exponentially and, as always, there is a trade off between simulation precision and simulation speed.

Both the TE and TM bands were calculated for the 3D case. The band structure produced by the 3D simulation are shown in Fig. 2.13.

2.4 Modes of a PhC Waveguide

In this section we will consider a waveguide created by removing a row of holes from the 2D PhC structure modelled in the previous section. This waveguide structure




Figure 2.13: Plot of the TE (purple) and TM (red) band structures of a 2D PhC embedded in a thin membrane surrounded by vacuum above and below. The band structure was calculated using a 3D FDTD simulation. Modes were filtered by comparing the phase at two symmetrically equivalent points within the unit cell. There is a band gap, indicated in green, for the TE bands but there is no corresponding gap in the TM bands.

is a building block for creating the cavities studied later in this chapter, so it is important that we understand the properties of the waveguide in isolation.

A diagram of the waveguide is shown in Fig. 2.14. The centre-to-centre spacing between the two rows of holes adjacent to the waveguide is actually $0.98\sqrt{3}a$ rather than the expected $\sqrt{3}a$. This is to match the physical devices fabricated by our collaborators.

2.4.1 Waveguide Band Structure

In order to find the guided mode(s) of the waveguide, it is necessary to calculate its band structure. The waveguide only has translational symmetry in one dimension, which makes the band structure calculation much simpler than it was for the 2D




Figure 2.14: Diagram of the waveguide structure studied in this section. The number of holes on each side of the waveguide and the length of the waveguide are chosen arbitrarily in this diagram.




Figure 2.15: Diagram of the waveguide structure with one possible choice of unit cell highlighted in blue.




Figure 2.16: Plot of the band structure of the PhC waveguide.




Figure 2.17: Plots showing the E_y component of the electric field in a 2D plane through the centre of the PhC waveguide for three modes with different frequencies but the same $\mathbf{k}$ vector. Red and blue correspond to positive and negative values of E_y respectively. The PhC structure is overlaid in grey and white, with white circles indicating air holes.

PhC. In particular the unit cell of the waveguide is a cuboid, as shown in Fig. 2.15, and hence it is not necessary to remove any folded bands from the simulation results.

As before, the E_y field was excited by 10 sources with random phase which were randomly located within a $1a \times 1a$ plane centred on the waveguide. The sources were spectrally broad enough to cover the entire photonic band gap plus a large margin either side. The simulation was run until the sources had finished emitting, and then for a further period to allow time-domain field data to be collected. The field data was collected from another 10 random points within the same region as the sources, then it was transferred to the Harminv package to extract the resonant modes. Even symmetry in the z direction was used to speed up the simulation, since the photonic band gap only existed for TE modes.

From the photonic band structure shown in Fig. 2.16 we can see three bright modes near the edge of the Brillouin zone. To understand the behaviour of these different modes it is useful to observe them in the temporal and spatial domains. To do this the previous simulation was repeated at a single $\mathbf{k}$ point, $k_x = 0.7 \pi/a$, and with a much narrower frequency source designed to excite only a single mode. Snapshots of the E_y field profile were recorded at time intervals




Figure 2.18: (a) is a plot of the spectrum of flux through a plane at one end of a waveguide (without any cavity structure) after exciting it with a broad pulse and allowing time for the short-lived modes to decay away. The images in (b) depict the E_y field component in a 2D plane through the centre of the PhC at regular intervals over a half period. Red and blue indicate positive and negative values of E_y , while the grey and white background is the PhC structure.

of $1/(20f)$. These frames were then arranged side-by-side such that the phases at joining edge were matched as closely as possible. The resulting images for the three modes are shown in Fig. 2.17

Fig. 2.17(a) is the most tightly confined mode with the smallest cross-sectional area. This is unsurprising since it is also the only mode with a frequency inside the photonic band gap of the surrounding PhC. The other two modes are significantly more lossy and less localised, with Fig. 2.17(b) exhibiting almost no confinement at all. Therefore the main guided mode of this waveguide is at a frequency of around $0.27c/a$.

2.4.2 Fabry-Pérot Modes

We have considered the modes of infinitely long PhC waveguides, but any real waveguide must occupy a finite volume. In this case we assume that the GaAs slab extends to the edge of the region of interest, but the lattice of holes within it terminates before the edge. As always, it's important to leave at least half a wavelength gap between the edges of any PhC structure and the start of the PML boundaries in order to avoid bleeding energy out of evanescent modes.

Fabry-Pérot modes are spectrally narrow resonant modes of the kind observed in a Fabry-Pérot interferometer. They occur due to constructive interference between a pair of reflecting elements. In the case of a finite PhC waveguide, the reflecting elements are the two ends of the waveguide. While the waveguide ends are not obstructed by any structures or interfaces, they still create reflections due to the mismatch in impedance between the PhC waveguide and the smooth GaAs membrane [17, 18]. This is similar to how acoustic reflections can occur at the ends of an open tube to produce resonances.

To model these Fabry-Pérot modes the E_y component of the electric field was excited by a short-lived broadband source in the centre of the waveguide. A large number of very short lived modes were excited, so the simulation was allowed to run for a time of $200 a/c$ to allow those to decay before beginning data collection. The simulation was then run for a further $600 a/c$ during which the fields were recorded in a plane perpendicular to the waveguide and offset a short distance from one end of it, outside the PhC structure. This time-domain data was then Fourier transformed to give a frequency-domain spectrum of the light escaping from the waveguide. Additionally the E_y field component in the plane of the PhC was recorded to allow the spatial shape of the modes to be observed.

The frequency spectrum is plotted in Fig. 2.18(a). The series of small undulations extending out from each side of the central peak correspond to Fabry-Pérot modes with different half-integer numbers of wavelengths spanning the length of the waveguide. If the transmissivity and reflectivity of the waveguide were independent of frequency then those modes would all have similar amplitude,




Figure 2.19: 3D render of the PhC cavity structure simulated in this chapter and fabricated in chapter 5. The PhC membrane is made of GaAs while the blue disk consists of SU-8 photoresist. Light is confined in the z direction by total internal reflection and in the x and y directions by the PhC band gap.

however the waveguide only transmits light within a narrow frequency band. This means that the guided mode acts as an envelope function to modulate the comb of Fabry-Pérot modes, giving rise to the large central peak in the spectrum.

Fig. 2.18(b) shows the E_y field profile in a plane through the centre of the PhC at a series of time steps corresponding to approximately half a period of the most significant frequency component. It is worth observing that the mode extends for most of the length of the waveguide but the intensity drops off gradually over a number of unit cells towards the ends. On shorter waveguides this can make it hard to distinguish between a localised cavity mode and Fabry-Pérot modes, but on sufficiently long waveguides the spatial profile is the most convincing way to distinguish between the two.

2.5 Simulating an SU-8 Cavity

2.5.1 Cavity Structure

The structure simulated in the following sections is the same as the one fabricated in chapter 5, and more details on the materials and fabrication are provided there. These simulations were performed to confirm the viability of this cavity design and determine some of its optical parameters.

Fig. 2.19 shows a 3D render of a cut-away section of the photonic crystal structure containing the cavity. Essentially the structure consists of a 2D PhC in a thin GaAs membrane with a waveguide formed by a missing row of holes in the Γ -K direction, and on top of the waveguide is a small disk of SU-8 photoresist. The disk of photoresist has a refractive index higher than the surrounding air or vacuum, and hence interacts with the evanescent fields of guided modes surrounding the waveguide. This interaction slightly changes the average refractive index which the guided mode ‘sees’. Changing the effective index within a region of the waveguide has a similar effect to changing the scale of the PhC and causes the frequency of the guided mode to shift. The photoresist is only applied over a small region of the waveguide, so the end result is a localised mode with a frequency which lies within the photonic band gap of the surrounding crystal and which also doesn’t match the frequency of the waveguide mode. It therefore acts as an optical cavity. This basic operating principle is very similar to the principle behind double heterostructure PhC cavities of which there are many examples in the literature [19, 20, 21].

2.5.2 Finding the Cavity Frequency

The first thing we want to find out about this cavity design is the frequency of the first order optical mode. We must do this first because in later stages we want to excite the cavity using a spectrally narrow source at exactly this frequency.

The basic procedure for calculating the cavity mode frequency involves exciting the cavity with a broadband source, then recording the spectrum of electromagnetic energy in the cavity as it decays away. The most long-lived frequencies correspond to the modes of the cavity. This is complicated slightly by the presence of Fabry-Pérot modes and the lossy nature of a finite PhC cavity, but at least in this case it was found that the frequency with the highest intensity corresponded to the first order cavity mode. As in the previous sections, the Harminv software package was used to decompose time-series data into a small number of exponentially decaying sinusoids with the greatest contribution to the total.




Figure 2.20: Diagram showing the PhC waveguide and SU-8 disk structure which was simulated to determine the frequency of the cavity mode.

When performing a large 3D simulation such as this, the execution time and memory requirements can very easily grow to impractical extremes for even moderately high simulation resolutions. For this reason it is very useful to be able to estimate the error associated with a particular resolution, and hence choose the correct balance between accuracy and computational efficiency. As the resolution is increased all of the physical properties of the system should converge asymptotically towards their ideal analytical values, and by observing this convergence it is possible to quantify how the error varies as a function of resolution.

The structure which was simulated is shown in Fig. 2.20. The lattice constant, a , of the PhC was defined to be 340 nm and the hole radius was $0.27a$. The PML was $1a$ thick and a gap of $2a$ was left between the edges of the PhC holes and the start of the PML region. The PhC slab was 200 nm thick, while the SU-8 disk was 100 nm thick and $1\text{ }\mu\text{m}$ in diameter. An E_y component source emitting with a Gaussian frequency spectrum over a range of free-space wavelengths from 1070 nm to 1570 nm was positioned at an arbitrary position close to the centre of the waveguide and beneath the SU-8 disk such that it would excite the fundamental mode of the cavity. Odd mirror symmetry about the xz plane was employed to speed




Figure 2.21: Plot of the frequency of the fundamental cavity mode as a function of simulation resolution. The solid blue line (left scale) is the frequency, while the dotted red line (right scale) is the absolute difference between adjacent points on the blue line. The corresponding vacuum wavelength is indicated for $a = 340$ nm.

up the computation and reduce memory usage. The simulation was run for a total time of $1200 a/c = 1.36$ ps. $200 a/c$ was required for the source to emit a pulse with the desired spectrum, and an additional $1000 a/c$ was allocated for measuring the resulting excited modes. After the source had turned off, the E_y field component was recorded at every time step at the same location as the source. The simulation was repeated for a range of different resolutions, keeping all other parameters the same.

The time-series data was analysed using Harminv which returned a list of frequency components and relative amplitudes. A good starting point is to look at the highest amplitude component at each resolution, and these frequencies are plotted in Fig. 2.21. The first thing to note is that, as expected, the frequency (solid blue line) appears to converge towards a particular value as the resolution is increased. The dotted red line shows the difference between each consecutive pair of points on the blue curve and serves as an estimate of the error. The red line is approximately straight when plotted on a log scale, suggesting that the error decreases exponentially. By fitting a straight line and extrapolating, we can estimate




Figure 2.22: Diagrams of the two PhC structures which were compared to determine the influence of size on the cavity wavelength. The corresponding vacuum wavelengths of the cavity modes are shown underneath. An identical SU-8 disk was present in each simulation, although they aren't shown in these diagrams.

the error associated with the highest resolution. If the blue line converged towards its final value from only one side then the error would likely be underestimated using this approach, but the blue line appears to oscillate around its final value so this approach can be expected to slightly overestimate the error.

From this we can arrive at a frequency of $(0.267145 \pm 0.000008) c/a$ or equivalently a vacuum wavelength of $(1272.72 \pm 0.05) \text{ nm}$. It was decided that the six significant figures of precision attainable at a resolution of $30 a^{-1}$ was somewhat unnecessary, so future simulations would be run at a resolution of only $20 a^{-1}$. This would reduce the execution time by a factor of 5.1 and reduce the memory requirements by a factor of 3.4 compared to the higher resolution.

Another factor to consider is the effect of the PhC size on the results. The PhC structures in these simulations are not infinite, so it is reasonable to expect that there would exist some size below which the model of an infinite PhC would break down and the frequency of the cavity mode would begin to deviate from its ideal value. In order to save computing resources, the PhC simulated in this section was smaller than the ones fabricated in chapter 5, so it is important to check that the results are still valid. To do this, the same simulation was run for an even smaller PhC at a resolution of $20 a^{-1}$ and compared to the previous results.




Figure 2.23: Example of subpixel averaging used to smooth out the edges of a circle.




Figure 2.24: Plot of the frequency of the cavity mode as a function of simulation resolution for the cases when subpixel averaging is enabled and disabled.

As can be seen from Fig. 2.22, the number of holes in the PhC structure has a negligible influence on the cavity wavelength, at least for the size of PhC studied here. This also strongly suggests that the mode we're observing is a cavity mode rather than a Fabry-Pérot waveguide mode, since if it were caused by reflections from the ends of the waveguide then changing the length of the waveguide would have more of an effect on it.

2.5.3 Subpixel Averaging, Stability and Numerical Accuracy

In some cases when running very long simulations, such as those used in the next section to calculate the Q factor, the magnitudes of the $\mathbf{E}$ and $\mathbf{H}$ fields were observed to diverge towards positive and negative infinity. Further investigation of this numerical instability revealed that it was exacerbated by two principal factors: the use of mirror symmetry to simplify the simulation, and subpixel averaging of the values of ϵ and μ .

As described in section 2.2.2, the FDTD simulation is run on a discrete grid of points which can be likened to pixels in a raster image. Each can only have a single value of ϵ and μ . A consequence of this is that it is generally impossible to accurately reproduce smooth functions such as the cylindrical holes in the PhC membrane. Subpixel averaging can somewhat mitigate this limitation by linearly interpolating between the properties of the two materials for any pixels which overlap the boundary between them. A graphical depiction is shown in Fig. 2.23. The implementation of subpixel averaging in MEEP is reported to give first order accuracy for fields directly on the boundary and second order accuracy elsewhere [8].

By default subpixel averaging is enabled in MEEP, and all of the simulations up to this point have taken advantage of it. Disabling it solves the problem of diverging fields but causes a reduction in accuracy. To quantify how much accuracy was lost by disabling subpixel averaging, the previous frequency convergence run was repeated without subpixel averaging. A comparison between the two sets of simulations is plotted in Fig. 2.24. The difference in the convergence is striking.

The same approach was used as in the previous section to estimate the error as a function of resolution. In both cases the error decreases exponentially with increasing resolution, but when subpixel averaging is disabled that convergence occurs more than three times more slowly. At a resolution of $30a^{-1}$ the error in frequency was still $2 \times 10^{-3} c/a$, two orders of magnitude greater than the equivalent when subpixel averaging was enabled.




Figure 2.25: Diagram showing a simplified model of how the Q factor was measured in the FDTD simulations. The sum of the Poynting vector, $\mathbf{S}$, through each face of the enclosing cuboid was recorded (only 3 sides are shown), as well as the total energy contained inside the cuboid.

There were two possible solutions to the problem of the simulation diverging: either disable subpixel averaging at the cost of accuracy, or reduce the symmetry at the cost of execution speed. The loss of accuracy from disabling subpixel averaging could be compensated for by increasing the resolution, but increasing the resolution by a factor of x increases the execution time by a factor of x^4 . In contrast, reducing the symmetry would only increase the execution time by a factor of 2 for each mirror plane which was removed. From the data plotted in Fig. 2.24 it is clear that to achieve the same accuracy with subpixel averaging disabled as can be achieved with it enabled would require increasing the resolution by much more than a factor of two. Therefore, given a fixed duration for the simulation, the best accuracy would be attained by reducing the symmetry rather than by disabling subpixel averaging.

Subpixel averaging was left enabled for the remaining simulations and the symmetry settings were adjusted as necessary to prevent instability.

2.5.4 Determining the Q Factor and Mode Volume

The Q factor, or quality factor, of a resonator is a measure of both how effectively it traps energy and how narrow its spectral linewidth is. These two definitions are not equivalent for very low Q cavities, but converge to the same value as the quality factor increases. The two definitions are as follows [22, 23],

$$Q = \frac{f}{\Delta f} \quad (2.7)$$

$$Q = 2\pi f \times \frac{\text{energy stored}}{\text{power loss}} \quad (2.8)$$

where f is the resonant frequency and Δf is the linewidth.

There are three basic methods for measuring the Q factor in an FDTD simulation. The first involves using the definition in (2.7) to calculate Q from the spectrum of the cavity mode. The field amplitude inside the cavity is recorded at each time step, then that time series is Fourier transformed to retrieve the frequency spectrum. The disadvantage of this method is that in order to resolve the linewidth of the resonant mode accurately, a large number of data points must be recorded in the time domain. For high Q cavities the linewidth can be extremely small, meaning that the duration and computational cost of the simulation becomes prohibitive, so this method is best suited to low Q cavities.

The second method is based on the definition in (2.8) and operates in the time domain rather than the frequency domain. An ideal decaying mode is described by an equation of the form $E = A \exp(-(t - t_0)/\tau) \exp(i\omega t)$, where A is some initial amplitude at time t_0 , and τ is the lifetime of the mode. It is possible to extract the $A \exp(-(t - t_0)/\tau)$ component from the time series data recorded by the simulation, essentially giving the envelope function containing the decaying sinusoidal mode, and from the decay constant, τ , it is possible to calculate the Q factor [24]. This method works much better than the frequency domain approach for high Q cavities since it is not necessary to record as many data points in the time domain [25].

The final method again takes advantage of (2.8) to calculate the Q factor using time domain data, however in this case the quantities of energy stored and

lost by the cavity are directly measured during the simulation. The energy stored in the cavity can be calculated by evaluating the integral,

$$\text{energy stored} = \frac{1}{2} \int \epsilon \mathbf{E}^* \cdot \mathbf{E} + \frac{\mathbf{H}^* \cdot \mathbf{H}}{\mu} dV, \quad (2.9)$$

over the entire cavity volume, including regions containing evanescent fields. The rate of energy loss can be calculated by integrating the Poynting vector, $\mathbf{S} = \mathbf{E} \times \mathbf{H}$, in a thin shell completely enclosing the cavity [26]. Given these two values, it is a simple matter to calculate the Q factor using (2.8). This method has the potential to be by far the fastest, and it also has the advantage of being able to decompose the Q factor into directional components based on the spatial distribution of the energy loss. For instance, by recording the integral of the Poynting vector through each wall of the cuboid computation cell separately the in-plane (xy) losses can be measured independently of the out-of-plane (z) losses, potentially providing information useful for improving the cavity design. For these reasons, I chose to use this method for studying the SU-8 cavity structures. The MEEP code used for this simulation is given in appendix B.

A diagram of this approach for measuring the Q factor is shown in Fig. 2.25. A very spectrally narrow E_y source was placed in the centre of the waveguide directly beneath the SU-8 disk, set to emit at the previously determined frequency of the cavity mode. The source had a frequency-domain FWHM of only $0.001 c/a$ to ensure that it would excite only the mode of interest. To produce such a narrow linewidth pulse it was necessary to run the simulation for a time of $10000 a/c$ with extremely gradual fade-in and fade-out of the source. After the source had finally stopped emitting, the simulation was run for a further time of $37.4 a/c$ (equal to 10 periods of the mode frequency) during which data collection occurred. The integral of the Poynting vector was recorded through planes positioned just inside the PML layers to avoid any undue attenuation, and the total energy was integrated within the volume bounded by those planes. The structure and fields possessed both x and y mirror symmetry (although exploitation of those mirror symmetries was not enabled in MEEP's settings due to the aforementioned instability), so the integral




Figure 2.26: Plot of the total energy contained in the electric and magnetic fields within a cuboid of dimensions $10a \times 6a \times 4.6a$ centred on the cavity as a function of simulation time.

was only performed over one x boundary and one y boundary then the values were doubled. Various field components were also recorded as a function of position at regular intervals during this data collection phase.

The total field energy within a small volume centred on the cavity is plotted as a function of time in Fig. 2.26. The initial Gaussian-like curve upwards is due to the source gradually turning on and adding energy to the cavity, while the much more gradual decay at the end shows the rate of energy loss from the cavity.

The total Q factor could then be calculated from the in-plane and out-of-plane components according to the formula [27],

$$\frac{1}{Q_{\text{total}}} = \frac{1}{Q_{\text{in}}} + \frac{1}{Q_{\text{out}}}$$

The E_y field of the energised mode is shown in Fig. 2.27. First of all, this provides absolute confirmation that we are looking at a localised cavity mode rather than a Fabry-Pérot waveguide mode since the field does not extend to the ends of the waveguide. The E_y field remains localised around the SU-8 disk




Figure 2.27: Plot of the E_y field component in a 2D plane through the centre of the PhC cavity at a single point in time after the cavity mode was excited by a narrow pulse. Red and blue indicate positive and negative values of E_y . The grey and white areas correspond to GaAs and vacuum respectively. Although it is not shown, the SU-8 disk is located directly above the midpoint of the waveguide.




Figure 2.28: Plot of the E_z field component in a 2D plane through the centre of the PhC cavity at a single point in time after the cavity mode was excited by a narrow pulse. Red and blue indicate positive and negative values of E_z . The grey and white areas correspond to GaAs and vacuum respectively. Although it is not shown, the SU-8 disk is located directly above the midpoint of the waveguide.

throughout the full period of the mode's oscillation, indicating that it is the SU-8 which is responsible for the confinement.

The Q factor was found to be 7.2×10^3 , with in-plane and out-of-plane components given by $Q_{\text{in}} = 8.0 \times 10^3$ and $Q_{\text{out}} = 7.0 \times 10^4$ respectively. It is particularly interesting that the value of Q_{in} is so much smaller than Q_{out} , since for any reasonably large PhC lattice (which this one is) the in-plane quality factor is almost never the limiting factor and typically exceeds at least 10^5 . Most practical

PhC cavities have their Q factors limited by the out-of-plane losses caused by overlap between the $\mathbf{k}$ -space mode distribution and the light cone [24].

The reason that Q_{in} is so low in this case can be understood by considering the separability of TE and TM solutions which was explained in the introduction to this chapter. So far we have only been analysing the TE modes of this cavity, and in any normal 2D PhC that wouldn't cause any problems since TE and TM modes are completely independent. However, in this case the z mirror symmetry is broken by the presence of the SU-8 disk and hence solutions cannot be neatly separated into TE and TM components. The modes in the vicinity of the SU-8 disk are instead free to be comprised of any combination of $\mathbf{E}$ and $\mathbf{H}$ field components. This is visualised in Fig. 2.28 which shows the spatial distribution of the E_z field component within the $z = 0$ plane through the centre of the PhC. If the TE and TM modes remained separable, we would not expect to see any excitation of the E_z component by the E_y source. The end result of this violation of z mirror symmetry is that the cavity mode can couple to both TE and TM modes radiating into the surrounding PhC. The TE modes are forbidden by the photonic band gap, however, as seen in Fig. 2.13, there is no similar band gap for TM modes so those are free to propagate. In terms of the band structure, the coupling between the cavity mode occurs at the points where the lowest TM band crosses the frequency of the cavity mode.

There are a number of different definitions for mode volume depending on the phenomenon being studied. In this case we are mainly interested in the Purcell effect arising from the cavity, so the relevant definition is as follows [28],

$$V_{\text{eff}} = \frac{\int \epsilon \mathbf{E}^2 dV}{(\epsilon \mathbf{E}^2)_{\text{max}}} \quad (2.10)$$

Equation (2.10) was averaged over approximately 100 time points to arrive at the final mode volume of $2.04 a^3$. Such mode volumes are normally expressed in units of $(\lambda_0/n)^3$ where $\lambda_0 = c/f$ is the vacuum wavelength and n is the refractive index of the primary material (in this case GaAs), resulting in a value of $V_{\text{eff}} = 1.44 (\lambda_0/n)^3$. In other words, the mode volume is barely larger than a single cubic

wavelength. This result is very promising for achieving Purcell enhancement or any strong coupling effects.

2.6 Conclusions

These results show that the design of PhC cavity studied here is at least theoretically sound. The Q/V ratio was calculated to be $5.0 \times 10^3 (\lambda_0/n)^{-3}$ which could, under ideal conditions, yield a Purcell factor of around 380. Assuming dot-cavity coupling can be achieved with reasonable reliability, this design is very promising for any applications where bright single photon sources are required.

3

Experimental Techniques

Contents

3.1	Introduction	45
3.2	Excitation Sources	46
3.2.1	Continuous Wave Lasers	46
3.2.2	Pulsed Lasers	47
3.3	Time-Integrated Micro-Photoluminescence	48
3.4	Micro-Photoluminescence Mapping	53
3.4.1	Standard Mapping	54
3.4.2	Excitation Mapping	57
3.5	Time-Resolved Micro-Photoluminescence	59
3.6	Microscope Imaging	62
3.7	White Light Interferometry	62
3.8	Automation Techniques	65
3.8.1	Automated Mapping and Pattern Writing	65
3.8.2	Automatic Position Optimization	66
3.8.3	Computer Controlled Shutters	67

3.1 Introduction

This chapter introduces the various techniques I employed to investigate the properties of both GaAs PhC structures containing QDs, and hybrid lead-halide perovskite materials. The overwhelming majority of measurements were optical in nature, generally involving using a laser to excite the sample and then passing

the photoluminescence emission into one of several different instruments. The different laser sources are listed and described in section 3.2. Section 3.3 covers the most common type of measurement, time-integrated micro-photoluminescence spectroscopy, which involves passing the emitted light through a spectrometer and recording the spectrum with a CCD camera. The other principal measurement mode, time-resolved micro-photoluminescence, makes use of a photomultiplying detector and time-correlated photon counting card to record the evolution of optical processes with sub-nanosecond resolution. That is detailed in section 3.5. Two different mechanisms were used to perform micro-photoluminescence maps over regular grids of points in a 2D plane, and those are described in section 3.4. Section 3.6 covers the apparatus used for capturing microscope images of the samples, while section 3.7 explains the white-light interferometry technique used to measure the thickness of thin films of photoresist. Finally, 3.8 describes a number of ways in which the data collection process was made more efficient through automation.

3.2 Excitation Sources

3.2.1 Continuous Wave Lasers

All of the micro-photoluminescence studies were performed using above-bandgap excitation. For the GaAs samples the band gap was greater than 1.5 eV so a 633 nm helium-neon (HeNe) gas laser with a maximum output power of approximately 1.7 mW was used.

The band gap of the perovskite samples was significantly larger, so for those it was necessary to use a neodymium-doped yttrium aluminium garnet (Nd:YAG) laser emitting at a wavelength of 532 nm. Internally this laser consists of an 808 nm AlGaAs diode which pumps the Nd:YAG crystal to produce emission at 1064 nm [29]. This 1064 nm light is then focussed through a crystal of KTiOPO_4 which doubles the frequency of the light through a nonlinear optical process [30]. Depending on the quality of the laser some of the pump beam can leak into the output beam, so it can be necessary to add a filter to the optical path for certain applications. In fact, two

different Nd:YAG lasers were used for this work: one which exhibited the 808 nm emission and one which did not. Both lasers emitted at a power of around 60 mW.

A continuous wave laser was also used for single-photon exposure of the SU-8 photoresist, as detailed in chapter 5. This was a blue 405 nm diode laser with an output power of 10 mW.

3.2.2 Pulsed Lasers

Significant parts of the work detailed in this thesis required the use of very fast sources with high peak power, both for nonlinear excitation and time dependent measurements. Two different pulsed lasers were used: a Ti:Sapphire laser for two-photon exposure of SU-8 photoresist, and a 450 nm diode laser for time-resolved micro-photoluminescence measurements of perovskite samples.

Ti:Sapphire Laser

The Ti:Sapphire laser gets its name from the Al_2O_3 crystal doped with Ti^{3+} ions which forms the gain medium [31]. This is pumped by a 532 nm diode laser with a Nd:YVO₄ gain medium producing 1064 nm light which is then frequency doubled.

Modelocking is achieved by the optical Kerr effect which is a non-linear optical process in which the refractive index of the medium increases with increasing optical power density [32]. This effect, exhibited by the Ti:Al₂O₃ crystal, causes sufficiently high intensity light to self-focus. Since the optical power density in the pulsed mode of operation is orders of magnitude greater than in continuous wave (CW) mode, the Kerr effect causes the two modes to have different beam diameters and divergences. Thus it is possible to design cavities which preferentially select for pulsed operation over CW operation [33].

The Ti:Sapphire laser used in this work was a Spectra-Physics Mai Tai emitting at a wavelength of 800 nm with a power of 0.75 W, a repetition frequency of 80 MHz, and a pulse duration of 100 fs. This gives a pulse energy of 9.4 nJ and a peak power of at least 94 kW.

450 nm Diode Laser

For time-resolved measurements of the perovskite measurements it was necessary to use a shorter wavelength laser to excite above band gap. A PicoQuant LDH-P-C-450B InGaN laser diode head was used as the laser source, with power supplied by a PicoQuant PDL 808 “Sepia” controller. Rather than pulses being generated by physical mode-locking within the laser cavity as was the case with the Ti:Sapphire, for this laser the pulses were generated by the controller sending pulses of power to the laser diode [34]. This allowed for much greater control over the timing of the laser pulses. The pulse rate could be adjusted on the controller in the range from 5 MHz to 80 MHz, with a pulse duration of less than 70 ps. Additionally the controller provided a secondary output from the crystal oscillator which could be connected to the “sync” input of the time-correlated single photon counting card, as detailed in Section 3.5. The maximum average power for this laser was around 8.3 mW, although reducing the repetition rate resulted in a proportional reduction in the average power since the energy per pulse was constant.

3.3 Time-Integrated Micro-Photoluminescence

The basic principle of micro-photoluminescence (μ PL) measurements is as follows. The sample is illuminated in a localised region by a focussed beam of light, electronic states are excited by absorption of this light, then those states decay and re-emit light which is collected for analysis.

In this work, excitation was provided by various different laser sources exciting above the band gap of the material being studied. The laser light was focussed onto the sample surface using an objective lens and the emitted light was collected back along the same path using the same lens. For all the μ PL studies we were exciting above the band gap, so the emitted light was always at a longer wavelength than the laser. This meant we could always separate the collected light from the laser light using a dichroic filter at an angle of 45° to reflect the laser but transmit the emission.

Most of the measurements were made using one of two different objective lenses, both manufactured by Mitutoyo. One was optimised for visible wavelengths while the other was optimised for near infra-red (NIR). Both objectives had a magnification of $100\times$ and a numerical aperture (NA) of 0.5.

In most cases when performing μ PL measurements it is desirable to restrict each measurement to as small an area (or volume) as possible on the sample. The way to do this is to use the technique of confocal microscopy. By spatially filtering the input beam using a single-mode optical fibre or pinhole it is possible to achieve a focal spot on the order of $1\text{ }\mu\text{m}$ in diameter. The same principle can be applied to the collection path to restrict the volume from which light is collected, however spatially filtering the collected light inevitably results in a significant loss of intensity - usually at least 50%. Depending on the exact nature of the sample being studied it is not always necessary to use strong spatial filtering on the output, and in some cases the increase in collection efficiency is more valuable than the loss in spatial resolution.

Before reaching the spectrometer, the collected light was directed through a long-pass filter to remove any stray laser light which had been reflected by the sample. This is very important because without the filter the high intensity laser light can scatter around inside the spectrometer and create false signals on the detector.

For non-confocal collection, light was input into the spectrometer by focusing it onto the input slit using a $5\times$ magnification lens. In fact, the slit still provided confocality in one dimension by restricting the width of the focused light which could enter the spectrometer. This was necessary in order to achieve high spectral resolution since the spectrometer imaged the input slit directly onto the CCD. If the input slit was set too wide, any given spectral line would be correspondingly spread out on the detector leading to a loss of resolution. In practice, while this configuration was technically only confocal in one dimension, it could still achieve reasonably high spatial resolution on some samples thanks to the small size of the excitation laser spot. If the carrier diffusion length was small, as it was in the GaAs samples, then the photoluminescence was tightly localised in the vicinity of the excitation spot. However, when doing measurements involving the scanning mirror

system or studying samples with very high carrier mobility, such as the perovskites, it was necessary to switch to a fully confocal collection configuration.

For fully confocal μ PL measurements the collected light was first coupled into an optical fibre with a $25\text{ }\mu\text{m}$ core diameter, then the other end of that fibre was positioned right up against the slit so that the core itself would act as the point source to be imaged onto the CCD detector. The optics used to couple the collimated light into the fibre were chosen such that a diffraction limited spot on the sample surface was imaged to fill the $25\text{ }\mu\text{m}$ core of the fibre.

The spectrometer itself was an Andor Shamrock 0.3 m imaging spectrograph in a Czerny-Turner configuration. It was equipped with three different diffraction gratings covering the visible and near-IR spectrum. In order to understand the properties of the diffraction gratings and hence the capabilities of the spectroscopy system, it is first necessary to cover some theory surrounding diffraction gratings.

For light incident at an angle of α with respect to the normal, a grating maximally reflects light of wavelength λ at an angle of θ_m according to the generalised grating equation [35],

$$d(\sin \alpha + \sin \theta_m) = m\lambda, \quad (3.1)$$

in which d is the distance between adjacent lines on the surface of the grating and m is an integer indicating the order of the reflection. For any practical system the diffracted beams have a finite angular width which is determined by the size of the illuminated region on the grating. In essence the finite size of the illuminated region acts as a finite aperture, restricting the Fourier components which can contribute to the outgoing beam and hence forcing it to have a non-zero divergence. From a quantum mechanics point of view the finite aperture extracts some information about the position of a given photon, therefore there must be some uncertainty in its velocity.

It can be seen from Eq. (3.1) that the zeroth order ($m = 0$) dispersed beam contains all wavelengths of light and is reflected at the same angle as one would expect from a plane mirror. The first order ($m = 1$) diffraction peak is the one

used by the spectrometer. When taking measurements of different wavelengths, the grating is rotated to focus the desired range of wavelengths onto the detector. Each grating is also blazed, meaning that rather than being a smooth plane interspersed with regularly spaced grooves, the surface has a saw-tooth structure of angled reflecting elements. The purpose of this is to concentrate as much of the incident light as possible into the first order reflection. Because the angle of the first order reflection is dependent on wavelength, a blazed grating only works optimally at a particular wavelength. Away from that wavelength the light collection efficiency smoothly degrades, although in practice the performance is still acceptable within a wavelength range of around $\pm 30\%$.

One subtle side effect of Eq. (3.1) is that when multiple wavelengths of light are present, one diffraction order of one wavelength can overlap with a different order of a different wavelength. For example, the second order diffraction peak from 600 nm light will occur at the same angle as the first order diffraction peak from 1200 nm light and hence produce a false reading on the detector. This was another reason to place a long-pass filter in the optical path before the spectrometer.

Two different types of detectors were used for recording the spectra. For the GaAs samples emitting in the near-IR, an Andor iDus InGaAs CCD array was used. It has a 1D array of 1024 pixels, each $25\text{ }\mu\text{m} \times 500\text{ }\mu\text{m}$ in size. For the perovskite samples which emit in the visible spectrum, an Andor NewtonEM silicon CCD camera was used which has a 2D array of 1600×200 pixels, each $16\text{ }\mu\text{m} \times 16\text{ }\mu\text{m}$. Both detectors were cooled down to a temperature of around -60°C using integrated thermoelectric coolers in order to reduce the dark count.

With the particular spectrometer configuration used in this research, the main factor limiting the spectral resolution was the CCD pixel size. Equivalently, the resolution could be said to be limited by the length of the optical path inside the spectrometer since that is what determines the ratio between angular dispersion from the grating and spatial dispersion on the CCD. The longer the distance between the grating and the CCD (and the longer the focal length of the internal focusing mirror), the more the different wavelengths spread out in space and hence the

line density (lines/mm)	blaze wavelength (nm)
600	1200
1200	1100
1200	500

Table 3.1: List of spectrometer gratings used in these studies.

further apart they are on the detector. In this limit making the spectrometer longer is equivalent to making the CCD pixels smaller.

The resolution of an array of sensor pixels is generally defined as the minimum separation, $\Delta\lambda$, between the centres of two narrow peaks while still being able to distinguish them. Resolving the separation between two peaks requires there to be a local minimum between them, which in the worst case requires a centre-to-centre separation of at least three pixel widths.

The parameters of the different gratings used in this research are given in Table 3.1. The top two gratings were blazed for wavelengths in the near IR and were mainly used when studying GaAs QDs and PhC cavities. They were paired with the InGaAs CCD detector, giving resolution limits of 0.33 nm and 0.084 nm for the 600 lines/mm and 1200 lines/mm gratings respectively when measuring emission at a wavelength of 1300 nm. The last grating was blazed for the visible spectrum and was used with the silicon CCD detector for studying the perovskite samples. Its effective resolution at a wavelength of 600 nm was 0.055 nm.

A diagram of the whole system is shown in Fig. 3.1, and a photograph is shown in Fig. 3.2. One thing worth mentioning from the photograph is the use of the ThorLabs cage system for mounting optical components. The basic cage system consists of four parallel metal rods to which other components can be affixed. Junction cubes and 45° mirror units allow different cage sections to be connected together, but all connected sections are constrained to angles of 90° relative to one another. The main advantage of having all optical paths running along three principal axes (east-west, north-south, up-down) is that it makes the whole system a lot simpler to design and maintain and makes it harder to accidentally introduce distortion to the beam through misaligned components.




Figure 3.1: A diagram of the entire micro-photoluminescence spectroscopy system. Multiple different lasers with different wavelengths were used as excitation sources, but those have been omitted for simplicity.

It also makes the system safer, since all laser beams are either parallel to the optical bench (and hence remain below eye level) or, in the case of vertical sections, blocked by a fixed mirror at the top.

3.4 Micro-Photoluminescence Mapping

Recording the PL spectrum of a single point on a sample is useful, however it is often desirable to be able to record an image of the PL spectrum as a function of position. It is also very useful to be able to move the excitation and collection spots around on the sample with sub- μm precision, for example to perform a μPL measurement at the centre of a PhC cavity. Coarse movements were made possible by mounting the cryostat on an XY stage with manual micrometers for positioning, but for fine positioning and computer control it was necessary to use piezoelectric actuators to adjust the optical path. Two different systems were implemented for different kinds of mapping.




Figure 3.2: A photograph of the μ PL system with some important parts labelled. The game controller on the optical table is used to manually control the XYZ piezo stage, scanning mirror, and shutter.

3.4.1 Standard Mapping

The simplest kind of positioning operation is to move the collection and excitation spots together in a way that is equivalent to moving the sample beneath a fixed μ PL system. The cryostat itself was too heavy to mount on a piezo-actuated stage, so instead the objective lens was attached to an XYZ piezo stage. The maximum displacement from the central axis was only $70\text{ }\mu\text{m}$ so the paraxial approximation could be applied. Since the angle of the collimated beam did not change, an ideal lens would always focus the light to the same position relative to itself. As the objective lens was moved, the focal spot moved with it. The same reasoning could be applied to the collected light, hence moving the objective over small distances was equivalent to moving the sample.




Figure 3.3: An SEM image showing one edge of the nanopillar sample which was used for testing the accuracy of the mapping system. Each pillar consists of a GaN core surrounded by 6 concentric shells of InGaN quantum wells. The pillars are arranged in a regular hexagonal grid with a pillar-to-pillar spacing of approximately $2\text{ }\mu\text{m}$.

Two different piezo stages were used over the course of this research. A PI P-611.3S NanoCube XYZ Piezo Stage with a range of $100\text{ }\mu\text{m}$ in all three directions was used for the SU-8 measurements, while a PI P-733.3CD XYZ Piezo Nanopositioning Stage with a range of $100\text{ }\mu\text{m}$ in XY and $10\text{ }\mu\text{m}$ in Z was used for the perovskite measurements.

In order to test the accuracy of the P-733.3CD piezo stage mapping system, a sample with a regular pattern of emitters was used. This sample consisted of a hexagonal lattice of nanopillars, shown in Fig. 3.3, each composed of a GaN core and 6 InGaN quantum well layers arranged in concentric shells around it. The nanopillars exhibited bright photoluminescent emission over a range from 420 nm to 560 nm while the areas of substrate between them were dark at those wavelengths.

Excitation was provided by the 450 nm pulsed diode laser at a power of 0.88 mW on the sample surface. The dichroic mirror was replaced with a 50:50 beam splitter so that the emitted light wouldn't be blocked from reaching the spectrometer. This allowed the reflected laser light through, so to block that a 500 nm to 580 nm band-pass filter was used. This allowed through the long wavelength tail of the quantum well emission which, when integrated over the




Figure 3.4: (a) Room temperature μ PL map of a $100\ \mu\text{m} \times 100\ \mu\text{m}$ area of the nanopillar sample from Fig. 3.3. The counts are integrated over the wavelength range 500 nm to 600 nm. The bright spots correspond to locations of nanopillars. (b) Plot showing the location of each auto-detected bright spot and the corresponding displacement from an ideal regular grid. The horizontal bands of greater displacement correspond to visible distortion in the μ PL map.

entire CCD, gave sufficient signal. The exposure time was kept short at 0.1 s to reduce the time taken to perform the map.

The μ PL map, shown in Fig. 3.4 (a), was taken over a $100\ \mu\text{m} \times 100\ \mu\text{m}$ area with a step size of $0.5\ \mu\text{m}$ on both axes. The regular grid structure on the sample is clearly visible, but careful visual inspection reveals a small amount of distortion. To quantify this distortion and understand its origin a more quantitative analysis algorithm was developed.

A computer program was written to automatically detect the locations of the pillars by looking for local maxima within the integrated map data. Thanks to the high contrast and high signal-to-noise ratio of the data, this simple approach proved to be very reliable and produced remarkably few false positives. For every local maxima found, its location was estimated with sub-pixel precision by taking a weighted average of the surrounding pixels.

The lattice spacing was estimated by simply counting the number of detected pillars and comparing it to the area of the map. The list of XY coordinates was also used to estimate the angle of the lattice relative to the mapping axes by averaging the angles from each pillar to its nearest neighbour, modulo $\pi/3$.

These lattice parameters were then further refined using a differential evolution optimization algorithm [36] to find the closest correspondence between points on an ideal hexagonal lattice and the observed positions of the pillars. It was then possible to iterate through the list of observed pillar locations and calculate the distance by which each one was offset from its corresponding ideal location.

These offset distances are plotted in Fig. 3.4 (b). The XY coordinates of each point correspond to its detected location, while the colour indicates the offset distance. Pillars very close to the edge are not included because their locations could not be accurately estimated.

All of the offset distances were small with an average of just $0.25\text{ }\mu\text{m}$, indicating that the positioning system had high precision. In fact this value represents an upper bound on the average positioning error of the system, since a certain amount of error would have been introduced during the automated pillar detection and fitting process. The other thing to note is that rather than increasing radially towards the edges of the map as one would expect from optical distortion in the lens, the offset distances appear to vary sinusoidally in horizontal bands. This suggests that the most likely cause is small temperature fluctuations in the room which produce thermal expansion and contraction in the metal framework supporting the sample and optical components. Each data point in the map is taken at a different time, so time-periodic fluctuations in the signal tend to manifest as spatially periodic anomalies in the data.

3.4.2 Excitation Mapping

The other sub- μm positioning operation which the system was designed to perform was moving the excitation spot while holding the collection spot in a fixed location. To change the location of the focal spot of the objective lens it was necessary to change the angle at which the collimated beam entered the lens. To do this a Conex AG-M100D piezo-actuated mirror was used to deflect the incident beam, then a 4f system was used to invert the beam path so that the position where the beam struck the objective remained constant while the angle changed.




Figure 3.5: A diagram of the 4f system used to precisely adjust the position of the excitation laser spot on the sample. f is the focal length of the two lenses. Whatever angle the beam is reflected off the scanning mirror at, the collimated beam always enters the centre of the objective at the same angle.

As shown in Fig. 3.5, the 4f system consisted of a pair of lenses of equal focal length, f , separated by a distance of $2f$. The first lens (left) acted as an image-space telecentric lens, collecting collimated light reflected from the piezo mirror and focusing it to a point in the plane at the halfway point in the system. The position of the focal spot in the middle was dependent upon the angle, θ , at which light was reflected from the mirror. The second lens (right) acted as an object-space telecentric lens, collecting the focused light from the halfway plane and re-collimating it before it entered the objective lens. Thanks to the symmetry of the system, the collimated beam entering the objective was exactly the same size as the beam reflected from the mirror and the angles of both beams were the same. Despite the angle of the output beam changing, its position at the objective lens always remained fixed in the centre of the optical axis.

As the angle of the beam entering the objective changed, the position of the focal spot on the sample also changed. Since the angle by which the beam was deflected was no more than 0.75° in each axis, the paraxial approximation could be assumed to hold and a linear scaling constant could be used to convert between angle and position. The value of this constant could be calibrated using any sample with visible features on its surface, even if the size of those features was unknown. The laser spot was positioned over a small feature (e.g. the corner of a PhC), the objective was moved by a known distance away from that point in the XY plane, then the mirror was adjusted until the laser spot was in its original position.

Compared to the piezo stage, the scanning mirror was significantly slower. The piezo stage used standard piezo-electric actuators which changed length as a function of the applied voltage. The settling time for short movements was less than 100 ms. In contrast, the scanning mirror used a proprietary slip-stick drive which moved in short increments by expanding and contracting a piezo element. To move any significant distance required many repeated cycles of controlled expansion and contraction, and hence took a lot more time than the simple voltage ramp required by the stage. Additionally the feedback loop used by the scanning mirror to move to a specified position seemed to be under damped, leading to a short period of oscillation after every movement. The average settling time for short movements of the mirror was closer to 500 ms, although it varied significantly. This meant that any excitation maps were slower than the equivalent conventional maps.

3.5 Time-Resolved Micro-Photoluminescence

As opposed to the time-integrated μ PL techniques discussed in Section 3.3, time-resolved photoluminescence (TRPL) measurements involve recording the arrival time of photons with a resolution on the order of a nanosecond or less. There are two main components which make up a TRPL system: the photodetector which turns photons into electrical pulses, and the time-correlated single photon counting card which records the arrival time of those pulses relative to pulses arriving at a synchronisation (sync) input.

For the TRPL studies presented here, a Hamamatsu H7422P-50 photomultiplier tube (PMT) was used as the photodetector. This detector is sensitive over a spectral range from 380 nm to 890 nm, covering all of the emission wavelengths of the perovskite samples studied in chapter 6. This was connected to a PicoQuant TimeHarp 260P time correlated photon counting (TCPC) card in a computer, and the sync input came from the 450 nm diode laser controller.

The TCPC card operates on the principle of recording the arrival time of each voltage pulse it receives. After a pulse arrives on the sync input an internal timer is started. This timer runs for up to 819 ns and is used to record the arrival time

of each subsequent pulse on the signal input accurate to the nearest 25 ps. When another sync pulse is received the timer resets, so all times are recorded relative to the last sync pulse. By repeating this process over millions of sync pulses, a histogram of relative arrival times can be built up. The TRPL system is fast enough to observe the decay of excited carriers created by the absorption of a laser pulse, and a lot of information about a sample's energy levels and carrier dynamics can be extracted by analysing the shape of the decay curve.

One noteworthy feature of the TimeHarp 260P TCPC card is that it has so-called multi-stop capability. Most older TCPC hardware could only record the arrival time of a single signal pulse for each sync pulse. The timer was started by the sync pulse then stopped by the signal pulse, and no further signal pulses could be recorded until the next sync pulse arrived. The multi-stop capability of the TimeHarp card means that multiple signal pulses can be recorded following a single sync pulse.

There are, however, still various factors which can limit the performance of this system. One important one is the dead time of the TCPC card. After detecting an electrical pulse from the PMT on the input channel, there is a delay of up to 25 ns before the recording operation is complete and the electronics are ready to receive another pulse. Pulses received during this dead time are lost. While it might seem reasonable to think that it is always better to have a higher photon flux on the PMT in order to minimise the time required to build up a smooth histogram, this dead time means that an excessively high count rate can lead to some data points being discarded and hence distortions in the shape of the decay curve. More specifically, the histogram tends to be biased towards earlier times.

Another consideration is the response of the PMT to incident photons. An ideal detector would emit a single electrical pulse at the instant a photon arrives, but real detectors can have flaws such as electrical reflections inside the processing circuitry which introduce random delays or secondary peaks to the electrical signal. All of these flaws can be encapsulated in an instrument response function (IRF) which describes the time-domain signal recorded by the TCPC card when the PMT is exposed to a very short pulse of light. An example of an IRF curve is plotted in




Figure 3.6: A plot showing a typical instrument response function of the TRPL configuration used for the studies in this thesis. The largest peak corresponds to the initial detection of the laser pulse, while the smaller secondary peaks are probably caused by electrical reflections.

3.6. This was recorded by using pulses from the 450 nm diode laser as the input to the PMT. The actual signal recorded by the TRPL system during a measurement is the convolution of the real time-response of the sample and the IRF of the detector, so it is necessary to perform deconvolution when analysing the data.

On its own the TRPL system is very useful, but in some cases it is desirable to be able to restrict the range of wavelengths being fed into the PMT, for instance to study the time-domain behaviour of excitons as they cool via phonon interactions or to measure a single emission peak out of many in a quantum well sample. To facilitate this, the spectrometer was used to diffract the light and an adjustable slit on one of its output ports functioned as a band-pass filter to select a narrow range of wavelengths. A full diagram of this configuration is given in Fig. 3.7. The maximum wavelength range which could be selected by the slit was approximately 10 nm when using the standard 600 lines/mm grating.




Figure 3.7: A diagram of the time-resolved photoluminescence measurement system used in this work. The spectrometer and adjustable shutter function as a narrow band-pass filter to provide wavelength selectivity. The time-correlated photon counting card is connected to a computer for recording the data.

3.6 Microscope Imaging

In addition to the standard excitation and collection paths in the system, there were additional optics added to allow the sample to be imaged directly with a silicon CMOS camera while it was in the cryostat. Illumination was provided by a white LED producing uncollimated light which was reflected off a beam splitter into the excitation path of the system. This illuminated the whole field of view of the camera, which was positioned on the collection path with a removable 45° mirror to redirect the light. An example of the kind of images which could be captured with this camera is given in Fig. 3.8.

The main purpose of this imaging system was to allow us to navigate around the sample and locate areas of interest before performing μ PL measurements. It was also essential for aligning the system, finding the correct focal plane, and investigating any problems which occurred. However, in cases where the sample wasn't required to be under the μ PL system, it was better to use a dedicated microscope to produce higher quality images. Most of the microscope images in this thesis were taken using a Keyence VHX-5000 microscope.

3.7 White Light Interferometry

Interferometry is a process in which light travels along two or more different paths before being recombined to produce an interference pattern. There are




Figure 3.8: An image of a PhC waveguide taken using the CMOS camera built into the μ PL system. The waveguide is approximately 20 μm long.

numerous different techniques for interferometry spanning a wide range of fields and applications [37, 38, 39, 40, 41]. In this work, interferometry was used to measure the thickness of a thin film of SU-8 photoresist by recording the reflection spectrum as a function of wavelength.

SU-8 has a refractive index of around 1.6 in the visible and NIR spectrum [42] compared to approximately 4 for GaAs [43] and 1.0 for air. This difference in refractive index means that there is reflection at both the top and bottom interfaces of the SU-8 film. The reflections from the two interfaces can interfere constructively or destructively depending on the phase difference between them, so the photoresist layer effectively acts as a Fabry-Perot etalon [44]. The phase difference is a function of optical path length and wavelength, thus by using a white light source a series of peaks and troughs are observed in the reflection spectrum. The position of those peaks and troughs is determined entirely by the optical path length through the SU-8, hence it is possible to use the spectrum to work out the thickness of the SU-8 layer.

In the ideal case the positions of the peaks in the interference spectrum would be given by the equation $2tn = m\lambda$, in which t is the thickness of the layer, n is




Figure 3.9: A plot of the spectrum of the white LED used for interferometry measurements. This spectrum was taken with the 450 nm long-pass filter in place, so it represents the light which actually reached the sample.

its refractive index, λ is the wavelength, and m is any integer. In practice the presence of the PhC made the situation considerably more complicated, and the solution to this problem is explained in chapter 5.

To measure all wavelengths simultaneously, a white LED covering most of the visible spectrum was used as the source and a 0.3m imaging spectrograph with a silicon CCD was used as the detector. The light from the LED was passed through a fibre before being collimated to a beam and focused onto the sample by a $20\times$ objective. This gave a spot size of approximately $5\text{ }\mu\text{m}$ diameter, allowing thickness measurements to be performed with relatively high spatial resolution. A 450 nm long-pass filter was used to block any photons with enough energy to expose the photoresist. The spectrum of the white LED (after being passed through the filter) is shown in Fig. 3.9

When measuring an interference curve it was actually necessary to record three different spectra: a background with the LED turned off, a calibration spectrum with the LED light reflected off a silver mirror, and the interference spectrum of the LED light reflected off the sample. First the background was subtracted from




Figure 3.10: A plot of an example interference spectrum from a thin film of SU-8 photoresist on top of a PhC waveguide. The data has been normalised by dividing it by the reference spectrum of the light source. The pale blue lines are error bars.

each of the others, then the interference spectrum was normalized by dividing it by the calibration spectrum. An example interference curve is shown in Fig. 3.10.

3.8 Automation Techniques

3.8.1 Automated Mapping and Pattern Writing

Performing μ PL mapping operations manually would be extremely tedious when the number of measurement locations can exceed 10^4 for a single map. Thus a piece of software was written to automate the process. Essentially the algorithm performs a 1D or 2D raster-scan over a grid of points with user-configurable dimensions and step sizes in each axis, recording a μ PL spectrum at each point. The recorded data is then flattened and written to a data file, along with some metadata such as the length scale and exposure time. Movements of both the positioning systems took a non-zero amount of time to complete, so to maximise the quality of the data it was necessary to implement a small delay (≈ 100 ms) between each point in the map and a slightly longer delay (≈ 500 ms) between each row.

As well as saving a lot of human effort, the automated design made it very easy to add a number of extra features. One such feature was an arbitrary angle by which the entire map would be rotated on the sample. This made it very easy to perform maps which were aligned with high symmetry directions on devices such as PhC waveguides. Another feature was the ability to choose different positioning hardware for each axis. For example, it was possible to use the piezo stage to move the objective in the X direction while Y axis movement was provided by the scanning mirror.

For the photolithography work it was necessary to expose various complex patterns on the sample and precise timing was even more important, so a similar system was used to automate that process. We didn't need to record spectra or drive the excitation spot independently of the collection spot, so in many ways this part of the software was simpler to implement. It simply took in an image or collection of vector-based instructions and converted them into movements of the objective lens with variable speed to control exposure conditions. More details of the photolithography process are given in chapter 5.

3.8.2 Automatic Position Optimization

In many cases it was desirable to position the μ PL system precisely over the field maximum of an emitter such as a quantum dot or a PhC cavity. The resolution limit of the optics was on the order of $1\text{ }\mu\text{m}$, however the focal spot still had the shape of an Airy disk so as long as the emitter was approximately point-like and the starting position was close enough, it was always possible to find the field maximum through a process of naive local optimisation. It was possible to do this process manually, however to save time a piece of software was written to automate the process.

The optimisation algorithm basically involves moving in one of the four cardinal directions until a local maximum is found, then changing the movement direction by 90° and repeating the process. The optimisation process ends when there is no change in the maximum position over 8 successive iterations. This is very similar to the way humans perform the same task. Pseudocode for this algorithm is included in appendix A.

3.8.3 Computer Controlled Shutters

At times, particularly when exposing samples of photoresist, it was necessary to block and unblock the laser beam path with timing accurate to a fraction of a second. Two different commercial shutter systems were used during various stages of this work – two Uniblitz shutters and a Newport Model 845HP shutter. In addition to these a third shutter system was constructed from an Arduino microcontroller, a stepper motor, and a light gate sensor.

All the different shutters were integrated into the computer control system. For general μ PL measurements the shutter was automatically closed when recording spectra of the background for later subtraction. The perovskite samples were easily damaged by prolonged exposure to focused laser light, so when doing measurements on those the shutter was configured to automatically open at the start of each measurement and then close immediately afterwards, minimising the potential damage. For exposing the SU-8 photoresist a similar procedure was used, only opening the shutter for the required duration.

4

Materials

Contents

4.1	Introduction	68
4.2	GaAs and other III-As Semiconductors	69
4.2.1	Crystal Structure, Growth and Chemistry	69
4.2.2	Electronic and Optical Properties	74
4.2.3	Low-Dimensional Semiconductor Nanostructures	74
4.3	SU-8 Photoresist	82
4.3.1	Overview	82
4.3.2	Photochemistry	82
4.3.3	Processing	84
4.3.4	Optical Properties	86
4.4	Hybrid Organic-Inorganic Metal-Halide Perovskites	86
4.4.1	Overview	86
4.4.2	Chemistry	87
4.4.3	Electronic and Optical Properties	89

4.1 Introduction

In this chapter I will discuss the three main materials used during the course of my research. The first section covers the III-V semiconductor GaAs and its binary alloys AlGaAs and InGaAs which comprised the photonic crystal samples. Details of the crystal structure, growth processes, and chemistry are given, followed by a discussion of the electronic and optical properties of bulk GaAs, quantum wells,




Figure 4.1: A diagram showing a unit cell of the zinc blende crystal structure of GaAs and its alloys. The red and blue atoms represent gallium and arsenic, but the symmetry of the crystal structure makes it irrelevant which colour represents which element. The structure consists of two inter-penetrating face-centred cubic lattices, one for each element, offset from one another.

and quantum dots. The second section covers the photoresist SU-8 which was used for performing optical lithography on top of photonic crystal waveguides. The chemistry and material properties of SU-8 are discussed, and the steps of the standard photolithography process are described. The final section covers perovskites, particularly the organic-inorganic hybrid metal-halide perovskites studied in chapter 6.

4.2 GaAs and other III-As Semiconductors

4.2.1 Crystal Structure, Growth and Chemistry

GaAs, InAs, and AlAs all take on a zinc blende crystal structure with space group $F\bar{4}3m$ under normal growth conditions. The crystal structure for GaAs is shown in Fig. 4.1.

The lattice constants for these pure semiconductors are listed in Table 4.1. When creating binary semiconductor alloys the lattice parameters can be closely approximated by linearly interpolating between the parameters of the two

Compound	Lattice Constant / Å	Band Gap / eV	Band Gap Type
GaAs	5.65	1.42 (298 K)	direct
InAs	6.06	0.36	direct
AlAs	5.66 (291 K)	2.15	indirect

Table 4.1: Lattice constants and band gaps for the pure III-As semiconductors used in this thesis [48]. All values are given at 300 K except where noted in brackets.

constituent compounds in proportion to their relative concentrations. This principle is enshrined in Vegard's Law [45, 46] which gives the lattice constant, a , for a binary semiconductor alloy, $A_xB_{1-x}Z$, as

$$a = xa_A + (1 - x)a_B \quad (4.1)$$

where a_A and a_B are the lattice constants of the pure semiconductor compounds AZ and BZ respectively. This formula can trivially be extended to alloys containing more than two compounds, although no such alloys are considered in this thesis. In practice there is always some small deviation from Vegard's law [47], and this is often accounted for by adding a bowing term of the form $bx(1 - x)$ to Eq. 4.1.

All of the photonic crystal samples studied in this thesis were grown via molecular beam epitaxy (MBE). First developed in the late 1960s [49], MBE has now become one of the cornerstones of the modern semiconductor industry [50]. It is a process for epitaxial growth of high purity single crystals while allowing the composition to be precisely controlled down to the scale of single monolayers. Fig. 4.2 shows a diagram of a basic MBE system. The elements required for growing the desired crystal (in this case aluminium, gallium, arsenic, and indium) are stored in the effusion cells and heated until atoms or molecules begin to sublime from the surface. The growth chamber is kept under ultra-high vacuum, so the mean free path for the molecules is very long and hence they can be shaped into beams by apertures on the effusion cells. These molecular beams impact the sample which is kept heated to a temperature of several hundred Celsius to improve adsorption. The sample is also rotated about the growth axis to improve the growth uniformity. Most MBE machines include a number of instruments to monitor the growth using




Figure 4.2: A simplified diagram of an MBE system. Each effusion cell contains a different element used in the growth process. The electron gun and fluorescent screen can be used for RHEED measurements of the sample surface as it is being grown.




Figure 4.3: A diagram showing the three basic modes in which epitaxial crystal growth can occur. The lattice mismatch between the substrate and the layer being grown increases from left to right across the three growth modes.

techniques such as reflection high-energy electron diffraction (RHEED) [51] and mass spectrometry. The composition of the layer being grown can be controlled in real time by adjusting the temperature of the elements in the effusion cells and by opening and closing the shutters.

The PhC devices are made up of multiple layers of III-As semiconductors with varying composition. Wherever there is a junction between two crystals with

differing lattice constants, some amount of strain is created at the interface as the two lattices are forced to conform to one another. Depending on the magnitude of the lattice constant mismatch ($\Delta\alpha$), three different growth modes are possible [52, 53]. In the Frank-van der Merwe (FM) mode the crystal grows monolayer-by-monolayer, forming a uniform lattice with a smooth surface and no defects. This is the standard growth mode when there is no lattice mismatch, and crystals can also be grown in this mode as long as $\Delta\alpha$ is small and the thickness of the strained layer does not exceed the critical thickness. Above the critical thickness, the crystal may develop defects or degrade into Stranski-Krastanov (SK) growth. In SK growth mode the crystal no longer forms complete monolayers but instead separates into numerous small islands, typically tens or hundreds of nanometres across and pyramidal in shape. The defining feature of SK growth is the presence of the so-called wetting layer – a uniform layer of FM growth, typically only a few monolayers thick, which occurred before the critical thickness was reached. On the other hand, if $\Delta\alpha$ is so large that even a single monolayer is unstable, then we get Volmer-Weber (VW) growth which is very similar to SK growth except the islands form immediately without any wetting layer. Diagrams of these three growth modes are depicted in Fig. 4.3.

The existence of a critical thickness for FM growth can be understood by considering what happens as each new monolayer is added to the crystal. If $\Delta\alpha$ is small, then the new atoms will be able to conform to the lattice parameter of the substrate without adding much strain energy to the system. As long as FM growth continues and no dislocation defects are introduced, each new monolayer must have almost exactly the same lattice parameter as the one below, thus increasing the total strain energy in the system. It might seem reasonable to think that the lattice parameter could smoothly transition from one value to another as a function of height, but that is only true within small isolated regions. When considering distances of micrometres or millimetres across a wafer, even a tiny variation in lattice parameter from one monolayer to the next would introduce a huge offset in atom positions. As the total strain energy stored in the structure increases, it

becomes more and more energetically favourable to relieve the strain by allowing the lattice to relax to its natural lattice constant [54]. This strain relaxation can only occur by either introducing dislocation defects or by forming islands, and the greater the difference in lattice constants, the lower the critical thickness becomes.

The island formation characteristic of FM and SK growth occurs because it allows each island to relax independently which requires far less energy than strain relaxation across an entire wafer. The distances involved are much shorter, so the absolute offset in atom positions from one monolayer to the next is smaller for the same difference in lattice parameter. Newly deposited atoms preferentially adhere to the top surfaces of existing islands instead of extending them laterally, so islands grow faster than the surrounding surface. At the high temperatures typically used during MBE processing it is also possible for atoms to diffuse across the surface to more energetically favourable locations, and in fact this is often the dominant process responsible for island formation [55].

Pure InAs on GaAs has a critical thickness of around 3 monolayers before the growth transitions to the SK mode [56]. This would increase slightly for InGaAs alloys, since the lattice constants would be more closely matched. In contrast, AlAs grown on GaAs has an exceptionally large critical thickness of $1\text{ }\mu\text{m}$ [57] before dislocation defects become energetically favourable. This makes AlGaAs alloys very useful for fabricating thick vertical heterostructures without introducing defects into the crystal.

The real value of AlGaAs for fabricating PhCs comes from its chemical properties. $\text{Al}_x\text{Ga}_{1-x}\text{As}$ is readily soluble in both hydrofluoric acid (HF) [58] and hydrochloric acid (HCl) [59], being etched at a rate on the order of $1\text{ }\mu\text{m}/\text{min}$ for values of x above about 0.5. In contrast pure GaAs is quite resistant to both HCl (etch rate $\approx 0.01\text{ nm}/\text{min}$) and HF (etch rate $\approx 0.1\text{ nm}/\text{min}$), and InGaAs is not etched at all by those acids even for blends with as little as 15 % indium [60]. That huge difference in etch rate of 4–5 orders of magnitude makes AlGaAs a very good material for growing sacrificial layers. In particular, our collaborators used this

technique to create the suspended membrane structures in all of the 2D PhCs studied in this thesis. More details of this are given in chapter 5.

4.2.2 Electronic and Optical Properties

Both GaAs and InAs are direct gap semiconductors and their band gaps are listed in table 4.1. The direct gap makes them efficient optical absorbers and emitters. The band gap of pure GaAs gives it a photoluminescence emission peak at around 870 nm, while pure InAs has a very small band gap and consequently emits at around 3.4 μm . The band gap of semiconductor alloys follows a similar near-linear relationship to that given in equation 4.1, so the band gap of InGaAs can be tuned over a wide range in the infrared by adjusting the composition.

4.2.3 Low-Dimensional Semiconductor Nanostructures

As with many physical systems, the macroscopic properties of bulk semiconductors are qualitatively different from the properties of semiconductor structures with nanometre length scales. In particular, in this section we will consider low-dimensional nanostructures made of a material with a small band gap, such as InGaAs, embedded inside a bulk material with a larger band gap, such as GaAs.

The low-dimensional structures of interest to this thesis are shown in Fig. 4.4(a). If we imagine reducing the length scale of a bulk semiconductor along a single axis, it must eventually become small enough that the movement of electrons and holes is effectively confined to a quasi-2D plane. Such a structure is called a quantum well (QW). This reduction in dimensionality alters the density of states available for electrons and holes to occupy, as shown in Fig. 4.4(b), and can therefore produce some very useful optical properties. Reducing the dimensionality of the structure further leads to quasi-1D quantum wires (not pictured) and quasi-0D quantum dots (QDs). The critical length scale at which this reduction in dimensionality occurs is approximately equal to the de Broglie wavelength of the carriers [61].




Figure 4.4: (a) Diagrams showing the spatial structure of bulk semiconductor, quantum wells, and quantum dots. (b) Idealised plots of the density of states (DOS) against energy for carriers in bulk semiconductor, quantum wells, and quantum dots.

Quantum Wells

The concept of the quantum well was simultaneously proposed in 1963 by Alferov *et al.* [62] and Kroemer [63] as a means of creating a more efficient solid state laser. Since then QWs have become a mainstay of solid state technology, being used not just in laser diodes but also in high electron mobility transistors [64], IR photodetectors [65], quantum hall effect sensors [66], and solar cells [67], among others. Quantum wells are quite easy to fabricate by using MBE or metalorganic vapour phase epitaxy (MOVPE) to deposit a thin layer of low band gap semiconductor between two layers of high band gap semiconductor.

Many of the properties of QWs can be understood by using the extremely simple 1D “particle-in-a-box” model, as detailed by e.g. Miller [68]. In this model the electrons and holes are assumed to be free to move in two spatial dimensions but confined in the third by an infinite potential barrier on both sides. We will begin by solving Schrödinger’s equation in 1D along the direction of confinement,

$$\frac{-\hbar^2}{2m} \frac{d^2 \phi_n}{dz^2} + V(z) \phi_n = E_n \phi_n, \quad (4.2)$$

where m is the effective mass of the electron or hole and E_n and ϕ_n are the eigenenergy and wavefunction of the n^{th} solution. The potential, V , as a function of position, z , is given by,

$$V(z) = \begin{cases} \infty & z \leq 0 \\ 0 & 0 < z < L \\ \infty & L \leq z \end{cases}, \quad (4.3)$$

for a QW of thickness L . In this extremely simplistic model the solutions are sinusoidal standing waves with energies of the form,

$$E_n = \frac{-\hbar^2}{2m} \left(\frac{n\pi}{L} \right)^2, \quad n = 1, 2, 3, \dots \quad (4.4)$$

This means that the motion of the electrons and holes in the direction perpendicular to the QW layer is constrained to a series of standing waves with discrete energies. However, in the plane of the QW the electron velocities are unconstrained so their total kinetic energy can take any value greater than the lowest discrete level. Therefore the density of states of this idealised quantum well consists of a series of steps as shown in Fig. 4.4. The optical absorption spectrum of a quantum well follows this step pattern quite well, but thanks to phonon-mediated relaxation the PL emission tends to take the form of a series of peaks corresponding to the low-energy end of each “step”. For shallow quantum wells there may only be a single confined state leading to a single peak in the PL spectrum.

Quantum Dots

Quantum dots, first developed by Ekimov *et al.* [69] and Rossetti *et al.* [70], are similar to quantum wells except the carriers are confined by heterojunctions in all three spatial dimensions. QDs have a number of very useful optical properties which have made them attractive for many applications including solar cells [71], light emitting diodes [72], quantum computing [73], biological imaging [74], and display backlighting [75].

The energy levels and density of states of a QD can be calculated in much the same way as we did for a QW by treating it as a simple square potential well

and solving the Schrödinger equation [76]. This is the same as Eq. 4.2, but for a QD the potential, V , must be a 3D function describing barriers on all sides. The solutions are once again a series of approximately sinusoidal wavefunctions with discrete eigenenergies, however, because the carriers are not free to move in any direction there is no continuum of states. The density of states of a QD is therefore a sequence of one or more delta functions, much like that of an atom. In fact, thanks to the similarities in quantum confinement and the resulting energy levels, QDs are sometimes referred to as artificial atoms.

The photoluminescence spectrum of a single QD is typically observed to be a single narrow spectral line from electrons in the $n = 1$ conduction band level recombining with holes in the $n = 1$ valence band level. While the simple “particle-in-a-box” model would suggest these lines are as spectrally narrow as those emitted by electronic transitions in atoms, in practice they can be considerably broader both due to shorter radiative lifetimes and due to effects such as spectral diffusion [77]. The energy gap between the lowest electron level and the lowest hole level in a QD always lies in the range between the band gaps of the two semiconductors forming the heterojunction. The exact value of this gap, and hence the principal emission wavelength, is determined by the spatial size of the QD. Smaller QDs exhibit stronger quantum confinement and hence emit at shorter wavelengths, although below a certain size and trapping potential there don’t exist any bound states. The emission energy of QDs tends to decrease as the temperature is increased [78], both due to the QD physically expanding and because the band gaps of the underlying semiconductors tend to decrease. At higher temperatures the carriers cease to be confined by the potential well and PL emission ceases. The temperature at which this happens depends on the potential well depth and is typically much less than 100 K for InGaAs dots [79].

The semiconductor quantum dots used in this research were all grown via self-assembly in the Stranski-Krastanov growth mode, as discussed in section 4.2.1. The small pyramidal islands of localised growth produced in the SK mode are typically within the size range required for quantum confinement effects to manifest, and their




Figure 4.5: A plot of the micro-photoluminescence spectrum of an ensemble of quantum dots excited by a continuous wave laser above the band gap.

average size can be adjusted within certain limits by altering the growth parameters. The wetting layer functions as a quantum well creating additional PL emission, although this is typically much thinner than the QDs and hence well separated spectrally. The QDs produced by SK growth can be very high quality and free of impurities and defects, leading to high quantum efficiency [80]. The disadvantage of self-assembled quantum dots grown on a uniform substrate is that both the size and position of the dots are random within a broad statistical distribution. Size variations mean that an ensemble of many self-assembled QDs exhibit significant inhomogeneous broadening, typically at least 20 nm wide for InGaAs QDs emitting at around 1300 nm. An example spectrum of some QDs is shown in Fig. 4.5. The random positioning and emission wavelength makes it very difficult to couple single QDs to photonic crystal cavities defined using e-beam lithography.

It is possible to grow self-assembled QDs in pre-determined locations by

etching small pits into the substrate prior to depositing the QD layer [81]. These site-controlled QDs can be positioned accurate to less than 50 nm and can have much reduced size variation, however the best achievable single-dot linewidths are still almost two orders of magnitude worse than those of randomly positioned QDs ($\sim 43 \mu\text{eV}$ [82] compared to $\sim 1.5 \mu\text{eV}$ [83]). Since high photon indistinguishability is a requirement of many quantum computation and cryptography schemes, this makes site-controlled QDs significantly less attractive as single photon sources.

There are two main schemes for optical excitation of QDs: resonant excitation, in which photons are absorbed by the QDs themselves, and off-resonant excitation in which photons are first absorbed by the bulk semiconductor then the resulting carriers are captured by the QDs. In resonant PL measurements it is often difficult to distinguish between the excitation laser and the emitted light, so all of the QD studies in this thesis were performed off-resonantly. Off resonant excitation can be performed with photons of almost any energy as long as it is higher than the band gap of the bulk semiconductor, in this case GaAs so $\lambda < 870 \text{ nm}$ at room temperature. Carrier relaxation from the continuous bands to confined QD states can occur either by phonon interactions or by multi-carrier Auger processes [84].

When electrons and holes are confined inside a QD, they are held in sufficiently close proximity for the Coulomb attraction to become significant. This electrostatic attraction causes them to bind into a neutral quasi-particle known as an exciton. These excitons can also exist in quantum wires, quantum wells and even bulk semiconductors at sufficiently low temperatures, but the full 3D confinement in a QD can allow them to exist even at room temperature. The type of excitons found in semiconductors such as GaAs are known as Wannier-Mott excitons and have a radius much larger than the lattice spacing. The binding energy of excitons in bulk semiconductor crystals is very small, typically no more than a few tens of meV, due to the large radius and the shielding effect of the valence electrons [85]. The strong confinement inside a QD can increase this binding energy up to almost 1 eV (for dots of approximately 1 nm diameter), with smaller dots producing higher




Figure 4.6: Various bright excitonic states of a quantum dot. (a) and (b) are neutral excitons. (c) is a charged exciton. (d) is a biexciton.

exciton binding energies [86]. The exciton binding energy manifests optically as a reduction in the photon energy required to create electron-hole pairs.

Quantum dots have another property in common with atoms: they are small enough for the Pauli exclusion principle to apply to their energy levels. Therefore the lowest energy level of the QD (equivalent to the 1s shell of an atom) can only contain up to two carriers, one with spin $+1/2$ and the other with spin $-1/2$. The convention used here is that a spin ‘up’ hole represents the absence of a spin ‘down’ electron in the valence band. Optical selection rules therefore require that when a photon is absorbed or emitted, it must involve the creation or annihilation of an electron-hole pair with opposite spins. This means there are only a limited number of possible exciton configurations in the lowest energy level of the QD, and only a subset of those are optically accessible (bright). Some examples are given in Fig. 4.6. The binding energies of the two neutral excitons, Fig. 4.6(a) and 4.6(b), are almost exactly the same, only separated by fine structure splitting [87] which could not be resolved with the instruments used in this thesis.

The biexciton shown in Fig. 4.6(d) is not simply two independent excitons, but is actually a bound state containing two electrons and two holes [88]. The binding energy of the biexciton is a function of the ratio of electron and hole effective masses, m_e^*/m_h^* , although the biexciton is stable for any ratio [89]. A diagram of the energy level structure of a biexciton is given in Fig. 4.7. One important




Figure 4.7: Energy level diagram of a biexciton. $|g\rangle$, $|e\rangle$, and $|b\rangle$ are the ground, exciton, and biexciton states respectively. B_{XX} is the biexciton binding energy.

feature to note is that there are two different pathways via which a biexciton can decay to emit photons. If the two intermediate states are degenerate then the two pathways produce pairs of cross-circularly polarised photons, and because the choice of decay pathway is not observed the emitted photons are polarisation-entangled [90]. Thus QDs can function as sources of entangled photon pairs for quantum information applications. If the intermediate states are non-degenerate (usually due to the exchange interaction between the electrons and holes) then the emitted photons are not entangled (because the which-path information is revealed by their energies) and they both have the same linear polarisation (either both horizontal or both vertical, in some basis) [91, 92].

Another important property of excitons in QDs is that only one photon at the lowest exciton energy can be emitted at a time. Two photons can be emitted during biexciton decay, but one will be at a different wavelength due to the biexciton binding energy. After one exciton decays from the lowest energy level there must necessarily be some small delay before another electron hole pair can relax down to replace it. This means that QDs inherently act as single photon sources [93].

4.3 SU-8 Photoresist

4.3.1 Overview

The main photoresist used throughout this work is SU-8 [94]. It is a negative photoresist meaning it polymerizes when exposed to light, as opposed to positive photoresists which become soluble when exposed to light. SU-8 photoresist is a solution of two active ingredients: EPON SU-8 epoxy resin and a photoacid generator, typically a triarylsulfonium hexafluoroantimonate salt. The solvent in commercially available SU-8 preparations is either gamma butyrolactone or cyclopentanone.

Probably the most common use of SU-8 is in high aspect ratio photolithography for fabricating very thick structures ($> 10\text{ }\mu\text{m}$) with high resolution [95]. Those thick SU-8 structures are typically used for applications such as microelectromechanical systems (MEMS) [96, 97, 98] and microfluidics [99, 100, 101]. SU-8 has also been used for processes such as nanoimprint lithography [102, 103] and as a mask for electroplating [104] and various types of etching [105]. Although much less common, SU-8 can also be used for fabricating sub-micron thickness structures [106, 107, 108]. SU-8 was designed to be a conventional photoresist, however it can also be exposed by e-beam lithography [109] and two-photon processes [110].

4.3.2 Photochemistry

Each molecule of the SU-8 resin is made up of a chain of Bisphenol A Novolac with multiple functional epoxy groups attached to it, as shown in Fig. 4.8. A typical molecule contains 8 epoxy groups, hence the name “SU-8”, however this is only a statistical average. A preparation of SU-8 photoresist will contain Bisphenol A Novolac chains of many different lengths with corresponding numbers of epoxy groups.

The polymerization of SU-8 actually occurs in two distinct stages, as shown in Fig. 4.9 [111, 112]. The first is the photochemical process in which light is absorbed by the photoacid generator and splits it to produce an acid catalyst. In




Figure 4.8: A diagram showing the chemical structure of a typical SU-8 epoxy oligomer.




Figure 4.9: A diagram showing the two main chemical reactions involved in the photopolymerisation process. The first reaction is the splitting of a photoacid generator by absorption of a photon to produce an acid catalyst. The second reaction is the actual polymerisation process in which that acid reacts with an arbitrary number of epoxy groups on adjacent SU-8 molecules to bond them together.

the second stage this acid then initiates the polymerization reaction, causing epoxy groups from nearby SU-8 molecules to bond together. A single absorbed photon can result in many SU-8 molecules being crosslinked together in a chain reaction, so for this reason SU-8 is said to be a chemically amplified photoresist [113]. The large number of epoxy groups on each SU-8 resin molecule leads to a very high degree of crosslinking in exposed SU-8 photoresist. This gives it good mechanical properties, heat resistance, and chemical stability [114, 115, 116].

4.3.3 Processing

The standard protocol for using SU-8 photoresist is shown in Fig. 4.10. The information in this section is largely based on a technical datasheet published by the manufacturer, MicroChem [117].

The recommended way to deposit a uniform layer of SU-8 onto a substrate is through spin-coating. A droplet of liquid SU-8 solution is placed in the centre of the sample, then it is rapidly spun at a defined speed. The centrifugal force causes the SU-8 to flow outwards and fall off the perimeter of the sample at a rate determined by its viscosity. At the same time the solvent is evaporating, causing the viscosity of the SU-8 to increase. The thinner the SU-8 layer becomes, the greater the relative rate of solvent evaporation. These processes compete until the viscosity becomes high enough to prevent any further thinning, leaving a uniform thin layer of non-volatile material [118]. The balance between these competing processes can be altered in various ways. Increasing the spin speed, decreasing the concentration of non-volatile compounds, and replacing the solvent with one with a lower viscosity or greater volatility will all decrease the thickness of the resulting layer.

Before exposing the photoresist it is common to perform a heating step known as a softbake or pre-bake. The spin coating process leaves a considerable amount of solvent in the SU-8 layer (4–24 wt%, depending on the thickness) and heating helps to remove this. This softbake step also improves adhesion and reduces any residual stress in the SU-8 layer [119]. For SU-8 the recommended softbake temperature is 95°C, although some groups report success using temperatures as low as 65°C [120].




Figure 4.10: A diagram showing the standard process used to pattern SU-8 photoresist on a substrate.

SU-8 photoresist is designed to be exposed using UV light at a wavelength of 365 nm, although it is sensitive to light at 405 nm with reduced sensitivity [108]. For larger patterns a mask aligner is typically used, however in this work exposure was performed using direct laser writing. It is also possible to expose SU-8 using a two-photon absorption process with light at a wavelength around 800 nm [121, 122]. Since this is a highly non-linear optical process it is necessary to use a pulsed laser with extremely high peak power density.

The initial photoacid generation process is irreversible, so once the acid catalyst is created it remains in the mixture. The polymerisation reaction proceeds extremely slowly at room temperature, but heating the SU-8 above its glass transition temperature allows the molecules to move much more freely and accelerates the process. The glass transition temperature for uncured SU-8 is around 50°C [114], so

a post-exposure bake (PEB) step is required. Care must be taken to ensure that both the softbake and PEB do not exceed 120°C, since higher temperatures can thermally activate the photoacid generator and cause the entire SU-8 film to crosslink [119].

The recommended developer for SU-8 is 1-Methoxy-2-propanol acetate, also known as propylene glycol monomethyl ether acetate (PGMEA) [123], however various other organic solvents such as ethyl lactate [124] and diacetone alcohol [125] have been reported to work. A typical development process involves immersing the sample in PGMEA for several minutes then rinsing it with isopropyl alcohol (IPA) to remove any residue.

4.3.4 Optical Properties

SU-8 exhibits very low optical losses at wavelengths longer than 400 nm, and the losses remain negligible all the way up to 1310 nm making it an ideal material for fabricating photonic structures [108, 126, 14]. Its refractive index at 1310 nm and room temperature has been reported as 1.574 [14], however this can be expected to vary significantly depending on the layer thickness and the details of the fabrication process [108].

4.4 Hybrid Organic-Inorganic Metal-Halide Perovskites

4.4.1 Overview

Perovskites are a large family of materials named after the mineral perovskite, CaTiO_3 , with which they share a common crystal structure. Oxide perovskites have been of scientific interest since at least the 1920s, but more recently a wide variety of different perovskite chemistries have been developed with interesting and useful properties. Of particular note among these are the hybrid organic-inorganic metal-halide perovskites which exhibit semiconducting properties and have proven remarkably effective as photovoltaic materials. Over the last decade they have seen a meteoric rise in efficiency and have garnered a great deal of scientific interest

as a result. Single-junction efficiencies as high as 23% have been reported [127], making them competitive with the best silicon (27%) and GaAs-based (29%) cells. Good quality perovskite films are comparatively easy to fabricate through standard solution processing and vapour deposition techniques [128, 129, 130], potentially allowing for much lower production costs than conventional silicon-based solar cells.

The perovskites studied in chapter 6 are mixed-cation lead mixed-halide perovskites with the formula $\text{FA}_{0.83}\text{Cs}_{0.17}\text{Pb}(\text{I}_{1-x}\text{Br}_x)_3$, where FA stands for formamidinium or $\text{HC}(\text{NH}_2)_2$.

4.4.2 Chemistry

Perovskites all follow the same basic chemical formula of ABX_3 in which A and B are cations and X is an anion. The A cation is the larger of the two while the B cation is smaller [131]. Many different choices of A, B, and X are possible, although not all combinations have desirable properties. Common choices for the A site cation include Rb^+ , Cs^+ , CH_3NH_3^+ (methylammonium or MA), and $\text{HC}(\text{NH}_2)_2^+$ (formamidinium or FA). B is a group IVA element in a +2 oxidation state, most commonly Pb^{2+} but Sn^{2+} and Ge^{2+} have also been used. X site anions are normally halides, and common choices include Br^- , I^- , and Cl^- . Just as with III-V semiconductors, it is possible to create perovskite ‘alloys’ containing mixtures of several different elements or compounds in the A, B, or X sites. For example, the perovskites materials studied in chapter 6 contain varying mixtures of iodine and bromine in the X site, and a mixture of FA and Cs in the A site.

The basic cubic perovskite crystal structure is shown in Fig. 4.11. The space group of this structure is $Pm\bar{3}m$, and it is often referred to as the α phase or $\alpha\text{-ABX}_3$. It is easiest to visualise it as a regular array of BX_6 octahedra, each joined to six others at its corners, with an A site cation located at the centre of each cube of eight B site cations. The sizes of the cations in the A and B sites (green and blue respectively) and of the anions in the X sites (red) are of critical importance in determining the crystal structure and hence many of the physical and electronic properties of the resulting material. If the ionic radius of B is too small compared to





Figure 4.11: A diagram showing the cubic perovskite crystal structure for a perovskite with the formula ABX_3 . (a) is a 3D rendering of the unit cell. (b) is a 2D cross-section through the crystal structure showing the BX_6 octahedra (pale blue diamonds) which are predominantly responsible for determining the electronic properties of the perovskite.

that of X or if the radius of A is too large or small compared to B then a perovskite crystal structure is unable to form. In fact the formability of perovskite structures can be predicted with very high reliability (96%) by considering only the ionic radii of the constituent components [132]. The two size dependent formability metrics are commonly known as the *tolerance factor* and *octahedral factor* in the literature.

Both pure $FAPbX_3$ and $CsPbX_3$ have tolerance factors either side of the optimal range and hence have a tendency to spontaneously transition to their hexagonal (δ_H) and orthorhombic (δ_O) phases respectively at room temperature. When this transition occurs the band gap increases significantly, degrading the photovoltaic performance, and the material changes to a translucent yellow. It is possible to combat this by alloying the two materials together, giving an average tolerance factor close to 1 and significantly improving the stability [133]. Ab initio structural simulations of various compositions have also shown that a mixture of FA and Cs cations results in greater stability of the α phase compared to either cation alone [134].


One of the main problems preventing perovskites from seeing more widespread use as photovoltaics is their relatively poor chemical stability. For illustrative

purposes the following discussion will focus on MAPbI₃, but analogous breakdown processes occur for other hybrid perovskites including FAPbI₃ [135, 136]. MAPbI₃ exists in a chemical equilibrium described by the following equation [137, 135]:



Under favourable conditions this equilibrium is dominated by the perovskite compound on the left hand side. The presence of moisture, however, causes the physical separation of the breakdown products on the right due to differences in solubility and hence suppresses the reverse reaction [138]. This causes the equilibrium to shift to the right and eventually results in the decomposition of almost all the perovskite.

As well as being highly sensitive to moisture, hybrid perovskites are also susceptible to breakdown from exposure to a combination of oxygen and light. When exposed to light the CH₃NH₃I produced by equation 4.5 reacts with oxygen according to the following equation [137]:



This not only drives equation 4.5 to the right, but it also produces water which further accelerates the breakdown of the perovskite material. Certain other oxygen-containing compounds such as CO₂ are similarly destructive to hybrid perovskites [137]. Replacing the Pb in the B site with lighter group IVA elements such as Sn makes the perovskite even more susceptible to chemical breakdown when exposed to oxygen [131]. On the other hand, adding Cs to FAPbI₃ makes it more resistant to moisture and photochemical breakdown [139].

4.4.3 Electronic and Optical Properties

Many perovskites including MAPbI₃ and FAPbI₃ have a direct band gap, making them efficient optical absorbers and emitters [131, 140]. This is particularly useful for photovoltaic applications since the same amount of light can be absorbed by

a much thinner layer of perovskite compared to silicon, reducing the material costs and weight.

The band gap energies of perovskites are highly tunable by varying the composition [141]. For instance FAPbI_3 has a band gap of around 1.48 eV ($\lambda \approx 840$ nm) at room temperature while that of CsPbI_3 is around 1.73 eV ($\lambda \approx 720$ nm) [142]. The band gaps of certain mixed methylammonium lead-tin halides can go as low as 1.17 eV [143] and the band gap of MAPbBr_3 is 2.3 eV [144].

The choice of halide can have a significant effect on the band gap because it influences the length of the bonds in the BX_6 octahedra [142]. Bromine has a significantly smaller atomic radius than iodine, hence bromine-rich perovskites tend to have a larger band gap than their iodine-rich counterparts [145, 146]. This allows for significant tunability of the band gap of the $\text{FA}_{0.83}\text{Cs}_{0.17}\text{Pb}(\text{I}_{1-x}\text{Br}_x)_3$ composition by simply altering the value of x . Chlorine has also been used in some perovskite formulations to achieve even higher band gaps due to its smaller atomic radius [147].

This high degree of band gap tunability is particularly useful for photovoltaic applications where the maximum theoretical efficiency of a single semiconductor junction is a function of the band gap. A smaller band gap means that photons with lower energy are absorbed and hence more of the solar spectrum can be converted to electricity. On the other hand, a smaller band gap results in more energy being wasted as heat when a high energy photon is absorbed and the ensuing carriers relax to the band edges via phonon interactions. This trade-off, as well as various other factors affecting efficiency, is described by the Shockley–Queisser limit [148]. The peak efficiency for a single p-n junction on the Earth’s surface is 33.77% at a band gap of 1.34 eV [149]. Higher efficiencies can be achieved by stacking multiple p-n junctions with different band gaps [150, 141], but this still requires fine control over the band gaps of the different layers.

Interestingly the electronic structure of the Cs, MA, or FA cation at the A site has only a minimal effect on the band structure of hybrid metal halide perovskites. The main contribution to the band structure comes from the electrons within the BX_6 octahedra, while the A site cation only has an indirect influence

by affecting the lattice constant and symmetry of the crystal structure [131, 151]. This is because the electronic transitions of the organic cation lie deep within the conduction and valence bands, far from any band edges [152].

One of the most attractive and surprising properties of hybrid metal-halide perovskites is their extremely high defect tolerance. The electronic states introduced by native point defects in the crystal lattice typically have energies either very close to the band edges or completely outside the band gap [153, 154]. These shallow defect states have much less of an adverse effect on device performance than mid-gap states.

Many hybrid metal-halide perovskites have been found to have very favourable carrier diffusion lengths for photovoltaic applications, however this strongly depends on both the chemical composition and the crystal quality. Electron and hole diffusion lengths in solution processed MAPbI₃ films have been shown to be on the order of 100 nm [155], while those of FAPbI₃ are about 3.1 μm [156]. Growing monocrystalline perovskites is much more challenging, but the carrier diffusion length in FAPbI₃ and FAPbBr₃ single crystals are an impressive 6.6 μm and 19 μm respectively [157]. There have even been reports of carrier diffusion lengths as long as 175 μm under 1 sun illumination in large single crystals of MAPbI₃ [158].

5

Lithographically Defined Photonic Crystal Microcavities

Contents

5.1	Introduction	93
5.2	Motivation and Background Theory	93
5.2.1	Single Photon Sources	93
5.2.2	Cavity Quantum Electrodynamics	99
5.2.3	Advantages of Photolithography for Cavity Fabrication	104
5.3	Device Design	105
5.4	Optimising SU-8 Photolithography	106
5.4.1	Thickness Optimisation on Smooth GaAs	107
5.4.2	Two-Photon Exposure and μ PL Apparatus	110
5.4.3	Cryogenic Exposure and Dose Optimisation	112
5.4.4	Spin-coating on PhC membranes	114
5.4.5	Problems With Two-Photon Exposure	117
5.4.6	Blue Exposure	120
5.4.7	White Light Interferometry	121
5.5	Cavity Fabrication and Optical Characterisation	125
5.5.1	Experimental Procedure	125
5.5.2	Results	129
5.6	Conclusions	134




Figure 5.1: A diagram of a photo-multiplier tube (PMT) detecting individual photons from a beam of light. The detection times on the right are described by the photon statistics of the incident light.

5.1 Introduction

In this chapter I discuss the experimental implementation of the photonic crystal (PhC) cavity design first considered in chapter 2. The cavities are created using a combination of both e-beam lithography and photolithography, however the focus of this chapter is on the photolithography stage of that process. The first section covers the background theory of single photon sources and coupled emitter-cavity systems, as well as explaining a bit about the motivation for creating this kind of cavity in the first place. Section 5.3 gives details of the target structure, while section 5.4 discusses the development and refinement of SU-8 processing techniques necessary to fabricate it in practice. In the final section I give details of the fabrication process which successfully produced high- Q cavities and discuss the micro-photoluminescence (μ PL) techniques used to characterise them.

5.2 Motivation and Background Theory

5.2.1 Single Photon Sources

Let us start by considering a beam of light shining on a detector such as a CCD or PMT. In the classical model, the beam of light can be described as a continuous electromagnetic wave obeying Maxwell's equations. In contrast, quantum theory tells us that the detector will produce a sequence of discrete signals as each quantised photon is absorbed, as shown in Fig. 5.1. The precise statistics of these photon detection events are known as the *photon statistics* of the light and will be the focus of this section.

In the ideal case of a monochromatic light source such as a single frequency laser emitting at constant intensity, the times of the photon detection events are completely uncorrelated. This is because the frequency (equivalent to energy) and arrival time of a photon are non-commuting quantum observables and hence cannot both be known to arbitrary precision [159]. If the frequency is known exactly, then the arrival time is completely unknown. In classical terms, this corresponds to the fact that the more you reduce the duration of a pulse, the broader it must become in the frequency domain and vice versa.

If we then separate the detection events into finite time bins, the number of photons in each bin will follow a Poissonian distribution with the probability, $P(n)$, of observing n photons within any single time bin given by the equation,

$$P(n) = \frac{\bar{n}^n}{n!} e^{-\bar{n}} \quad (5.1)$$

where $\bar{n}$ is the mean number of photons per bin [160]. The standard deviation, σ , of a Poisson distribution is equal to $\sqrt{\bar{n}}$, which is a well known result for the purposes of error estimation. The standard deviation also serves as a useful way of characterising other photon statistics.

Most light sources in the real world do not have an exactly constant intensity. The classic example is the random fluctuation in the intensity of light from a black-body source, but almost every conventional light source suffers from some intensity variations. These fluctuations in intensity mean that the photon detection events are no longer uncorrelated. More photons are detected during periods when the intensity is higher, so knowing that one photon has been detected at a given time allows you to infer that there is a slightly higher probability of other photons being detected close to that time. Similarly, knowing that there have been no photon detections in a particular time window allows you to guess that the intensity is lower at that moment and hence there is a lower probability of other detection events. In other words, the photon detection events tend to occur in bunches.

The effect of this photon bunching on the statistical distribution of photon counts in each bin is to broaden it. Effectively the Poisson distributions at times of




Figure 5.2: Plots showing examples of Poissonian, sub-Poissonian, and super-Poissonian functions. All are normalised to the same total area. n indicates the number of photons observed in a particular time bin, $\bar{n}$ is the average, and σ is the standard deviation.

low and high intensity get averaged together, producing a total distribution which is no longer Poissonian. σ becomes greater than $\sqrt{\bar{n}}$, so we describe such light sources as super-Poissonian. This is visualised in Fig. 5.2.

So far all of the light sources we have considered have been semi-classical; The photon detection events have been quantised but the light itself could be fully described by classical electric and magnetic fields. However, in order to explore the third type of photon statistics – sub-Poissonian statistics – we need to consider a beam of light made up of quantised fields with no classical equivalent.

The amplitude (in other words the photon flux), n , and the phase, ϕ , of a beam of light are linked by an uncertainty relation which places a lower bound on the product of their uncertainties, $\Delta n \Delta \phi \geq 1/2$. The light from a classical coherent source distributes this noise equally between the amplitude and phase, however in the quantum regime it is possible to ‘squeeze’ the photon statistics such that the uncertainty in the measurement of one observable is much lower than the other [161]. In the case of single photon sources, we’re interested in squeezed light with reduced uncertainty in the amplitude. If this amplitude-squeezed light is

measured by a detector, the spacing of the photons appears more regular than a classical light source and hence the standard deviation of the photon count per bin is sub-Poissonian. An example plot of such photon statistics is given in Fig. 5.2. In the limit of maximal amplitude squeezing (and assuming a perfectly efficient detector) the photons would arrive in a regular train with constant spacing and exactly the same number would be detected in each time bin. A single photon source is simply the special case where the number of photons per time bin is never greater than one.

Quantum Key Distribution

Probably the greatest interest in single photon sources has come from researchers studying quantum key distribution. Quantum key distribution (QKD) is technique which allows data to be shared securely between physically separated entities while guaranteeing with arbitrarily high confidence that the data has not been read by a third party [162]. The name comes from its principal use case wherein two individuals use it to agree on a private encryption key which can then be used to securely transmit much larger quantities of data over a classical communication link.

Currently the internet relies upon the mathematics of public key cryptography to allow encrypted communication [163], however Shor's algorithms running on a quantum computer could solve the underlying prime factorisation and discrete logarithm problems in polynomial time [164] – exponentially faster than a classical computer could solve the same problems. This means that in the relatively near future it may be feasible for entities with sufficient resources to break public key cryptography. QKD is one possible replacement for the current algorithms which would be resistant to both classical and quantum attacks. After agreeing on an encryption key via QKD, symmetric encryption algorithms such as AES could be used which, given a key size of at least 256 bits, is not vulnerable to any known quantum algorithms [165, 166]. A one-time pad could also be used, giving an even stronger guarantee of security at the cost of requiring a much larger key.

There have been a number of different protocols developed for QKD, but they all depend upon the basic principle of quantum mechanics that measuring




Figure 5.3: A diagram showing an idealised photon number splitting attack against a quantum key distribution system. Eve uses a beam splitter to capture only some of the photons from multi-photon states, allowing her to intercept the communication. Single photon sources are one way of preventing such attacks.

a superposition state will destroy it. More specifically, QKD protocols can be separated into two categories: prepare-and-measure protocols which rely on the quantum no-cloning principle [167] to guarantee secrecy, and entanglement-based protocols which rely on the monogamy of entanglement [168] to do the same. In this section I will briefly explain one of the best known QKD protocols and its relevance to single photon sources.

The BB84 protocol, developed by Bennett and Brassard [169], was the first example of a QKD protocol. It falls into the category of prepare-and-measure protocols, since the sender (traditionally labelled Alice) prepares a photon in a specific state and then sends it to the recipient (Bob) who measures it. When preparing each photon, Alice randomly chooses one of two conjugate bases to work in then maps 0 and 1 to two orthogonal states within that basis. The mapping within each basis is public knowledge, but the choice of basis is known only to Alice. Typically the two bases used are the linear polarisations ‘+’ (0° and 90°) and ‘ $\times$ ’ (45° and -45°), although other choices are possible. The intended recipient, Bob, independently chooses a random basis to measure each photon in. If he uses the same basis as Alice the measured value should be the same as the one sent by Alice, however if he measures the polarisation in the other basis then the value will be

completely random. An eavesdropper, Eve, can intercept photons and perform her own measurements on them, but without prior knowledge of the correct basis she can only guess with 50% probability of success. After the quantum transmission has finished, Bob then sends a list of which photons he received and which basis he used to measure them (but not the result of that measurement) to Alice over an insecure classical channel. They can use this information to discard the photons which were transmitted and measured with different bases. Photons which were not received by Bob do not form part of the final key either, so Eve is forced to re-transmit any photons she intercepts. The quantum no-cloning principal prevents Eve from preserving the full state of those re-transmitted photons, so only 50% of them are sent with the same polarisation basis as originally chosen by Alice. Of the remaining photons, Alice and Bob randomly choose a small subset and compare their values over the classical channel. Any photons intercepted and re-transmitted by Eve have a 25% chance of giving the wrong value when measured by Bob, so comparing values in this way gives a statistical estimate of approximately how many photons might have been intercepted by Eve. The final cryptographic key is derived from the remaining photons whose values have not been disclosed over classical channels.

The effectiveness of the entire protocol relies upon the fact that Eve cannot measure the full quantum state of any given photon and cannot perform any measurements without the risk of introducing detectable errors. This in turn relies upon each binary bit being encoded into just one photon. If there were two photons carrying the same polarisation information then Eve could capture one while allowing the other to continue undisturbed to Bob, thus remaining undetected. After Bob revealed his measurement basis, Eve could then measure the stored photons in the same basis and thus obtain full knowledge of the secret key. This kind of attack is known as a photon number splitting attack [170]. Fig. 5.3 shows a diagram of this kind of attack. Additionally, photons with the four different polarisations must be indistinguishable. This means that they must have the same frequency and be travelling in the same optical mode.

There have been many successful demonstrations of practical QKD implementations [171, 172], but so far none of them have used true single photon sources. Instead they have used coherent sources with Poissonian photon statistics attenuated such that $\bar{n} \ll 1$ and $P(n > 1) \approx 0$, however this does not give the same guarantee of secrecy that a true single photon source would provide. There have been a number of proposals of ways to mitigate photon number splitting attacks so it is possible to build a secure QKD system using an attenuated coherent source [173, 174, 175, 176], but the fact that the majority of pulses must contain zero photons limits the key generation rate in such designs. If an on-demand single photon source were to be developed it would make the implementation simpler and more efficient [177].

Other Quantum Information Applications

Single photon sources also have applications in various other quantum information architectures. For instance, it has been proven theoretically that it is possible to build an all-optical quantum computer using only linear optical components and single photon sources [178]. This scheme uses the photons as flying qubits and encodes information into the photon number states, while the linear optical components can be arranged into quantum logic gates which operate on the qubits via interference effects.

5.2.2 Cavity Quantum Electrodynamics

Cavity Quantum Electrodynamics (cavity QED) is the study of the quantum interaction between light and matter in an optical cavity. It is this light-matter interaction which gives rise to many of the most useful properties of optical microcavities such as the ones developed in this chapter. The mathematical analysis in this section is largely based on that by Mark Fox [179], and a more in depth derivation of all the relevant formulae can be found in that book.

To begin with we will consider a very simple, idealised system consisting of a two-level emitter such as an atom or quantum dot resting inside an optical cavity. A diagram of this system is shown in Fig. 5.4. There is some coupling strength,




Figure 5.4: A diagram showing the basic model of a two-level emitter coupled to an optical cavity. g_0 is the coupling strength, κ is the rate of energy loss from the cavity mode, and γ is the rate of photon emission into modes other than the cavity mode.

g_0 , which can be thought of as the rate of energy transfer between the electronic states in the emitter and the optical modes in the cavity. The cavity has a finite loss rate, κ , and similarly the emitter loses energy into unconfined modes with rate γ . All three parameters have units of frequency.

The coupling strength is given by the following equation [179],

$$g_0 = \sqrt{\frac{\mu_{12}^2 \omega}{2\epsilon_0 \hbar V_0}}, \quad (5.2)$$

where μ_{12} is the electric dipole matrix element for the transition and V_0 is the volume of the cavity mode. The cavity loss rate depends on the Q factor and is defined as,

$$\kappa = \frac{\omega}{Q}. \quad (5.3)$$

The emitter loss rate, γ , will be ignored for the moment, since the main concern in this chapter is the cavity design.

The behaviours of emitter-cavity systems fall into two different regimes depending on the relative values of the coupling strength and loss rates [179]. When $g_0 \ll \max(\kappa, \gamma)$, the system is said to be in the weak coupling regime and the lifetime of the source is simply reduced by the Purcell effect. In contrast, when $g_0 \gg \max(\kappa, \gamma)$, the system enters the strong coupling regime and the photonic and excitonic modes hybridise to form quasi-particle states called polaritons.

By combining (5.2) and (5.3) we can rewrite the condition for strong coupling as,

$$\frac{Q}{\sqrt{V_0}} \gg \sqrt{\frac{2\omega\epsilon_0\hbar}{\mu_{12}^2}}. \quad (5.4)$$

Therefore, to achieve strong coupling it is desirable to maximise this ratio of $Q/\sqrt{V_0}$. Interestingly, we will see that the optimum trade-off between Q factor and mode volume is different for the Purcell effect than for strong coupling.

The Weak Coupling Regime

In the weak coupling regime the majority of photons produced by the emitter escape from the cavity before they have time to be re-absorbed. Therefore a weakly coupled system can be closely approximated by assuming that the cavity merely alters the density of states into which photons are emitted. Fermi's golden rule gives the rate of spontaneous emission,

$$W \equiv \frac{1}{\tau} = \frac{2\pi}{\hbar^2} |M_{12}|^2 \rho(\omega), \quad (5.5)$$

as a function of the density of final states, $\rho(\omega)$, and the transition matrix element, M_{12} . Here τ is the radiative lifetime.

From this it is possible to derive equations for the spontaneous emission rate in free space [179],

$$W_{\text{free}} = \frac{\mu_{12}^2 \omega^3}{3\pi\epsilon_0 \hbar c^3}, \quad (5.6)$$

and inside the cavity, assuming that the cavity has a Lorentzian lineshape and is exactly on resonance with the emitter,

$$W_{\text{cav}} = \frac{2Q\mu_{12}^2}{\epsilon_0 \hbar V_0}. \quad (5.7)$$

The Purcell factor is a measure of the relative increase in spontaneous emission rate due to the presence of the cavity and is defined as the ratio of (5.7) and (5.6),

$$F_P \equiv \frac{W_{\text{cav}}}{W_{\text{free}}} = \frac{6\pi c^3 Q}{\omega^3 V_0}. \quad (5.8)$$

The Purcell factor is more usually written in the following form,

$$F_P = \frac{3Q(\lambda/n)^3}{4\pi^2 V_0}. \quad (5.9)$$

where λ is the free space wavelength and n is the refractive index. Likewise, the mode volume is usually given in units of $(\lambda/n)^3$, simplifying the calculations.

A large Purcell factor is particularly desirable for enhancing the emission of single photon sources. In the case of an emitter which acts as a single photon turnstile device, such as a quantum dot, the rate of single photon production is ultimately limited by the spontaneous emission rate. Thus increasing the Purcell factor leads to brighter single photon sources which could then enhance the performance of any technologies based upon them. For instance, the maximum rate of data transfer in a QKD system is limited by the rate at which single photons can be generated [162]. The decreased lifetime also helps to reduce the timing jitter caused by the random nature of spontaneous emission.

Since the photons are preferentially emitted into the cavity mode, it is also possible to engineer the cavity design to produce directed beams of single photons [180]. Having more photons emitted into a single mode improves their indistinguishability, which is very useful for all kinds of quantum information applications.

It is apparent from (5.9) that the Purcell factor is proportional to Q/V_0 . Hence it is this ratio which is typically used as a figure of merit when evaluating cavities for use in the weak coupling regime.

The Strong Coupling Regime

In the strong coupling regime the rate at which energy is lost from the system is much less than the rate of interaction between the cavity and the emitter. This means that a single quantum of energy may be emitted as a photon and re-absorbed to create an exciton many times before it eventually escapes. This reversible transfer of energy gives rise to fundamentally different behaviour than the irreversible emission process in the weak coupling regime.




Figure 5.5: An energy level diagram of the dressed states in a strongly coupled emitter-cavity system.

The full mathematical treatment of strongly coupled light-matter interaction in a cavity is given by the Jaynes-Cummings model [181], however the details of this are beyond the scope of this thesis. Instead this section will briefly cover the most important features and applications of the strong coupling regime.

The ladder of possible states in the cavity can be described by wavefunctions of the form $|e, n\rangle$ and $|g, n\rangle$ where $|e\rangle$ and $|g\rangle$ represent the excited and ground states of the emitter respectively and $|n\rangle$ indicates that the cavity contains n photons. $|g, 0\rangle$ is the ground state of the system and has energy $\hbar\omega/2$ due to the vacuum field in the cavity. As the emitter is brought into resonance with the cavity the pairs of states $|e, n-1\rangle$ and $|g, n\rangle$ become degenerate, both having energy equal to $(2n+1)\hbar\omega/2$. In a coupled system this degeneracy violates the Pauli exclusion principle, so an energy level splitting is induced to lift the degeneracy. It can be shown that this splitting is equal to $\sqrt{n}\hbar g_0$, so more strongly coupled systems exhibit greater splitting. These mixed states are known as dressed states, and the corresponding quasi-particles are known as polaritons.

Experimentally this can be observed as an anti-crossing of the energies of the excitonic and photonic modes as they are brought into resonance. See, for instance,




Figure 5.6: A graph showing an idealised anti-crossing between the cavity (photon) and emitter (exciton) modes of a strongly coupled system. As the detuning approaches zero, the photonic and excitonic states merge to form two polaritonic states with an energy splitting which depends on the coupling strength.

Brossard *et al.*[182]. A diagram of this is shown in Fig. 5.6.

One of the most exciting applications for strong coupling in cavities is the creation of quantum logic gates [183] and quantum memories [184]. The strong coupling regime facilitates the highly sensitive non-linear interactions between photons necessary to implement quantum logic gates which can be controlled by the presence or absence of a single photon [185].

5.2.3 Advantages of Photolithography for Cavity Fabrication

By far the best way to create optical microcavities with a high Q factor and small mode volume is using photonic crystals. PhC structures operating at visible or near-infrared (NIR) wavelengths have features on the scale of 100 nm and tolerances on the order of 10 nm, meaning they are too small to be patterned using conventional photolithography processes. Therefore the standard method of fabricating these PhCs uses e-beam lithography.

In order to achieve coupling between a QD and a PhC cavity it is necessary to position the QD near to one of the antinodes of the cavity mode. Again this requires accuracy on the order of 100 nm or less. It is very difficult to grow good quality

QDs suitable for use as single photon sources in a site-controlled manner, so the QDs are grown in a random distribution over the wafer surface (see chapter 4.2.3). For 2D PhCs the cavity mode antinodes are always located near the middle of the semiconductor membrane, so the QDs are capped with additional semiconductor material to position them at the right depth. Unfortunately when the QDs are buried beneath 100 nm of semiconductor crystal, it's almost impossible to locate them using a conventional e-beam system. Some groups have built specialised hybrid e-beam and cathodoluminescence systems to locate buried QDs optically [186], however such a systems are prohibitively expensive in most cases. On the other hand, it is relatively easy and inexpensive to add photolithography capabilities to an existing μ PL system.

In prior work [187], Lee *et al.* have used photolithography to create tungsten registration marks a known distance from QDs, allowing PhC cavities to then be fabricated in the right location using e-beam lithography. In this chapter I present an alternative approach in which cavities are created by patterning photoresist on top of existing PhC waveguides. Since the location of each cavity is defined via photolithography, it would be possible to achieve good spatial overlap between the cavity mode and a QD located somewhere in the waveguide.

5.3 Device Design

One fairly common design of optical microcavity involves taking a PhC waveguide and modifying the PhC structure along a short length of it to create a heterostructure [19, 20, 21]. The modification can take many forms (altered hole size, altered lattice spacing, etc.) but the result is always to shift the frequency of the waveguide mode within the affected region. This then means that certain frequencies of light are allowed to propagate within a short section of the waveguide but are reflected from the sides and ends by the photonic band gap, thus forming an optical cavity.

The cavity design studied here falls into the same category of waveguide heterostructure cavities and operates by much the same principle. The PhC consists of a 200 nm thick suspended membrane of GaAs with air (or vacuum) above and




Figure 5.7: 3D render of the PhC cavity structure simulated in chapter 2 and fabricated in this chapter. The PhC membrane is made of GaAs while the blue disk consists of SU-8 photoresist. Light is confined in the z direction by total internal reflection and in the x and y directions by the PhC band gap.

below it and a lattice of holes etched through it. A missing row of holes in the middle forms the waveguide itself, and on top of that is deposited a disk of SU-8 photoresist with a diameter of roughly $1\text{ }\mu\text{m}$ and a thickness of around 100 nm . The presence of the SU-8 changes the effective refractive index experienced by light inside the waveguide, and hence alters the frequency of the guided modes. Since the SU-8 only covers a small section of the waveguide, it creates a localised cavity mode with a slightly different frequency to the guided mode. A more detailed theoretical treatment of PhCs and simulations of this cavity design are covered in chapter 2.

Fig. 5.7 shows a 3D render of the cavity design. A step-by-step description of the SU-8 fabrication process is given in chapter 4.

5.4 Optimising SU-8 Photolithography

Before attempting to create any cavities it was necessary to first develop procedures for working with SU-8. Simulations performed by our collaborator, Frederic Brossard, indicated that the highest Q factors could be achieved by making the SU-8 disks as thin as possible. A thickness of around 50 nm was agreed upon as our target for optimal device performance. Unfortunately almost all of the literature covering

SU-8 fabrication deals with much thicker layers (typically $> 1\text{ }\mu\text{m}$), so we could not simply copy an existing spin-coating or exposure recipe. As discussed later in this section, additional complexity was introduced by attempting to spin-coat SU-8 on top of PhC membranes.

5.4.1 Thickness Optimisation on Smooth GaAs

As described in chapter 4, the best way to cover a sample in a thin, uniform layer of SU-8 is through spin-coating. The thickness of the resulting layer is determined by various factors including the viscosity, concentration, and volatility of the solution, and the speed at which it is spun. The spin-coater used for this work could run at speeds up to 7000 rpm with adjustable acceleration. A solution of SU-8 2000.1 was chosen since that was the most dilute commercially available preparation, giving a nominal thickness of $0.1\text{ }\mu\text{m}$ when spun at 3000 rpm. Because of the significant cost (in both time and money) of creating PhC cavities via e-beam lithography, the initial samples used for testing the spin-coating process were unadorned GaAs.

The GaAs wafer was cleaved into a number of square samples, each approximately $8\text{ mm} \times 8\text{ mm}$ in size. These samples were cleaned, first by immersing them in acetone, then by rinsing with isopropanol (IPA) to remove the acetone residue. The samples were allowed to dry before being transferred to the spin coater. Each sample in turn had $20\text{ }\mu\text{L}$ of SU-8 2000.1 solution deposited onto it using a micro-pipette and was then immediately spun for 60 s with an initial acceleration of 500 rpm s^{-1} and a final speed in the range from 3000 rpm to 7000 rpm. After all the samples had been spin-coated in this way they were transferred to a hotplate at 95°C for 5 minutes for the pre-exposure bake step.

A Karl Suss MJB-3 mask-aligner with a 200 W UV lamp was used to expose the samples. Each sample was half covered with aluminium foil, leaving the other half exposed to the UV light. The exposure lasted 30 s, then the samples were transferred back to the hotplate for 5 minutes at 95°C to accelerate the cross-linking process. After the post-exposure bake was complete the samples were developed by immersing them in propylene glycol methyl ether acetate (PGMEA) for 10 seconds




Figure 5.8: A photograph of one of the samples used for thickness testing. The left half is bare GaAs substrate, while the right half is covered in a layer of exposed SU-8 photoresist. The sample is approximately $8\text{ mm} \times 8\text{ mm}$ in size.

with gentle agitation, then they were rinsed with IPA and allowed to dry. This was sufficient to remove all of the unexposed SU-8 without affecting the exposed half. An photograph of one of the samples is shown in Fig. 5.8.

To measure the thickness of the SU-8 layer a Veeco Dektak 150 surface profilometer with a $12.5\text{ }\mu\text{m}$ stylus was used. This had a relatively poor lateral resolution, however the vertical resolution was on the order of 1 nm making it suitable for this task. The stylus force was set to its lowest value of 1 mg to avoid deforming the SU-8. The scans were performed over a length of $300\text{ }\mu\text{m}$ perpendicular to the masked edge, starting on the bare GaAs side and ending on the side covered in SU-8. Each sample was scanned in 5 different locations to account for thickness variations, excluding a 0.5 mm border around the sample edges where significant non-uniformity was present. The surface baseline was determined from the profile on the bare GaAs side, then the SU-8 thickness could be measured by comparing the two sides.

The thickness measurements from the first run with 500 rpm s^{-1} acceleration are plotted in Fig. 5.9. The most interesting thing to note is that there is no significant trend in the data. Spinning at 7000 rpm appears to give the same thickness as spinning at 3000 rpm . This can be explained by considering the spin speed as a function of time. At an acceleration of 500 rpm s^{-1} it takes 6 seconds to




Figure 5.9: A plot of the thickness of the SU-8 layer produced by spinning SU-8 2000.1 on smooth GaAs at various speeds with an initial acceleration of 500 rpm/s. The error bars show the standard deviation of the 5 measurements taken on each sample.




Figure 5.10: A plot of the thickness of the SU-8 layer produced by spinning SU-8 2000.1 on smooth GaAs at various speeds with an initial acceleration of 5000 rpm/s. The error bars show the standard deviation of the 5 measurements taken on each sample.

reach 3000 rpm and 14 seconds to reach 7000 rpm. When spin-coating, the thickness of the layer stops decreasing once enough solvent evaporates to make the viscosity overcome the centrifugal force. The thinner the layer, the greater the surface area to volume ratio and hence the faster the viscosity increases. The results suggest that this happened within the first 6 seconds, so further acceleration had little or no effect.

To investigate this, the same procedure as above was repeated with fresh samples and with an acceleration of 5000 rpm s^{-1} . The results of this second run are plotted in 5.10. Here we can see a much more obvious trend of decreasing

thickness as spin speed increases. The thickness at 3000 rpm was still higher than the nominal thickness for SU-8 2000.1, but this could be because the solution had been stored for more than a year when these data were taken, allowing some of the solvent to escape. It is also interesting to note that the variation in thickness across the surface, indicated by the size of the error bars on the graph, decreased with increasing spin speed.

5.4.2 Two-Photon Exposure and μ PL Apparatus

It was decided that a two-photon exposure process would initially be used for patterning the photoresist. SU-8 is normally only sensitive to light with a wavelength of around 400 nm, but with sufficiently high energy density it is possible to trigger the same photochemical processes by the absorption of two 800 nm photons. The main advantage of this becomes apparent when we consider the combined μ PL and exposure process.

The band gap of GaAs ranges from 1.52 eV at 0 K to 1.43 eV at 300 K [188], meaning that 800 nm photons have sufficient energy to excite the substrate (and hence the QDs) via a one-photon interaction. This means that as the power is decreased, the rate of SU-8 exposure decreases with the power squared while the PL intensity only decreases linearly with power. Therefore it should be possible to find a power at which μ PL measurements are possible using 800 nm light without exposing the SU-8. By using the same laser both to locate the QDs and to expose the SU-8 it would eliminate a possible source of alignment errors.

Another advantage lies in the limitations of optical components. The QDs were grown to emit at a wavelength of around 1.3 μ m to match the telecoms O-band – a region of the spectrum with particularly low losses in optical fibres – for QKD applications. This meant that the μ PL system would need to be able to collect and collimate light at this wavelength. The same optics would need to be used for both μ PL and exposure, so it would need to be able to focus light at both 1.3 μ m and either 800 nm or 400 nm. It is quite easy to acquire objective lenses which have similar performance from 800 nm to 1.3 μ m, but the difference between 400 nm




Figure 5.11: A diagram of the optical apparatus used both to perform μ PL measurements at cryogenic temperatures (using the HeNe laser) and to expose SU-8 photoresist via a two-photon process (using the Ti:Sapphire laser). The $4f$ system, while present, was not used for the work in this chapter.

and $1.3\,\mu\text{m}$ is so great that there is inevitably some difference in the position of the focal plane. While it is possible to compensate for this (as described later in this chapter) it adds extra complications.

The 800 nm light was provided by a Ti:Sapphire laser producing pulses of 100 fs duration and with a repetition rate of 80 MHz. The same experimental system was also used for more conventional μ PL measurements so it was desirable to have a continuous wave (CW) excitation source – a 633 nm HeNe laser. In order to use both at once it was necessary to have the beams from the two lasers travelling along the same path so that they were focused to the same position on the sample – the same position from which light was collected into the spectrometer. Fortunately both lasers emitted linearly polarised light, making it relatively simple to combine them into the same beam path using a polarising beam-splitter. A half-wave plate was positioned in the path of each laser and then rotated until the transmitted (or reflected) power was maximised. A full diagram of the system is shown in Fig. 5.11.

Additionally, a pair of lenses and a pinhole were added to the Ti:Sapphire beam path in order to give the beam a clean Gaussian profile. The first lens focussed the collimated beam on the pinhole which acted as a spatial filter, then the second lens re-collimated it into a beam. By selecting lenses with different focal lengths it

was possible to change the beam diameter to match the size of the aperture on the objective lens, thus achieving the smallest possible spot size on the sample.

Coarse positioning and focus on the sample were achieved by way of three manual micrometers: two to move the stage holding the sample and cryostat horizontally and a third to move the objective lens in the vertical direction. These micrometers were only accurate to approximately $10\text{ }\mu\text{m}$, so an XYZ piezo-actuated stage was used for nanometre-scale control of the objective position. This piezo stage was controlled by computer software which allowed for the automation of both μPL maps and SU-8 exposure routines. A computer controlled shutter was also added to the Ti:Sapphire beam path to allow for precise timing of exposures. When combined, these systems allowed for the automatic writing of almost any pattern which fit within a $100\text{ }\mu\text{m} \times 100\text{ }\mu\text{m}$ bounding box. The exposure power could be adjusted by way of a neutral density filter wheel with a continuous gradient, although unfortunately that part was never automated.

A 50:50 beam splitter was used to reflect the lasers down into the microscope objective while still allowing half of the collected light to pass up into the spectrometer. This had a helpful side effect of providing a convenient place to measure the laser power. By placing an optical power meter in the path of the transmitted laser light it was possible to measure the power without blocking either the excitation or collection paths. This also improved repeatability, since the power meter was not moved between measurements. A simple linear factor could be used to convert the measured power to the power reaching the sample.

5.4.3 Cryogenic Exposure and Dose Optimisation

Since the eventual goal was to perform nearly simultaneous μPL measurements of QDs and patterning of SU-8, we needed to test the ability of the system to expose SU-8 at very low temperatures. InGaAs QDs are highly sensitive to phonons, so the optimal temperature for μPL measurements is around 4 K. As the temperature is increased the PL intensity decreases and spectrally broadens, and by 50 K it is




Figure 5.12: (a) A schematic of the dose plot which was patterned into the SU-8 using a Ti:Sapphire laser at a temperature of 4 K inside the cryostat. The red dots indicate the location of exposures performed with specific laser power and duration settings. The powers and durations are shown numerically as axis labels. The surrounding box, the large arrows, and the ‘X’ were all written onto the sample to serve as position and orientation markers. (b) An AFM image of the plot after it was patterned and developed. The colour scale indicates the surface height at each point.

no longer possible to identify individual QDs. To reach the temperatures required, the sample was placed in a Janis continuous flow liquid helium cryostat.

To achieve adequate thermal contact between the sample and the cryostat’s cold-finger it was necessary to use an adhesive. Initially we tested an adhesive containing silver nanoparticles to increase the thermal conductivity, however it was found that during the SU-8 development step the silver particles tended to migrate to the top surface of the sample, severely contaminating it. In all subsequent experiments we switched to using Copydex, a water-based latex rubber cement. This had sufficient thermal conductivity to allow for μPL measurements while leaving minimal residue after development.

In order to determine the dose required to expose the SU-8 at a temperature of 4 K, an array of spots was written on SU-8 coated GaAs using different laser powers and durations. Various alignment markers were also written (at a speed of $0.2 \mu\text{m s}^{-1}$ and a power of 51 mW) so that it was possible to find and orient the array of spots under a microscope later. Fig. 5.12(a) shows a schematic of the




Figure 5.13: SEM images showing the design of the PhC structure used both for initial testing and for fabricating cavities using SU-8. The L3 cavities were included to act as alignment markers, allowing the centre line of the waveguide to be accurately located via μ PL measurements.

pattern, while Fig. 5.12(b) shows a height map of the sample surface recorded using an atomic force microscope (AFM) after patterning. The powers listed next to the schematic correspond to the laser power reaching the sample surface. It is worth noting that the time axis is approximately logarithmic while the power axis is linear.

We can see from these results that the power requirements were not insignificant. The lowest power which produced any kind of visible feature was 20 mW, and then only at the longest duration exposure. This corresponded to a pulse energy of 0.25 nJ and an instantaneous optical power density of at least $2.5 \times 10^{15} \text{ W m}^{-2}$ (this is a lower bound, since we don't know the time-domain shape of the laser pulses). This is not surprising, since two-photon interactions scale as energy density squared and typically only manifest at much higher powers than their one-photon equivalents. The target thickness of 50 nm was only achieved with a significantly higher dose of 40 mW for 20 s.

5.4.4 Spin-coating on PhC membranes

One of the challenges of patterning SU-8 on PhC waveguides was positioning. It was possible to achieve coarse alignment by viewing the camera image of the sample




Figure 5.14: An AFM image showing the fluctuations in surface height around PhC membranes coated in SU-8 photoresist.

surface and adjusting the micrometers and piezo-actuators until the laser spot was centred over the waveguide, however this only gave accuracy on the order of $1\text{ }\mu\text{m}$. To increase the accuracy, registration markers were added to the PhC design during the e-beam lithography stage. These consisted of a pair of smaller PhC structures containing L3 cavities in line with the waveguide but separated from it by $> 10\text{ }\mu\text{m}$. A scanning electron microscope (SEM) image of these structures is shown in Fig. 5.13. The L3 cavities had very small mode volumes and hence provided very localised emission peaks when observed with the μPL system. By moving the piezo-actuators in very small increments (typically $0.2\text{ }\mu\text{m}$) it was possible to locate the centres of these L3 cavities with very high precision. Once the L3 positions were known, any other locations on the PhC structure could be calculated relative to those.

The other main challenge presented by the PhC membranes was achieving the correct SU-8 thickness. Some property of the surface in the vicinity of the PhCs caused the SU-8 to avoid covering them during spin-coating. Examples of this are shown in Fig. 5.14 and Fig. 5.15(a). While the flat parts of the sample were covered in a relatively uniform layer of the desired thickness, the PhC membranes were left uncoated.

One possible explanation for this is the presence of the PhC holes themselves. The tendency for a liquid to ‘wet’ a surface can be characterised by the contact angle which naturally forms between a droplet of liquid and the surface. A small contact angle (and hence a large contact area) indicates that a surface is highly wettable (termed hydrophilic, in the case of water), while a large contact angle (and




Figure 5.15: Optical images showing SU-8 layers of various thicknesses in the vicinity of PhC membranes. The colours are real, as observed by the camera attached to the microscope, and are caused by thin-film interference effects from the SU-8. More closely spaced bands of colour indicate larger gradients in the SU-8 thickness. (a) shows a very thin layer of SU-8 which failed to adhere to area around the PhC membrane. (b) shows a thicker layer of SU-8 which successfully covered the PhC but with reduced thickness. (c) shows a very thick layer of SU-8 which uniformly covered the PhC without any visible reduction in thickness.

small contact area) corresponds to a surface with low wettability (hydrophobic). The contact angle for smooth surfaces is generally a property of the solid-liquid pair and determined by their individual surface chemistries, however sub-millimetre scale physical features on the surface can drastically alter the interaction. There are numerous examples, both artificial and natural, of nano-structured surfaces which exhibit superhydrophobic behaviour [189, 190, 191, 192]. This structural hydrophobicity occurs when the liquid cannot fully penetrate the nano-scale features on the surface, leaving pockets of trapped air. The surface seen by the liquid can then be thought of as a composite formed of a mixture of solid and air, and the contact angle, θ_c , is given by the Cassie-Baxter equation [193],

$$\cos \theta_c = f(\cos \theta_s + 1) - 1,$$

where f is the solid fraction of the contact area and θ_s is the contact angle formed between the liquid and a smooth surface of the same solid. As f approaches zero, θ_c approaches 180° and the force holding the droplet onto the surface vanishes. For the PhCs studied here f is no less than 0.74, however this may still be sufficient to significantly affect the balance of forces during spin coating.

The other possible explanation for the poor SU-8 adhesion around the PhCs is some kind of contamination left by the etching process used to cut the holes and dissolve the sacrificial AlGaAs layer. We attempted to mitigate this by placing the sample in an oxygen plasma asher to clean the surface, however no improvement was observed.

During testing it was observed that nominally much thicker layers of SU-8 (> 200 nm) did successfully cover the PhCs, as seen in Fig. 5.15(b) and (c). In order to have finer control over the SU-8 thickness a new recipe was developed based on diluting SU-8 2007 (nominal thickness $7\text{ }\mu\text{m}$ at 3000 rpm) with cyclopentanone. Cyclopentanone was chosen because it is the same solvent used in the commercial SU-8 2000-series preparations. Various dilutions and spin recipes were tested with volume ratios typically around 1 part SU-8 2007 to 7 parts cyclopentanone and spin speeds around 3000 rpm. In order to create the edges required for accurate thickness measurements in well-defined locations the μPL system was used to expose disks and lines on the PhCs. The thickness measurements themselves were made using an AFM instead of the DekTak due to the high positioning accuracy required. Due to the significant time required for each measurement and the fact that we only had a very small number of PhC samples to work with, this optimisation was an ongoing process performed in parallel with other exposure tests and measurements.

5.4.5 Problems With Two-Photon Exposure

When it came time to test the two-photon exposure procedure on PhC membranes, two serious problems became apparent. First, in an attempt to save time and helium, initial exposures were tested while the sample was at room temperature. The Ti:Sapphire laser was attenuated to a power of 20 mW and a series of 10 s duration exposures were performed on 16 PhC waveguides. The devices were then examined under an SEM. More than half of the devices were found to have suffered significant damage from the laser, many having large holes melted in the membrane, rendering them completely unusable as PhCs. An example of this is shown in Fig. 5.16.




Figure 5.16: An SEM image showing the damage caused to a PhC waveguide by the 100 fs pulsed Ti:Sapphire laser during an attempted SU-8 exposure at room temperature.

The GaAs comprising the PhC melted (or possibly was vaporised) because of the extremely high power density combined with poor heat dissipation. Since the PhC membranes are suspended over an air gap and perforated with holes, there are far fewer paths through which heat can escape from hot spots. To try and prevent this damage, further tests were performed at a temperature of 4 K.

When exposing using similar power at 4 K no immediate damage was visible. To test for invisible damage, μ PL measurements were required. Notomi cavities in PhC waveguides were chosen as test samples for this purpose since they most closely resembled the cavity design we were trying to create. The exposure conditions were set at 25 mW for 20 s, which was towards the lower end of useful exposure powers measured in Fig. 5.12. No SU-8 was present on these samples during exposure.

The spectrometer was configured to use the grating with the highest spectral resolution available, since peaks with high Q factors have correspondingly narrow linewidths. The pixel-to-pixel resolution at the wavelength of the peak was 0.03 nm. Excitation was provided by the 633 nm HeNe laser at a power of $30 \mu\text{W}$. Once the Notomi cavity was located coarsely the piezo-actuators were used to adjust the position in small increments until the PL intensity was maximised. The PL spectrum was recorded, as well as a background spectrum with the laser blocked




Figure 5.17: Plots of the μ PL spectra recorded from the same Notomi cavity before (blue) and after (orange) being exposed with the Ti:Sapphire laser at a power of 20 mW for 20 s. The corresponding Q factors are listed in the legend.

to allow the background noise and dark count to be subtracted. The exact same μ PL measurement procedure was followed both before and after the exposure to ensure that the results were comparable. The same experiment was performed on 6 different Notomi cavities.

Fig. 5.17 shows an example of the background-corrected spectra recorded before and after exposure. The Q factors were calculated by an algorithm which performed automatic peak detection followed by fitting a Lorentzian function to the data. The Lorentzian parameters were initially estimated from the height and width of each detected peak, then a least-squares minimisation algorithm was used to refine the initial estimate. A linear offset was included in the fitting function to account for the emission from the QD ensemble. The Q factor was then given by the equation, $Q = \lambda/\Delta\lambda$, where λ is the centre wavelength and $\Delta\lambda$ is the full width at half maximum (FWHM).

In all cases there was a very significant reduction in PL intensity of around 75% and a reduction in Q factor of 20% to 30% after exposure. Interestingly there was no significant decrease in the PL intensity from the QD ensemble, suggesting that the QDs themselves were not being damaged. Rather the reduction in emission from the Notomi cavities was almost certainly caused by the reduction in Q factor leading to reduced Purcell enhancement. The over all cause is most likely to be slight

distortions in the PhC structure created when the GaAs membrane approached its melting point. These were not visible under SEM, but were sufficient to reduce the performance of the PhCs.

To verify that this was not a sample-dependent anomaly, the same experiment was run on the L3 cavities adjacent to the waveguides we were using for SU-8 fabrication. Those showed an even more significant reduction in PL intensity, typically exceeding 90%.

It may have been possible to mitigate this by increasing the exposure time and decreasing the exposure power, however due to the non-linear nature of two-photon processes the exposure times would have needed to be increased enormously. To decrease the power by a factor of x , it would have been necessary to increase the exposure time by a factor of x^2 to maintain constant dose. Additionally, it's unlikely we could have performed room temperature exposures in any reasonable time without damaging the sample. For this reason we decided to switch to a one-photon exposure process using 405 nm CW light.

5.4.6 Blue Exposure

The 405 nm light was provided by a 10 mW diode laser, and the beam was made co-axial with the HeNe beam using a dichroic mirror which reflected short wavelengths while transmitting longer wavelengths. It was essential for the lasers to be exactly co-axial in order to accurately position the SU-8 spots, so fine adjustment was performed while viewing the camera image of the two laser spots on the sample. The NIR objective was replaced with a 100 $\times$, 0.5 NA objective optimised for visible wavelengths. As mentioned earlier, the focal planes of the HeNe and 405 nm diode lasers were offset by a few microns due to the difference in wavelengths, so a computer program was written to allow for rapid automatic switching between the two focal planes using the piezo-actuators. This allowed the laser spot to be positioned using the HeNe as a guide, then we could switch to the diode laser to perform the exposure itself.




Figure 5.18: An optical microscope image of the dose plot for exposure of SU-8 using 405 nm light at room temperature. The exposure powers and durations are listed next to each row and column respectively. In this case the SU-8 was roughly 250 nm thick.

In order to find the optimal exposure conditions for 405 nm light, another dose plot was written onto a flat region on the surface of an SU-8 coated sample. This is shown in Fig. 5.18. As expected, the power required for exposure was much lower than the two-photon process. The spot diameters ranged from roughly 0.6 μm up to 2 μm with increasing power and duration. A dose of around 3 μW for 5 s was chosen for further experiments since that produced a spot very close to the target diameter of 1 μm .

5.4.7 White Light Inteferometry

Unfortunately, even when spinning with exactly the same conditions, small variations in average SU-8 thickness were amplified over the PhCs producing large variations in local thickness. The same sample could have some PhCs with no SU-8 on while others had a layer as thick as 0.5 μm . We wanted to be able to target devices with the correct SU-8 thickness, but it was not possible to use the AFM to measure the thickness of the layer before exposure since it covered the entire sample. Interferometry presented a possible alternative which would allow the

thickness of a continuous layer to be measured. An optical thickness measurement system was particularly appealing since it could be integrated into the existing μ PL system and used for rapid, in situ measurements.

Illumination was provided by a white LED, the spectrum of which is shown in chapter 3 Fig. 3.9. The light from the LED was collimated by a lens, and to avoid exposing the SU-8 a 450 nm long-pass filter was added to block the high energy photons. The light was then focussed through the objective, reflected off the sample surface, and collected by the spectrometer. In order to calculate the reflectivity a reference spectrum was taken with an aluminium mirror in place of the sample. A background spectrum was also recorded with the LED turned off. For initial testing a $20\times$ magnification objective was used, producing a spot size of $5\text{ }\mu\text{m}$ diameter. This was done to ensure that the wavefront arriving at the sample surface was as flat as possible, and to average out very localised variations in thickness.

The reflection spectrum, R , could be calculated from the recorded sample (S), mirror (M), and background (B) spectra using the equation, $R = (S - B)/(M - B)$.

For SU-8 on smooth GaAs the wavelengths of the maxima and minima in the spectrum could easily be calculated analytically. For instance, for SU-8 of thickness t and refractive index n , the maxima would fall at wavelengths of $\lambda_{\text{max}} = 2tn/m$, where m is any non-zero positive integer. However, in this case the SU-8 was on top of PhC structures which were themselves suspended above the substrate by air gaps. It was unclear exactly how much this would affect the results, so a number of electromagnetic FDTD simulations were performed using the proprietary software Lumerical.

The 3D FDTD simulations consisted of a plane wave source positioned $2\text{ }\mu\text{m}$ above the top surface of the GaAs and a detector plane positioned $0.5\text{ }\mu\text{m}$ above that. The source was uni-directional, producing a broad plane wave spanning the wavelength range from 400 nm to 900 nm in the downward direction only. In one simulation the structure was simply a flat substrate of GaAs covered in SU-8, while in the other two it consisted of a PhC waveguide with the same parameters as in chapter 2 suspended $1.4\text{ }\mu\text{m}$ above the solid GaAs substrate. The simulation

volume was roughly $1\text{ }\mu\text{m} \times 1\text{ }\mu\text{m} \times 5\text{ }\mu\text{m}$ in size and was surrounded on all sides by PML boundaries. A series of simulations were run on each structure, varying the SU-8 thickness and recording the reflection spectrum from the flux through the detector plane. Of the two simulation sets run on the PhC, one was performed over the waveguide while the other was performed over another part of the PhC. The reflection spectra were corrected by dividing by the spectrum of the source, then normalised into the range $[0, 1]$ for comparison with the experimental data.

The normalised spectra calculated for various SU-8 thicknesses on the smooth GaAs are shown in Fig. 5.19(a). While the positions of the peaks and valleys in the simulated spectra generally matched the corresponding features in the experimental results, there was significant variation in amplitude which made fitting difficult. Additionally it was not simple to determine the locations of the peaks and valleys from the experimental data since there were secondary fluctuations over many different scales. To smooth out the noise in the experimental data a discrete Fourier transform was performed, the high frequency components were discarded, then it was transformed back into real space. The amplitude discrepancies still made conventional fitting approaches difficult, so instead both the simulated and real data was transformed into the set of discrete values $-1, 0, 1$ by taking the sign of the gradient at each point. Thus the amplitude information was discarded and the only information remaining was whether the gradient was positive, zero, or negative at each point. The fitting was performed by iterating through the set of simulated spectra, taking the inner product of each spectrum with the experimental data, and finding the thickness at which the inner product was maximised. A successful example of this fitting is shown in Fig. 5.19(b). The data in that figure has been normalised, but no smoothing or gradient quantisation has taken place.

To evaluate this technique a set of 16 devices were measured, then they were exposed using the 405 nm laser and the thicknesses of the resulting spots were measured using an AFM. These thickness measurements were only accurate to around $\pm 10\text{ nm}$ due to variations across the top surface of each exposed spot.




Figure 5.19: (a) A waterfall plot of the simulated reflection spectra for a large range of different SU-8 thicknesses. (b) An example of the measured reflection spectrum (orange) from 400 nm thick SU-8, and the simulated spectrum (blue) which was automatically identified as the best fit. All the spectra in this figure are normalised by dividing by their maxima.

Of the three simulated datasets, the off-waveguide PhC simulations were found to produce the best fitting results with a root mean squared (RMS) error of 92 nm. A plot of the correlation between the interferometry and AFM results is shown in Fig. 5.20. Although this error was large, it still provided a rough estimate of the SU-8 thickness which would be useful for device fabrication.

We noted that when viewed under an optical microscope there were significant colour variations across the SU-8 layer. Out of curiosity, a free and open source library, `python-colormath`, was used to convert the simulated reflection spectra into the corresponding colours as seen by the human eye. An exact conversion was not possible without knowing the spectrum of the LED illumination in the microscope, so the standard CIE illuminant D65 [194] (corresponding to daylight with a colour temperature of 6504 K) was assumed. A plot of the colour as a function of thickness is shown in Fig. 5.21. This provided an even more convenient way to estimate the SU-8 thickness at the cost of some accuracy.




Figure 5.20: A plot showing the correlation between thickness estimates made using interferometry and the real thickness of the resulting SU-8 disks measured by AFM. The red line shows where the points should lie if the interferometry and AFM measurements agreed exactly.




Figure 5.21: A plot showing the apparent colour of the SU-8 as a function of thickness, based on simulations of the reflection spectrum from PhC membranes coated in SU-8. Conversion from spectra to RGB colour space was done using the python-colormath library.

5.5 Cavity Fabrication and Optical Characterisation

5.5.1 Experimental Procedure

Having developed a reasonably reliable set of procedures and techniques for patterning SU-8 on PhCs, we were ready to begin fabricating optical cavities. A new sample was used, since the previous ones had become quite contaminated from repeated processing. This sample had embedded InGaAs QDs with a density of approximately $100 \mu\text{m}^{-2}$ and the PhC design was scaled to match the wavelength of the low-energy tail of the QD distribution at around 1260 nm. The PhCs were arranged in a rectangular grid of 6 regions, each containing a grid of roughly 200 devices.

In order to be able to confirm the creation of optical cavities it was necessary to record the optical modes of the waveguides before adding the SU-8. As discussed in chapter 2, PhC waveguides of finite length support Fabry-Pérot modes at a frequency very close to the cavity mode. These Fabry-Pérot modes can have Q factors comparable to those of SU-8 cavity modes, making them difficult to tell apart. For these μ PL measurements the sample was placed inside the cryostat (attached to the cold-finger with Copydex glue) and cooled to a temperature of 4.2 K.

To ensure that the μ PL maps were exactly comparable, the L3 cavities offset from each end of the waveguide (see Fig. 5.13) were used as registration markers. The sample orientation was calculated from the vector between the PL maxima of the two cavities, and the midpoint of the waveguide was calculated from the position half way between them. Once the orientation of the sample was known, it was only necessary to measure the position of a single L3 cavity for each waveguide, making the procedure somewhat faster. Excitation for the μ PL measurements was provided by the 633 nm HeNe laser attenuated to a power of $6\mu\text{W}$ on the sample surface. Automated mapping software developed by the author was used both to control the XYZ piezo-actuator holding the objective lens and to read data from the spectrometer. The objective was moved along the length of the waveguide over a distance of $20\mu\text{m}$ in steps of $0.2\mu\text{m}$. After each movement the system waited for half a second before recording a spectrum to allow any vibrations from the movement to settle. The integration time was 0.5 s per point, and the lower resolution grating (600 lines/mm) was used giving a pixel-to-pixel spectral resolution of 0.1 nm at the cavity wavelength.

The SU-8 recipe used for the main cavity fabrication run was a dilution of 1 part SU-8 2007 to 8 parts cyclopentanone, spin-coated at a speed of 2800 rpm with an initial acceleration of 1000rpm s^{-1} . As mentioned earlier, there was significant thickness variation across the area of the sample. This was partly caused by edge effects due to the small size of the sample ($5\text{ mm} \times 5\text{ mm}$) and partly due to the non-uniform surface created by the grid of PhC structures. Optical

microscope images were used as a guide to find the devices with non-zero SU-8 coverage for later exposure.

Exposure on the candidate devices was performed at room temperature using the 405 nm diode laser at a power of $3.4\,\mu\text{W}$ for 5 s. Unfortunately the combination of the temperature and the SU-8 layer significantly degraded the performance of the L3 cavities, making it difficult to precisely determine their locations. Instead the camera image was used to manually position the HeNe laser spot over the centre of the waveguide. The diameter of the waveguide ($0.5\,\mu\text{m}$) was on the same order of magnitude as the resolution of the optical system, so it was possible to identify it as a slightly brighter line in the middle of the PhC. Once the HeNe laser was in position, another piece of custom software was used to automatically lower the objective to match the focal plane of the blue light, open the shutter for the desired duration, and then return the objective to its original position. This procedure was performed on a total of 40 devices in one section of the sample.

After development in PGMEA and rinsing with IPA, the sample was returned to the cryostat for the second set of μPL maps. The exact same procedure was followed as before, registering the locations of the waveguides using the L3 cavities and recording μPL spectra at the same points.

In addition to the μPL maps, point spectra were also recorded using a grating with a higher pixel-to-pixel resolution of $0.05\,\text{nm}$ in order to more accurately calculate the Q factors of the modes. For these measurements the laser spot was positioned over the centre of the cavity mode by manual optimisation to maximise the PL intensity. A longer integration time of 1 s was used to compensate for the reduction in brightness from using a higher resolution grating.

For comparison to the simulation results it was desirable to know the dimensions of the SU-8 disks. To this end the exposed devices were measured using an AFM. A custom script was developed in order to automatically extract estimates of the thickness and diameter of the SU-8 disks. The AFM introduced some low-frequency distortion to the data, so to correct for this a linear background was calculated from the height at the four corners of the scanned area (away from any




Figure 5.22: AFM images showing the surface profile of two successful PhC cavities. (a) is a typical device, while (b) is the cavity with the highest measured Q factor.

features) and subtracted. The individual rows of the raster scan also needed to be levelled by measuring the height along two edges of the area and altering the height of each row to have the same average height at those points.

The SU-8 disks were found to typically have a slightly elliptical shape (visible in Fig. 5.22), most likely due to slight astigmatism in the laser spot used to expose them, so the shape was approximated as an elliptical top-hat function. Unfortunately a true top-hat function with vertical walls performed extremely poorly when fitting was attempted because, when the input was a discrete grid of data points, the derivatives of all the parameters in the X-Y plane became delta functions. To avoid this problem the side walls of the function were made to slope, making it a conical frustrum. Least-squares minimisation was performed by the Levenberg-

Marquardt algorithm from the SciPy library. The radius was simply taken as the average radius of the ellipse, while the error on that value was calculated by taking the RMS difference between the position of the half-maximum points at various angles and the fitted radius. The error on the height was given by the RMS difference between the fitted height and the heights of all the points within the half-maximum region of the function.

5.5.2 Results

The results of μ PL measurements from 4 of the more successful devices are plotted in Fig. 5.23. Successful cavity creation is evident from looking at the plots on the left hand side of the figure. Fig. 5.23(a,c,e,g) show the μ PL maps recorded before the addition of SU-8 and (b,d,f,h) show the maps over the same devices afterwards. In each case a new bright mode has been created near the centre of the waveguide. The other modes visible in those plots are Fabry-Pérot modes which can be seen to shift slightly when the SU-8 disk is added. These maps allowed us to immediately see which modes were created by the SU-8 disks, ensuring that we didn't mistake a pre-existing Fabry-Pérot mode for a cavity mode.

The plots on the right, Fig. 5.23(i,j,k,l), show high resolution spectra recorded at the position of maximum PL intensity for the same devices. These data have already had the background (taken with the laser blocked) subtracted, so the sloped background visible in those plots is the low-energy tail of the QD ensemble. The Q factors of these cavities were found using the same automated fitting algorithm mentioned earlier in section 5.4.5. It is also noteworthy that the wavelengths of the cavity modes varied over a range of roughly 10 nm, although this was not found to correlate significantly with any physical parameters of the SU-8 disks. The wavelengths of the Fabry-Pérot modes followed the same variations, so it is likely caused by imperfections in the underlying PhC structures.

Fig. 5.24 shows an SEM image of the cavity which yielded the highest Q factor of 7.4×10^3 . The SU-8 actually appears slightly translucent in this image because the thickness of the SU-8 (70 nm) was comparable to the scattering length




Figure 5.23: Plots showing μ PL measurements of the cavities. Each row corresponds to a single device. (a-h) show the 1D μ PL maps recorded before (left) and after (right) cavity fabrication. (i-l) show high-resolution spectra recorded at the maxima of the cavity modes, with annotations listing the Q factors.




Figure 5.24: SEM image of the cavity with the highest quality factor, $Q = 7400$. The SU-8 photoresist was only 70 nm thick on this device.

of the electron beam. For comparison, Fig. 5.25(a) shows a similar SEM image of a lower Q cavity ($Q = 5.6 \times 10^3$) with thicker SU-8 (230 nm). In both cases there is a very small (≈ 100 nm) misalignment between the centre of the SU-8 disk and the centre line of the waveguide.

Fig. 5.25(b) shows a single slice through a 2D μ PL map at a wavelength of 1254.1 nm. This is the same device as the SEM image in (a) and covers the same area (to an accuracy of around $0.5 \mu\text{m}$ in each direction). The bright spot in the middle is the cavity mode which, in addition to being quite localised, can be seen to correspond well with the position of the SU-8 disk in the SEM image.

The spatial alignment of the two images in Fig. 5.25 was achieved by looking at the QD ensemble emission around the edges of the PhC. The QD density inside the area of the PhC is effectively lower because of the presence of the holes. Given a hole radius of $0.27a$, the PL intensity could be expected to drop by around 26% when inside the PhC. This dip in PL intensity is visible in the μ PL map and was used to determine the correct alignment of the two images. The μ PL map in (b) actually shows a slightly longer wavelength than the maximum intensity peak at 1253.7 nm, but was chosen because it demonstrates the contrast in PL emission between the PhC and the surrounding area more clearly.




Figure 5.25: (a) SEM image of one of the cavities with $Q = 5.6 \times 10^3$, SU-8 disk thickness of 230 nm, and disk diameter of 1.2 μm . (b) Plot of the PL emission at a single wavelength as a function of position on the same device. The step size for this map was 0.5 μm .




Figure 5.26: (a) Plot of cavity Q factor against SU-8 disk thickness. (b) Plot of SU-8 disk diameter against thickness.




Figure 5.27: Plots of the theoretical quality factor, Q , and mode volume, V , for cavities created using SU-8 disks of diameter $1\ \mu\text{m}$ and various thicknesses. These data are taken from simulations performed by our collaborator, Frederic Brossard.

A total of 40 devices were selected for patterning. Of those, 34 could be confirmed to have cavity modes after processing, giving a success rate of 85%. The remaining devices typically failed either due to a lack of adhesion between the SU-8 and the PhC surface or due to damage to the PhC membrane during processing. The distribution of Q factors and SU-8 disk thicknesses and diameters are plotted in Fig. 5.26. Looking at (a), there is a clear correlation between the thickness and the Q factor. For comparison, Fig. 5.27 shows theoretical results from FDTD simulations of the Q factor and mode volume as a function of SU-8 thickness. The same trend is visible in both, with the highest Q factors achieved by the thinnest SU-8 disks.

One of the main differences in the experimental results is the large spread of Q factors among devices with very similar SU-8 thicknesses. This is most likely due to imperfections in the fabrication process such as misalignment from the waveguide and irregularities in the shape of the disks. A small systematic positioning error on the order of 200 nm was observed after fabrication, with most SU-8 disks located slightly off centre from the axis of the waveguide. The systematic offset was almost certainly caused by a slight misalignment between the blue and red lasers, while the random variation came from human error when manually

positioning the laser spot. The latter would not be a problem when using μ PL measurements to centre the red laser spot over a QD, but the systematic error could still lead to sub-optimal Q factors.

The other important difference is that for thicker SU-8 the experimental Q factors actually exceed those predicted by FDTD. This is because the simulations assumed a disk diameter of $1\text{ }\mu\text{m}$ while many of the fabricated SU-8 disks were significantly larger. More recent simulations have since confirmed that Q factor is positively correlated with disk diameter. The SU-8 thicknesses and disk diameters are plotted in Fig. 5.26(b), and we can see a clear positive correlation. The dose was constant across all exposures, so the best explanation for this is the increased quantity of photo-acid catalyst produced when thicker layers are exposed.

5.6 Conclusions

The main result of this chapter is the experimental demonstration of the theoretical results presented in chapter 2. We successfully fabricated 34 cavities, all with Q factors in excess of 2×10^3 . The maximum Q factor of 7.4×10^3 is comparable to the Q factors of typical L3 and Notomi cavities on similar samples, showing that this cavity design has the potential to be used for the same kind of light-matter coupling applications. Further improvements on this performance, particularly in regards to reliability, are expected.

The photo-lithographically defined cavities presented here also have the significant advantage of allowing μ PL-based positioning techniques to be employed, potentially allowing cavities to be positioned directly over buried QDs. The samples used for the measurements in this chapter had a relatively high density of QDs so the Purcell enhancement was observed for an ensemble of many QDs rather than for a single QD. With only ensemble μ PL data it wasn't possible to calculate the exact magnitude of the Purcell enhancement, but we know Purcell enhancement must be occurring to produce the observed spectral lines. As of the writing of this chapter we are in the process of acquiring some low-density QD samples which will be used to test coupling to single QDs.

While there are some limitations to this fabrication approach, we believe that many of them will be overcome in time. For instance, the problem of achieving good alignment between the 405 nm laser used for exposure and the longer wavelength μ PL excitation laser could potentially be solved by using a fibre optic splitter to combine the two input beams into a single output beam. Work on such improvements is ongoing.

6

Photoluminescence from Hybrid Perovskites

Contents

6.1	Introduction	136
6.2	Micro-Photoluminescence Studies	137
6.2.1	Temperature Dependence	138
6.2.2	Objective Maps	141
6.2.3	Macroscopic Maps	146
6.2.4	Mirror Maps and Carrier Diffusion	148
6.3	Time-Resolved μPL Studies	157
6.3.1	Measurement Procedure	157
6.3.2	Modelling and Fitting	161
6.3.3	Limitations on the Fitting Algorithm	167
6.3.4	Lifetime Fitting Results	172
6.3.5	Measuring the Diffusion Coefficient	175
6.4	Conclusions	179

6.1 Introduction

Hybrid organic-inorganic perovskites have received a huge amount of attention over the last decade due to their promising optoelectronic properties. The solar conversion efficiencies of perovskite-based single-junction photovoltaic cells have rapidly risen from just 3.8% in 2009 [195] to more than 23% in 2018 [127]. While this is still

significantly lower than the best single-junction efficiencies achievable with conventional silicon-based and GaAs-based photovoltaics (28% and 29% respectively) [127], perovskite cells have the potential to be manufactured at scale for around half the cost per square metre of multicrystalline silicon cells [196] which exhibit similar performance. Historically the biggest hurdle to the commercialisation of perovskite-based photovoltaics has been their chemical instability, however recent work has demonstrated significant progress towards developing highly stable perovskites [197]. Despite all this attention, however, the fundamental physical properties of many members of the perovskite family are still quite poorly understood.

In this chapter two main experimental approaches are used to study perovskite samples. The first section covers the results of time-integrated micro-photoluminescence (μ PL) measurements in which the emission spectrum is recorded with exposure times on the order of seconds. This includes various types of μ PL mapping where the excitation and collection positions are moved around the sample independently.

The second section of this chapter covers time-resolved photoluminescence (TRPL) measurements in which the emission intensity is recorded as a function of time with sub-nanosecond resolution. In combination with a fast pulsed laser source these TRPL measurements facilitate the study of carrier recombination processes. This section also includes a detailed description of the computer algorithms used to analyse the TRPL results.

6.2 Micro-Photoluminescence Studies

The organic-inorganic hybrid perovskites studied in this chapter are composed of a mixture of formamidinium (FA), caesium, lead, iodine, and bromine with the formula $\text{FA}_{0.83}\text{Cs}_{0.17}\text{Pb}(\text{I}_{1-x}\text{Br}_x)_3$. This same composition has previously been studied by McMeekin *et al.* [141]. Four different compositions were investigated with x values of 0.4, 0.3, 0.2, and 0.1. The samples were fabricated by spin-coating a solution onto 1 inch squares of glass substrate. The resulting layer was polycrystalline with crystal size estimated to be on the order of 100 nm. More details of these materials can

be found in chapter 4. No other structures (e.g. hole transport layers or protective coatings) were added to the perovskite layer.

While in most practical situations photovoltaic materials would be expected to operate at temperatures around 300 K, measurements performed at cryogenic temperatures can often provide more useful information about the fundamental physical properties of materials. In fact, in the case of some of the samples studied in this chapter cooling was necessary in order to get a useful level of PL emission while keeping the laser power low enough to avoid causing damage. As in previous chapters, cooling was achieved using a Janis continuous flow liquid helium cryostat with an integrated heater element and external PID controller providing precise temperature control over the range from 4.2 K to 300 K.

Due to the high chemical instability of these perovskite samples it was necessary to store them in a nitrogen atmosphere while not in use. Once loaded into the μ PL system, the samples were kept under vacuum inside the cryostat at a pressure of no more than 100 mPa. The vacuum did not appear to have any adverse effects on the samples even over periods of several hundred hours, however the samples eventually deteriorated due to atmospheric contamination of the storage chamber.

6.2.1 Temperature Dependence

To begin with, simple μ PL spectra were recorded at a single point on each sample as a function of temperature. Maintaining a consistent position between measurements was difficult because the visual appearance of the sample surface was quite uniform¹ and there was no way to fabricate clearly visible registration marks. Positioning was performed using the manual micrometer stage which held the cryostat, and the four measurement locations (one on each sample) were recorded by reading the micrometer scales. This gave a positioning accuracy of around 10 μ m. While it would have been possible to eliminate this error by measuring one sample at a time and keeping the position constant, thermal expansion during temperature changes would have introduced a similar error.

¹The imaging system has since been improved, allowing subtle features on the sample surface to be clearly identified and enabling much more accurate position registration.

The laser source for these measurements was a 532 nm continuous wave laser produced by frequency doubling the output of a 1064 nm Nd:YAG laser which itself was pumped by an 808 nm laser diode. The beam was coupled through an optical fibre to produce a clean Gaussian profile, and neutral density (ND) filters were used to attenuate it down to a power of 33 nW as measured immediately above the cryostat window. Given a spot size of about $1 \mu\text{m}^2$, this power was roughly 33 times the standard solar irradiance measured on the Earth's surface (1 kW m^{-2}). Initially higher powers were tested, however those caused the PL emission to gradually decrease, suggesting that the sample was being damaged.

Measurements started with the samples at 4.2 K, then the temperature was increased. After each temperature change the cryostat was allowed to settle for around 15 minutes while the PID controller stabilised, then the four samples were measured in sequence. The integration time for each μPL measurement was adjusted from a minimum of 1 s up to 5 s as the emission rate dropped. The wavelength range recorded by the spectrometer was also adjusted as the emission spectrum of the samples changed. A 600 lines/mm grating was used in the spectrometer. The resulting spectra are plotted in Fig. 6.1 and Fig. 6.2.

One interesting feature of these data is that the band gap (and hence emission energy) tends to increase with increasing temperature. This behaviour is the reverse of the usual trend for semiconductors, but has been observed previously in other perovskites such as MAPbI_3 [198], $\text{MAPb}(\text{I}_{1-x}\text{Cl}_x)_3$ [199], and CsSnI_3 [200]. It has been hypothesized, at least in the case of MAPbI_3 , that this anomalous temperature dependence is caused by the lengthening of the Pb-I bonds with decreasing temperature [198]. It is those bonds that are primarily responsible for determining the electronic properties of the perovskite semiconductors. Whether this explanation is also applicable to the FA-based perovskites studied in this work remains unknown.

These samples also follow the previously observed trend [145, 147, 146] of increasing band gap as the fraction of bromine is increased in mixed halide perovskites. This can be explained by the fact that bromine has a smaller




Figure 6.1: Plots of μ PL spectra recorded at a single point on each perovskite sample as a function of temperature. The four plots correspond to samples with I:Br ratios of: (a) 60:40, (b) 70:30, (c) 80:20, and (d) 90:10.

atomic radius than iodine and hence forms shorter bonds to the surrounding lead atoms [142].

The PL intensity increases dramatically as the temperature is reduced, particularly for the bromine-rich samples. This is most likely caused by a suppression of non-radiative recombination pathways at low temperatures. There is also a slight broadening of the emission at higher temperatures, probably due to phonon interactions.

The small peak visible at 808 nm in some of the spectra was not caused by emission from the perovskite samples. Rather it came from leakage of the pump laser into the output beam of the Nd:YAG laser. The main 532 nm beam was prevented from entering the spectrometer by a narrow notch filter, but that did nothing to block the 808 nm leakage. In later experiments a narrow band-pass filter centred on 532 nm was added to the excitation path to remove that peak.




Figure 6.2: Plots of μ PL spectra recorded at a single point on each perovskite sample as a function of temperature. These are the same as the plots in Fig. 6.1, but all the spectra have been normalised and are plotted on a linear scale for ease of comparison.

6.2.2 Objective Maps

In order to get a better understanding of the uniformity and micron-scale features in these lead-halide perovskite samples we took advantage of the 2D μ PL mapping capability of the system. These were what we termed “*objective maps*” because they involved using the three-axis piezo-actuated stage to raster scan the objective lens over the sample. This had the effect of moving both the excitation laser spot and the region from which light was collected in unison so that they remained aligned. This was the simplest mapping mode and the same one that was used in chapter 5.

The main difference between this system and the one used in chapter 5 was the addition of an optical fibre to the collection path. Instead of relying on the input slit of the spectrograph to restrict the area on the sample from which light was collected, the collected light was focussed into a fibre and then the other end of




Figure 6.3: Plots of μPL data from all four samples at a temperature of 4.2 K. (a-d) show the normalised PL spectrum after integration over the whole $50 \mu\text{m} \times 50 \mu\text{m}$ mapping area. (e-h) show the spatial distribution of PL emission from two different wavelength ranges. The wavelength integrals corresponding to the red and blue colouration are shown in the spectra on the left.




Figure 6.4: Plots of μPL data from all four samples at a temperature of 290 K. (a-d) show the normalised PL spectrum after integration over the whole $50 \mu\text{m} \times 50 \mu\text{m}$ mapping area. (e-h) show the spatial distribution of PL emission from two different wavelength ranges. The wavelength integrals corresponding to the red and blue colouration are shown in the spectra on the left. The extraneous signal (due to laser pump leakage) between 805 nm and 810 nm in (d) has been excluded from the integral.

the fibre was used as the point-source for the imaging spectrograph. If we imagine sending light in the reverse direction and imaging the slit onto the sample surface, it would look like a line with a width on the order of $1\text{ }\mu\text{m}$ but a length of many microns. In contrast, the core of the fibre would be a uniform circle (or technically an Airy disk thanks to diffraction). With a $100\times$ magnification objective, a $5\times$ fibre-coupling lens, and a $25\text{ }\mu\text{m}$ diameter core in the fibre, the resulting collection spot would have a diameter of around $1.25\text{ }\mu\text{m}$ which is only slightly larger than the minimum achievable spot size in the system.

Additionally, in an attempt to keep the sample as pristine as possible, automatic control over the laser shutter was added to the mapping software. A map could be started with the shutter blocking the laser, then it would be opened only for the duration of the map. After the map was completed it would automatically close again, preventing damage to the sample from prolonged laser exposure. This system had the added benefit of allowing background spectra to be recorded with a single button press.

Maps of each of the four samples were performed at a temperature of 4.2 K where the PL emission was at its brightest for the 40% and 30% bromine samples. While 50 K would have resulted in greater emission from the 10% and 20% samples, those were already significantly brighter. Maximising the brightness allowed the exposure time at each pixel of the map to be minimised, and hence significantly reduced the time taken to perform the whole map. Each map covered a $50\text{ }\mu\text{m} \times 50\text{ }\mu\text{m}$ area with a step size of $0.5\text{ }\mu\text{m}$ in each axis. The laser power was 33 nW with an exposure time of 1 s .

Fig. 6.3 shows the four maps of the samples along with the corresponding spectra produced by integrating all the points in each map. The maps are coloured with two additively mixed colour scales. The red channel shows the relative intensity of the longer wavelength region indicated in the spectrum above each map. The blue channel shows the relative intensity of the shorter wavelength region. Darker colours indicate an absence of emission, while purple would indicate emission of

equal intensity at both long and short wavelengths. Each spectral integration window is 10 nm wide, and they are separated by 10 nm.

In all cases there is very obvious variation across the mapped area. The surface seems to be divided into areas which are predominantly red and areas which are predominantly blue. The anti-correlation is strong, with relatively few areas where the emission is equal across both wavelength bands. This indicates that there is significant non-uniformity in the emission spectrum, most likely caused by variations in composition. The phenomenon of halide segregation in mixed-halide perovskites has previously been directly observed by scanning helium-ion microscopy and secondary ion mass spectrometry [201], although the authors of that paper only report segregation on the order of hundreds of nanometres. The PL-differentiated domains in these samples range from the scale of 1 μm in the $\text{I}_{0.6}\text{Br}_{0.4}$ sample to around 10 μm in the $\text{I}_{0.9}\text{Br}_{0.1}$ sample. It is also possible that the formamidinium and caesium cations segregate to produce FA-rich and Cs-rich regions, or there could be some combination of both types of segregation.

Similar objective maps were also performed at a temperature of 290 K with the same resolution, though not at the same positions. Due to the low luminosity of the $\text{I}_{0.6}\text{Br}_{0.4}$ sample, that was measured with twice the laser power and a 5 s exposure time per pixel. These results are plotted in Fig. 6.4.

It is clear that at 290 K there is much less variation in emission intensity between the two integration windows. This can partly be explained by the broader emission spectrum at this temperature. We can imagine the surface being segregated into two compositions with approximately Gaussian emission spectra. For the same centre-to-centre wavelength separation between the two Gaussians, the difference in intensity at any given wavelength decreases as the width of the Gaussians is increased. An attempt has been made to compensate for this by increasing the separation of the integration windows in proportion to the spectral broadening. Even after this compensation there is still significantly less variation, suggesting reduced segregation of halides or cations. This reduction in segregation at higher

temperatures persisted after several cycles of repeated cooling and heating, indicating that the segregation is a reversible process.

6.2.3 Macroscopic Maps

The area over which maps could be performed using the piezo-actuated stage was limited to only $100\text{ }\mu\text{m} \times 100\text{ }\mu\text{m}$. To investigate the compositional uniformity at larger scales it was necessary to use the manual micrometer stage which supported the sample cryostat. Because all movements had to be done manually, only a 1D line across each sample was mapped. μPL spectra were recorded every $20\text{ }\mu\text{m}$ over a total distance of 1 mm . These measurements were repeated on each sample at temperatures of 4 K and 300 K . A laser power of 110 nW was used for the measurements at 4 K while a power of 275 nW was used at 300 K due to the reduced brightness.

The resulting μPL spectra are plotted in Fig. 6.5. The 4 K maps are on the left and the 300 K ones are on the right. Each plot has been background corrected and normalized so that the maximum emission corresponds to a value of 1. The spectral broadening at higher temperatures is much more obvious in these plots, as is the variation in emission spectrum with location. Of course the sampling points were sparse, so these data only give an estimate of the variation of the emission spectrum.

One interesting feature is that in some of the maps there is a variation in brightness from one end to the other. This isn't consistent with compositional variation, but could indicate a variation in the thickness of the spin-coated layer. Alternatively it may have been caused by slight drift of the focal plane if the samples weren't completely level.

The magnitude of the wavelength variation could be quantified by extracting the wavelength of the maximum intensity emission in each spectrum then taking the standard deviation of those wavelengths within each map. These standard deviations are listed in Table 6.1. This quantitatively confirms that there is much greater variation at 4 K than at 300 K . Interestingly, despite the physical scale




Figure 6.5: Plots of the normalised μ PL emission spectra as a function of 1D position from hybrid perovskite samples of four different compositions at temperatures of 4 K and 300 K. The position was stepped in increments of 20 μ m, and the positions were different at the two temperatures.

I:Br ratio	$T = 4\text{ K}$	$T = 300\text{ K}$
60:40	9.2 nm	3.8 nm
70:30	7.7 nm	3.4 nm
80:20	8.2 nm	3.4 nm
90:10	3.8 nm	2.6 nm

Table 6.1: Standard deviations of the wavelengths at which the maximum μPL intensity was measured. The rows correspond to different samples, while the columns are different temperatures. The raw data used to calculate these values is plotted in Fig. 6.5.

of the segregation being greatest in the $\text{I}_{0.9}\text{Br}_{0.1}$ sample, the standard deviation of its emission wavelengths was the smallest.

6.2.4 Mirror Maps and Carrier Diffusion

One of the properties of perovskites (and semiconductors in general) which is very important for photovoltaic applications is the carrier diffusion length. In order to extract electrical power from photon absorption events, the generated carriers must reach the conductive contacts before recombining. Thus the carrier diffusion length of the absorptive material places constraints on the device design. Short diffusion lengths can be overcome, for instance by adding electron/hole transport layers [202] or using nano-structured surfaces for carrier collection [203], but longer diffusion lengths are generally preferable. There have been a number of reports of very long carrier diffusion lengths in perovskites [204, 158, 157], which bodes well for these materials.

Carrier diffusion is driven by the random thermal motion of carriers in the presence of a carrier concentration gradient. Electrons (or holes) undergo random walks which, on average, transport them from areas of high electron (or hole) concentration to areas of lower concentration. The rate of diffusion is described mathematically by Fick's laws [205]. Care should be taken to avoid confusing it with carrier drift which is driven by electric fields. Because electrons and holes can have different effective masses and interact differently with defects, the electron and hole diffusion lengths are not generally identical.

The equations describing carrier diffusion are largely the same as the equations describing any other kind of diffusion due to concentration gradients. The rate of

carrier movement at any point is proportional to the magnitude of the concentration gradient at that point and is directed towards the lower concentration. The current, J , in one dimension is given by equation (6.1) in which q is the charge on each carrier, ρ is the carrier concentration, and D is known as the diffusion coefficient and depends on the material. This equation is sometimes called Fick's first law.

$$J = -qD \frac{d\rho}{dx} \quad (6.1)$$

It is important to point out that the current, J , in equation (6.1) is only the contribution from one type of carrier. If both electrons and holes are diffusing, as is usually the case, the total current would be given by the sum of the individual currents.

In most cases we are more interested in knowing how the carrier concentration changes as a function of space and time, and this is given by Fick's second law. Equation (6.2) shows the one dimensional case, but it is easy to extend it into higher dimensional systems.

$$\frac{d\rho}{dt} = D \frac{d^2\rho}{dx^2} \quad (6.2)$$

While it is not generally easy to find exact solutions to equation (6.2) under arbitrary boundary conditions, simple cases such as step functions have been thoroughly studied. An example of the time evolution of a top-hat density distribution under diffusion is plotted in Fig. 6.6(a). Interestingly, despite starting from a decidedly non-Gaussian shape, the density distribution rapidly converges towards a Gaussian. This convergence is quantified in Fig. 6.6(b) where Gaussian functions have been fitted to each of the curves and the total squared difference is plotted against time. This convergence towards a Gaussian distribution will be important when we come to analyse the experimental data.

In practice, the reason we can talk about a diffusion length is because the carriers only have a finite lifetime before recombining. Electrons and holes are created in one location, diffuse down their respective concentration gradients, then randomly encounter one another and recombine. The typical distance travelled before recombining is termed the diffusion length. A caveat of this is that the




Figure 6.6: (a) is a plot of the diffusion of a simple top-hat density distribution according to equation (6.2). The values of D , x , and t have been chosen arbitrarily to show the general time evolution behaviour. (b) is a plot of the total squared difference between each of the curves in (a) and the corresponding best fit Gaussian.

diffusion length can vary hugely depending on the local carrier concentration, since that determines the carrier lifetime.

The diffusion coefficient, D , is related to a more fundamental property: the electrical mobility, μ . In equation (6.3), known as the Einstein relation, k_B is the Boltzmann constant, T is the temperature, and q is the charge on the carrier species. However, this equation only applies if the energies of the electrons can be described using a single value of T . If the energy distribution is non-thermal then the relationship between mobility and diffusion coefficient becomes more complicated.

$$D = \frac{\mu k_B T}{q} \quad (6.3)$$

The mobility also determines the drift velocity, $\mathbf{v}_d$, of carriers when an electric field, $\mathbf{E}$, is applied. This is given by equation (6.4).

$$\mathbf{v}_d = \mu \mathbf{E} \quad (6.4)$$

It is difficult to directly observe carrier diffusion, but we can observe the light emitted when carriers recombine. The $4f$ system and scanning mirror which were




Figure 6.7: Diagram depicting the carrier diffusion process which contributes to increasing the spatial extent of the PL emission.




Figure 6.8: Diagram showing the procedure for performing μ PL maps with the excitation spot moving independently of the collection spot. The position of the excitation spot was controlled using a $4f$ system to angle the beam entering the top of the microscope objective. The emitted light was collected from a fixed point at the centre of the mapping area. The angle of the excitation laser and the size of the mapping area and working distance have been greatly exaggerated in this diagram.

already integrated into the μ PL system allowed the excitation laser spot to be moved independently of the light collection spot. This meant that carriers could be generated via photon absorption in one location while the PL emission was measured at a different location many microns away. As long as both the laser spot size and collection area were kept small, this arrangement allowed us to indirectly map the diffusion of carriers. Fig. 6.7 shows a simplified diagram of the process. Note that this only allowed us to see the areas where both electrons and holes were present, so we couldn't measure differences between electron and hole diffusion lengths.

The procedure for performing μ PL maps using the scanning mirror is illustrated in Fig. 6.8. These are referred to as “*mirror maps*”. The objective was held stationary over the centre of the mapping area while the excitation laser was raster scanned back and forth. At each measurement location a PL spectrum was recorded using the spectrometer.

Prior to doing any mirror maps the coordinate systems of the excitation and collection spots had to be aligned. This was done by setting both the scanning mirror and piezo stage to their central positions and then manually adjusting the alignment of the system until the light intensity reaching the spectrometer was maximised. As long as a uniform sample was used for alignment (typically a GaAs quantum well), this ensured that the excitation and collection spots were in the same position on the sample.

A series of μ PL mirror maps were performed at temperatures from 4 K to 300 K in increments in 50 K on each of the four samples. While an effort was made to keep the maps in the same location on each sample, some drift was inevitable due to thermal expansion of the sample and cryostat internals. Each map spanned an area of $50\text{ }\mu\text{m} \times 50\text{ }\mu\text{m}$ in steps of $1\text{ }\mu\text{m}$. This size was chosen because there was negligible emission for excitation-collection separations beyond $30\text{ }\mu\text{m}$. The Nd:YAG laser power was maintained at 33 nW with an exposure time of 1 s at each position in the map. Fig. 6.9 shows some examples of these maps after integrating over the whole wavelength range collected by the spectrometer.




Figure 6.9: Plots showing the total integrated PL intensity measured at the centre ($x = 25 \mu\text{m}$, $y = 25 \mu\text{m}$) as a function of the excitation location for all four samples. All these measurements were performed at a temperature of 4.2 K using 33 nW excitation power. The colour corresponds to integrated PL intensity on a log scale, clipped to a range of two orders of magnitude.

Before the diffusion length could be estimated, the data had to be transformed from the Cartesian coordinates of the mapping system to polar coordinates centred on the collection spot. Despite the system being calibrated before starting measurements, there was inevitably some small ($< 1 \mu\text{m}$) offset between the collection spot and the centre of the mapping area. If not corrected for, this offset would introduce errors into the distance measurements. To determine the real location of the collection spot with sub-pixel precision, a least-squares minimisation algorithm was used to fit a Gaussian function to a 1D line of pixels $10 \mu\text{m}$ in length taken from the centre of the data. This fitting was performed along both the X and Y directions and the centres of those Gaussian functions were then used as the coordinates for the centre of the polar transformation. The angular (θ) polar coordinate was then




Figure 6.10: A typical example of the variation in PL intensity as the excitation-collection separation is increased. These data (blue) were taken from the 20% Br sample at a temperature of 150 K. The fitted curve (orange) is a Gaussian added to an exponential.

discarded, leaving just the separation and integrated PL intensity information.

The blue points in Fig. 6.10 are a typical example of the data from one of the mirror maps after applying this transformation. Interestingly there are two distinct regions in the plot. At short distances the PL intensity drops off in a Gaussian shape, but at longer distances an exponential term dominates until the background noise level is reached. It is tempting to say that the Gaussian part comes from the finite size of the laser spot, but even when convolved with the collection area the diameter is too small to produce the observed effect. Instead the Gaussian region must come from carrier diffusion.

The exponential region, however, is not caused by carrier diffusion alone. The dominant process at longer distances is photon recycling. This process, depicted in Fig. 6.11, involves the re-absorption and re-emission of previously emitted photons [206, 207]. In each emission event there's a significant chance of the photon being emitted in a direction close to parallel with the perovskite layer. Total internal reflection at the perovskite-air interface can trap these photons, allowing them to be re-absorbed by the perovskite. Because the photons are travelling through an




Figure 6.11: A diagram illustrating the photon recycling process responsible for the exponential section of Fig. 6.10.

absorptive medium and each absorbed photon has a chance of being lost (either via out-of-plane emission or non-radiative decay), the photon flux (and hence the PL emission) decreases exponentially with distance.

A simple model consisting of a Gaussian plus an exponential was fitted to the data using a least-squares minimisation algorithm. The precise form of the model was,

$$I = A_g \exp\left(\frac{-r^2}{2c^2}\right) + A_e \exp\left(\frac{-r}{l}\right),$$

where I was the PL intensity, r was the excitation-collection separation, c was the Gaussian width parameter, l was the exponential length parameter, and A_g and A_e were amplitude parameters. The full width at half-maximum (FWHM) of the Gaussian was related to c by the expression $\text{FWHM} = 2\sqrt{2\ln(2)}c$. This model was found to fit the data well in all cases, however at higher temperatures the exponential component was not visible due to the low signal-to-noise ratio. The curve fitted to the data in Fig. 6.10 is indicated by the orange line.

The same transformation and fitting procedure was performed for all of the maps on the four samples at different temperatures. The Gaussian FWHM values are plotted in Fig. 6.12(a) and the values of the exponential length, l , are plotted in Fig. 6.12(b). The data is too noisy to resolve any significant differences between the different compositions, but there is a visible trend of increasing FWHM with




Figure 6.12: Plots showing the variation of the model parameters fitted to the mirror map data as a function of temperature for all four samples. (a) is a plot of the full width at half-maximum of the Gaussian component, while (b) is a plot of the characteristic length of the exponential tail (corresponding to a factor of $1/e$ change). Some points have been omitted from (b) because the noise level was too great to resolve the exponential component.

increasing temperature. This indicates that the carrier diffusion length also increases with temperature. There's no significant trend in the length of the exponential, which is consistent with the optical absorption coefficient of the perovskite layer remaining roughly constant.

From these data we can estimate a diffusion length, L_D , of around $2.5 \mu\text{m}$ at 4 K rising to around $4 \mu\text{m}$ at room temperature. This range of values is significantly shorter than the extremely long diffusion lengths (as much as $100 \mu\text{m}$) measured in single crystal perovskites, however it is very close to the diffusion lengths previously reported in polycrystalline FAPbBr_3 ($1.3 \mu\text{m}$) and FAPbI_3 ($3.1 \mu\text{m}$)

at room temperature [156]. This diffusion length is also significantly longer than the optical absorption length (typically a few hundred nm) which makes it comparatively easy to design photovoltaic cells with good charge extraction efficiency.

The carrier diffusion length is dependent upon the carrier density in the material. In the case of an extremely low carrier density (or a large disparity in the density of electrons and holes), the carriers can be made to diffuse for almost arbitrarily long distances because the recombination rate is so low. Conversely, if the carrier density is very high then the diffusion length will be limited by various recombination processes. The maximum carrier density in these measurements was estimated to be around $(8 \pm 2) \times 10^{16} \text{ cm}^{-3}$ using the extinction coefficient at 532 nm ($\kappa = 0.4 \pm 0.1$) measured by Tejada *et al.*[208] and the recombination rate coefficients measured by McMeekin *et al.*[141].

The data shown in Fig. 6.10 is far noisier than would be expected from simple $\sqrt{N}$ Poissonian statistics. The source of this extra noise is most likely the non-uniformity of the sample composition leading to anisotropy in the carrier diffusion and photon recycling. Because recycled photons are emitted at energies very close to the band gap, small variations in the band gap energy result in large changes in the absorption coefficient and could therefore have a large effect on the range of the photon recycling effect. The apparent level of the noise could be reduced by restricting the analysis to a narrow slice of points instead of the full 2D plane, but this would not provide any more precision because the composition would still be varying along the length of that slice.

6.3 Time-Resolved μ PL Studies

6.3.1 Measurement Procedure

To learn more about the carrier dynamics within these perovskite samples during photoexcitation, the μ PL system was expanded to also be capable of carrying out time-resolved photoluminescence (TRPL) measurements. The necessary pulsed excitation was provided by a 450 nm diode laser with a base repetition rate of




Figure 6.13: Diagram showing the spectrometer output slit and photo-multiplying detector used for TRPL measurements. The width of the slit could be adjusted to restrict the range of wavelengths reaching the PMT. The spectrum could be moved relative to the slit by rotating the diffraction grating, allowing different parts of the spectrum to be measured.

80 MHz and a pulse duration of less than 70 ps. A setting on the laser controller allowed the base repetition rate to be divided by powers of 2 up to a maximum of 16, yielding a minimum repetition rate of 5 MHz.

The spectrometer already had a second output port with a computer controlled mirror which could redirect the diffracted light. A photomultiplier tube (PMT) detector was attached directly to this output port and opaque shielding was used to block stray light from the room. No additional optics were required to focus the light on the PMT since the imaging spectrograph already provided a small spot size (technically a horizontal line of length ≤ 2 mm and height ≤ 50 μ m) at the output port and the distance to the PMT was not sufficient for the light to spread out much. The output port of the spectrometer had a built-in slit, the width of which could be adjusted electronically. Thus, by combining movements of the grating and the slit, it was possible to select any arbitrary range of wavelengths (up to a maximum width of around 10 nm) to project onto the PMT. This arrangement is depicted in Fig. 6.13.

The electronic output from the PMT was connected to the signal input of a time-correlated photon counting (TCPC) card in a desktop computer. The laser controller was also connected to this card and provided electronic synchronisation (sync) pulses every time a laser pulse was emitted. The TCPC card recorded, to an accuracy of 25 ps, the delay between each sync pulse and the subsequent signal from the PMT. Over many millions of laser pulses a histogram could be built up showing the average photon emission rate as a function of time. The signal was periodic with the same period as the laser pulses, so the maximum duration which could be observed was 200 ns when using the lowest repetition rate.

In order to properly analyse the TRPL data it was essential to first record the instrument response function (IRF) of the system. Every component in the system, particularly the electronic parts, had the potential to introduce small distortions and echoes in the signal. An ultra-short pulse of light arriving at the PMT could easily get stretched to a duration of several hundred picoseconds and given a complex modulation with multiple smaller peaks in the time domain. To record the IRF the laser was focussed on a reflective sample (smooth GaAs) and the spectrometer was adjusted so that the reflected laser light reached the PMT. While the duration of one laser pulse was slightly longer than the temporal resolution of the TCPC card, it was still sufficiently short to approximate a delta function when determining the IRF.

When performing TRPL measurements, and especially when measuring the IRF, it was important to always limit the rate at which photons arrived at the PMT. After each detection event the system took a certain amount of time, known as the dead time, to reset itself in preparation for another detection event. Any photons arriving during this dead time would be ignored, leading to errors in the resulting data. The dead time for the TCPC card used in this work was 25 ns. This meant that, assuming a perfectly Poissonian light source and a perfectly efficient detector, a photon rate of $6 \times 10^6 \text{ s}^{-1}$ would result in approximately 1% of the photons failing to be counted. Of course a pulsed laser is very definitely not a Poissonian light source, nor is the emission from the sample when illuminated with such a laser, so the actual loss rate would be higher. Each laser pulse was much

shorter than the dead time of the detector so under laser illumination (which is the worst case scenario) any laser pulse containing more than 1 photon would result in missed detection events. The distribution of photons among laser pulses should be Poissonian, so the probability of having n photons in a pulse is given by,

$$P(n) = \exp(-\bar{n}) \frac{\bar{n}^n}{n!},$$

where $\bar{n}$ is the mean number of photons per pulse. Given a laser repetition rate of 5 MHz, we can calculate that a signal rate of just $7 \times 10^5 \text{ s}^{-1}$ would result in an error rate of around 1%.

The control software displayed the current count rate, so to ensure accuracy in the collected data we were careful to always keep it below $1 \times 10^5 \text{ s}^{-1}$. This was done both by adding neutral density (ND) filters to the optical path and by narrowing the output slit on the spectrometer. To prevent reflected laser light from reaching the spectrometer during normal measurements a 600 nm long-pass filter was inserted into the collection path.

Before doing any measurements, the TRPL system was also used to measure the background. The laser was switched on, providing a sync signal, but the shutter was closed so that it could not excite the sample. Because the background noise was not periodic all of the time bins could be average to produce a very precise estimate of the noise level. This average was then subtracted from every time bin in the measured histogram.

TRPL measurements were made on three of the samples with bromine percentages of 10%, 20%, and 40% at a temperature of 4.2 K. A range of excitation powers were used, from a minimum of 25 nW up to a maximum of 1600 nW. It was hoped that by doing measurements at a range of powers, the behaviour could be probed over a wider range of carrier densities, making it easier to build a model of the system.

Due to the long photoluminescence lifetimes of these samples, the laser was set to its lowest repetition rate of 5 MHz. A pristine site on each sample was chosen for measurements and the lowest power was used first so as to minimise

damage. Time-domain data was collected for time periods ranging from 60 s to 600 s depending on the intensity of the PL emission. Immediately after each TRPL measurement a μ PL spectrum was also recorded at the same location and power. The recorded TRPL data are plotted as solid lines in Fig. 6.15. Note that the plots on the right side of this figure show the same data as those on the left but with the time axis changed to better show the fast decay processes.

Because some of the TRPL data were recorded using different output slit widths, and hence different wavelength ranges, a correction was applied to the data before fitting. The μ PL spectra from the same series of measurements were used to calculate this correction. The counts in each set of TRPL data were multiplied by the fraction of the counts in the μ PL spectrum which fell within the wavelength range transmitted by the slit. This produced values which approximated what would be expected to be measured if the TRPL measurements included the entire wavelength range.

6.3.2 Modelling and Fitting

Analysis of this kind of TRPL data typically involves developing a theoretical model of the system and fitting this to the data. The physical properties of the material can then be extracted from the parameters of the fitted model. In some systems the model can be extremely simple (e.g. a single exponential decay), however in the case of these perovskites the decay shape is evidently much more complicated. For instance, the group studying the same type of mixed-cation lead mixed-halide perovskites (albeit using a different method) used a model involving three different orders of decay rate terms [141].

In order to model the charge carrier dynamics of the perovskite samples, a number of approximations were made. The laser pulse duration was similar to the time resolution of the TCPC card, so it was safe to treat the initial excitation as instantaneous. Since the measurements were made at cryogenic temperatures it was likely that excitons would play a significant role in the photoluminescence process, so populations of both free carriers and excitons were included in the model.

However, in order to limit the number of fitting parameters, interactions between excitons and free carriers were not included in the model. The two populations were assumed to evolve entirely independently. The perovskite was also assumed to be undoped, meaning that the populations of free electrons and holes could be described by a single free carrier density parameter, n_f .

Three different decay pathways were modelled for free carriers: trap-assisted (Shockley-Read-Hall) recombination which scales as a monomolecular process [209, 210], conventional band-to-band recombination which scales as a bimolecular process, and Auger recombination which scales as a trimolecular process. Of those decay pathways, trap-assisted and Auger recombination were assumed to be entirely non-radiative while photons were only emitted from band-to-band recombination. The full rate equation for free carrier decay is given in Eq. 6.5, where k_{f1} , k_{f2} , and k_{f3} are rate coefficients which were varied to fit the TRPL data.

$$-\frac{dn_f}{dt} = \underbrace{k_{f1}n_f}_{\substack{\text{carrier} \\ \text{trapping}}} + \underbrace{k_{f2}n_f^2}_{\substack{\text{radiative} \\ \text{recombination}}} + \underbrace{k_{f3}n_f^3}_{\substack{\text{Auger} \\ \text{recombination}}} \quad (6.5)$$

The model for exciton decay was much simpler, containing only a monomolecular radiative recombination term. The rate equation for the exciton population density, n_x is given in Eq. 6.6, where k_x is the rate coefficient.

$$-\frac{dn_x}{dt} = k_x n_x \quad (6.6)$$

Under normal solar illumination the carrier density usually remains so low that the contribution from Auger processes is insignificant, however in this case the excitation arrived in short pulses with a far greater power density. At an average power of 200 nW and a repetition rate of 5 MHz each pulse delivered 4×10^{-14} J of energy which is equivalent to roughly 9.1×10^4 photons at 450 nm. Assuming an excitation volume of around $0.1 \mu\text{m}^3$ and a conversion efficiency of unity, the carrier density could reach as high as $9 \times 10^{17} \text{ cm}^{-3}$. The Auger rate coefficient, k_{f3} , has been measured to be around $10^{-28} \text{ cm}^6 \text{ s}^{-1}$ in other perovskites [211]. If

the value was similar for these samples the Auger contribution could easily be significant, hence its inclusion in the model.

Computer code (provided in appendix C) was written to numerically solve these differential equations when provided with values for all of the parameters, returning values for n_f and n_x as a function of time. This code utilised a Runge-Kutta method of order 5 with error estimation provided by an order 4 solver running in parallel [212], as provided by the SciPy library. In an attempt to keep the execution time of the fitting algorithm reasonable, the differential equations were only solved up to a time of 1 μ s after the initial excitation. By this point the emission rate was found to be negligible ($\ll 1\%$ of its value at time $t = 200$ ns), so that limitation shouldn't significantly affect the results. Examples of the resulting population density curves are shown in Fig. 6.14(a).

The TRPL measurements only observed the emitted light, so the two population density curves had to be transformed into a single curve tracking the total radiative recombination rate, R_{rad} , over time. This was done using the equation,

$$R_{\text{rad}} = k_{x1}n_x + k_{f2}n_f^2.$$

To emulate the time-binning performed by the TCPC card, the radiative recombination rate was integrated over the duration of each bin. This yielded a histogram of photon counts per time bin (up to a scaling factor), as plotted in Fig. 6.14(b).

The next step was to convolve this signal with the IRF in order to accurately reproduce the distortions induced by the measurement system. Rather than using the full 200 ns long IRF signal, it was cropped to include only the parts which exceeded a certain threshold. Shortening the IRF reduced the extent of the edge effects produced by incomplete overlap during convolution. The previously recorded background was also subtracted from the IRF before convolution. Fig. 6.14(c) shows the general shape of the signal after convolution with the IRF, although in that figure the IRF has been artificially stretched to make it more visible.

Finally the periodic nature of the laser pulses had to be accounted for. This wouldn't normally be necessary, but the perovskite materials studied here had such




Figure 6.14: Sequence of plots showing the process by which carrier decay and PL emission was modelled in order to fit the TRPL data from the perovskite samples. The duration of the IRF has been increased in these plots to make its effect more clearly visible. All of the vertical scales are in arbitrary units.

long luminescence lifetimes that the carriers produced by one pulse were not fully decayed by the time the next pulse arrived. This was approximated in the model by simply summing the PL emissions from 5 time-offset copies of the single-pulse decay. This approximation wasn't perfect since it didn't account for interactions between the carrier populations produced by different pulses, but the populations were so low after 200 ns that their contribution to non-linear effects should be negligible. The majority of emission at later times actually came from excitons which were assumed not to interact with each other at all. The final result is plotted in Fig. 6.14(d).

All of these steps were incorporated into a software library function which would return the simulated TRPL signal corresponding to the given initial conditions and decay constants. This function could then be called by a fitting algorithm in order to find the set of parameters which produced the closest match to the measured data.

However, in order to fit the model to the data the algorithm needed to be able to calculate a numerical estimate of the difference between the two curves. This meant that the arbitrary time offset in the measured data had to be corrected for. This could be done by adding a 'time offset' parameter to the modelling code and allowing the fitting algorithm to find the correct value by trial and error, however this approach led to poor convergence and significantly slowed down the fitting algorithm. It was more efficient to use a fast deterministic algorithm to detect the time offset and automatically apply this adjustment to the model. The algorithm detected the rising edge corresponding to the excitation pulse by scanning along the measured data and looking for the first occurrence of a value $\geq 60\%$ of the maximum. This threshold was chosen because it reliably located a consistent point on the rising edge in all cases. The same algorithm was used to find the corresponding point in the simulated data, then the time offset was simply given by the difference. Applying the time offset involved a simple adjustment of the periodic summation algorithm.

Various different error estimation functions were tested, including a simple sum of squared differences. The main difficulty arose from the fact that the functions contained both very intense, short-lived components and very weak, long-lived




Figure 6.15: Plots showing the recorded TRPL data (solid lines) and the corresponding models fitted to that data (dashed lines) for a range of excitation powers on each sample at a temperature of 4.2 K. The right hand column of plots (b, d, f) show the same data as (a, c, e) but with the time axis zoomed in to better show the fast decay immediately after the arrival of the excitation pulse.

components. When compared on a linear scale the high intensity part dominated the error function, leading to a poor fit to the long tail. When a logarithmic scale was tried instead the long tail dominated because it contained far more data points, leading to a poor fit to the short-lived components. The eventual solution was to divide the time axis into three regions corresponding to high, intermediate, and low intensity regions with durations of 3 ns, 24 ns, and 173 ns respectively. The data was log scaled and then the mean squared difference was calculated within each region. The errors for the three regions were multiplied by weights of 1, 2, and 2 respectively, then added together to produce a total error value. This error estimation algorithm was found to result in consistently good fits to all parts of the data.

Finally we get to the choice of fitting algorithm. In order to avoid problems caused by local minima, an approach known as basin-hopping was used. Basin-hopping is a two-stage process which alternates between gradient descent stages to find local minima and ‘hopping’ steps in which the parameters are abruptly changed in an attempt to find a different minimum [213, 214]. The hopping steps are random, so basin-hopping falls into the category of Monte Carlo minimisation algorithms. The implementation was provided by the SciPy library and the number of hopping steps was fixed at 100.

The fitting results are plotted as dashed lines in Fig. 6.15. In every case the fitted model was a close match to the experimental TRPL data. In fact, in an attempt to estimate the uncertainty in the various parameter values, each curve fitting was repeated at least 5 times with different initial parameters randomly chosen over a log-scaled range from 0.1 to 10 times the ‘typical’ values. The results were then compared to check for convergence.

6.3.3 Limitations on the Fitting Algorithm

Surprisingly, despite all ten simulated decay curves being very similar for each set of data, some of the fitted parameters differed by multiple orders of magnitude. To understand why this happened it was helpful to plot the various model parameters against one another. Most of these plots showed little or no correlation, but




Figure 6.16: Log-log plots of two of the rate parameters against the initial free carrier population parameter in the theoretical model after fitting it to one set of experimental TRPL data. Each data point corresponds to a separate fitting run with different initial values for the parameters. The red dashed line shows the trend and has the form $k \propto n^{-2}$.




Figure 6.17: A series of theoretical model results using various different parameters. The base parameters were taken from one of the fitting results to the TRPL data from the 40% Br sample at 800 nW. In each variation the initial free carrier population, n_{f0} , is multiplied by s and the bimolecular and Auger recombination constants, k_{f2} and k_{f3} respectively, are divided by s^2 . $s = 1$ corresponds to the original fitting results.

two in particular stood out. Examples of these for one particular dataset are shown in Fig. 6.16.

Here we see two of the rate parameters, k_{f2} and k_{f3} , plotted against the initial free carrier population density parameter, n_{f0} . Both plots show an extremely high correlation following the same power-law relationship. This correlation, coupled with the fact that the models produced almost identical curves, strongly indicates that the theoretical model is invariant under a particular transformation. Stated mathematically, if we define a scaling factor s which can take any real value, the resulting simulated TRPL data does not change when n_{f0} is multiplied by s and k_{f2} and k_{f3} are both divided by s^2 . The existence of such an invariance suggests that the model parameters are under-constrained by the available experimental data.

Upon further investigation it was found that the model was not strictly invariant under that transformation. Instead it contained two asymptotes corresponding to the $s \rightarrow \infty$ and $s \rightarrow 0$ limits. The $s \rightarrow 0$ asymptote corresponds to the case where the free carrier emission was negligible compared to the exciton emission and is not applicable to most of the data from this experiment. The $s \rightarrow \infty$ asymptote is more interesting, and an example of this asymptotic behaviour is plotted in Fig. 6.17. The inset shows the curves getting closer and closer together as s is increased. This asymptote corresponds to the limit where Auger recombination is the dominant free carrier decay process. Most of the experimental data lay very close to this Auger-dominant limit for at least a few nanoseconds after excitation.

At later times, however, the free carrier density dropped low enough for simple bimolecular recombination to become the dominant process. If the model only included free carriers then the shape of this tail region would allow the fitting algorithm to converge towards a single solution, but there were also excitons to take into account. There was no way to distinguish between photons from free carriers and photons from excitons in the experimental data. By varying the initial exciton population, n_{x0} , as well as the free carrier parameters the fitting algorithm was able to make quite a wide variety of parameter sets produce almost identical curves. Fig. 6.18 shows an example of this trade-off between free carrier and exciton emission in the model.

Another way of looking at this problem is to consider the shape of the error function. The asymptotic behaviour manifests as an infinitely long, shallow valley in the 6 dimensional parameter space. Somewhere in this valley is the global minimum, but the gradient in the surrounding area could be arbitrarily close to zero. If the error function was exact then this would only be a slight problem, but the error function was based on noisy data and hence had an associated uncertainty. The shallow shape of the valley results in a very large uncertainty in the position of the global minimum.

All is not lost, however. As we will see later, there are still useful pieces of information which are also invariant under that scaling, such as the ratio k_{F3}/k_{F2} . These do not suffer from the same uncertainty. Additionally, even some of the




Figure 6.18: A plot of the 5 different fitting results from the 20% Br sample at 200 nW excitation power. The green curve is the total PL emission calculated by the model. The blue and red curves show only the PL from excitonic and free carrier decay respectively. The green curves are all nearly identical, but there's significant variation in the blue and red curves.

parameters which do vary (such as the free carrier to exciton emission ratio) still have relatively small uncertainty.

The other limitation comes from the difficulty of defining the model parameters in terms of physically meaningful units. The recorded TRPL data is an accurate representation of the photons detected by the PMT, but it only represents a small fraction of the light emitted from the sample. The optical losses were not only dependent upon the optics of the measurement system but were also affected by the precise geometry of the perovskite-air interface. Furthermore, the absorption of the excitation laser was not known with any great precision, nor was the thickness of the perovskite layer at the measurement location. Without all of these pieces of information the fitting was limited to use arbitrary units for the photon counts and carrier concentrations. This meant that the only physically meaningful parameters which could be extracted from the model were dimensionless ratios or ones with units of time. It was still possible to identify trends in other parameters, as long

Description	Expression	Dimensions
Exciton decay rate constant	k_{x1}	$[\text{Time}]^{-1}$
Trap-assisted decay rate constant	k_{f1}	$[\text{Time}]^{-1}$
Ratio of free carrier and Auger recombination rate constants	$\frac{k_{f3}}{k_{f2}}$	$[\text{Length}]^3$
Ratio of free carrier and exciton radiative recombination rates	$\frac{k_{f2}n_f^2}{k_{x1}n_x}$	dimensionless
Ratio of Auger and trap-assisted non-radiative recombination rates	$\frac{k_{f3}n_f^3}{k_{f1}n_f}$	dimensionless

Table 6.2: A list of the physically meaningful values which could be calculated from the theoretical model parameters after fitting. All of these parameters are invariant under the transformation by the scaling factor, s .

as they were invariant under the previously discussed transformation, because a trend is nothing more than a series of ratios.

Working within these limitations, five values were identified as having fairly well defined values. These are listed in Table 6.2, along with their dimensions.

6.3.4 Lifetime Fitting Results

First of all, from Fig. 6.15 we can see that the model is a very good fit to the data, suggesting that it is a fairly accurate description of the underlying carrier dynamics. Models which did not include any excitonic contribution did not produce good fits to the data, showing that both excitons and free carriers play an important role in photoluminescence from these perovskites under the experimental conditions used in this work.

The trap-assisted recombination rate was found to be negligible in all of the samples studied. The effective lifetime of free carriers due to trap-assisted recombination, given by $1/k_{f1}$, was always on the order of 1 ms or more. Because




Figure 6.19: Plots of values derived from the model parameters which were fitted to experimental TRPL data. All measurements were made at a nominal temperature of 4.2 K, but at high laser powers the local temperature may have been higher. The colours of the lines are the same in the two plots. (a) is a lin-log plot of the exciton lifetime (inverse of k_{x1}). (b) is a log-log plot of the ratio between the Auger rate constant and the bimolecular free carrier rate constant.

this decay process was so insignificant, the fitting algorithm produced a large spread of possible values making it difficult to identify any trends with respect to composition or excitation power.

In contrast, quite clear trends could be identified in the exciton decay rate parameter, k_{x1} . The inverse of that decay rate, i.e. the exciton lifetime, is plotted

against excitation power in Fig. 6.19(a). The error bars indicate the standard deviation across different fitting runs on the same experimental data. While there is not much variation with power in the 20% and 40% Br samples, the exciton lifetime in the 10% Br sample decreases logarithmically as the power is increased. In all cases the exciton lifetime is very long.

The reason any of these parameters vary with respect to laser power is most likely due to local heating of the sample surface. While the exact relationship between laser power and temperature is not known, higher laser powers can be assumed to correspond to higher temperatures within the excitation volume. An estimate of the temperature change could be extracted from the μ PL spectra recorded at each power. The centroid wavelength was calculated from each spectrum, and they were found to only change by about 5 nm as the power was increased from 25 nW to 1600 nW. Comparing this to the spectra recorded at different temperatures in Fig. 6.1, we can say with confidence that the temperature was below 50 K at all times. Interestingly the wavelength shift with laser power follows the same logarithmic relationship as the free carrier lifetime. Given that the wavelength shift appears to be roughly linear with temperature, this suggests that the exciton lifetime may also decrease linearly with increasing temperature in the 10% Br sample.

Moving on to the free carrier decay coefficients, the ratio of the Auger rate constant to the bimolecular rate constant is plotted in Fig. 6.19(b). Interestingly the 20% Br sample appears to have generally lower values of this ratio than either of the others. There is a clear trend of decreasing Auger rate constant with respect to laser power, suggesting that heating the samples actually improves their luminescence efficiency.

Plotted in Fig. 6.20 is the total number of photons emitted from exciton decay divided by the total number emitted from free carrier recombination. This was calculated by breaking the model down into the two independent carrier populations, calculating the PL from each following a single excitation pulse and then integrating over a period of 1000 ns of decay. Values > 1 indicate that the emission is dominated by excitons while values < 1 indicate that free carrier emission dominates.




Figure 6.20: Log-log plot of the ratio between the total number of photons emitted from exciton decay and the number emitted from free carrier recombination during one excitation pulse and subsequent decay. The error bars show the standard deviation across multiple fitting runs.

We can see from this that, once again, there isn't a simple linear trend with respect to bromine fraction. Emission from the 20% Br sample is dominated by free carriers at all powers while excitons always dominate in emission from the 10% Br sample. The decrease in the exciton-to-free-carrier ratio with increasing laser power isn't surprising, since excitons tend to dissociate into free carriers at higher temperatures.

6.3.5 Measuring the Diffusion Coefficient

It was also possible to combine the scanning mirror capability of the system with the TRPL measurement capability. The sample could be excited in one location while TRPL measurements were made in a different location. While this was a significantly slower procedure than time-integrated μ PL mapping, it had the advantage of providing both spatial and temporal information.

By measuring the spatial distribution of the PL emission as a function of time,




Figure 6.21: Plots of the PL emission as a function of time and distance from the excitation source. Each line is a TRPL trace recorded with a different separation between excitation and collection spots. Separation distances are listed in the legend.

it is possible to directly observe the carrier diffusion process and extract an estimate of the diffusion coefficient. This kind of spatio-temporally resolved approach is an established technique which has been used to measure diffusion coefficients in many different materials [215, 216, 217] including perovskites [218].

The carrier distribution can be modelled as a Gaussian with FWHM w which expands by diffusion according to equation (6.7) [217],

$$\rho(r, t) \propto \left(\frac{w_0^2}{w^2} \right) \exp \left(\frac{-4 \ln(2) r^2}{w^2} - \frac{t}{\tau_r} \right) \quad (6.7)$$

where τ_r is the radiative lifetime. The time evolution of w is described by the equation,

$$w^2 = w_0^2 + 16 \ln(2) D t \quad (6.8)$$

in which D is the diffusion coefficient. Note that the radiative lifetime does not appear in equation (6.8) so it shouldn't matter that these perovskite samples exhibit a more complicated carrier decay behaviour.

TRPL measurements were carried out on the 10% bromine sample at a temperature of 4.2 K. The excitation power was kept at 33 nW while the laser spot




Figure 6.22: Plots of the PL emission as a function of time and distance from the excitation source. The filled circles show the measured data, while the dashed lines are Gaussian fits to the data.

was moved away from the collection spot in increments of $1\ \mu\text{m}$. At each location TRPL data was collected for 300 s, integrating over the wavelength range from 810 nm to 820 nm. Some of these TRPL data are plotted in Fig. 6.21.

Fig. 6.22 shows some examples of the spatial PL profiles which were extracted by taking slices from the TRPL dataset at particular times. Gaussian functions (dashed lines) have been fitted to each set of data. The squared widths of these Gaussians have been plotted as a function of time in Fig. 6.23.

Equation (6.8) is linear with respect to w^2 , so by fitting a straight line to the data in Fig. 6.23 it is possible to find a value for D from the gradient.

It is reassuring that the values of w^2 mostly lie quite close to the best fit line. This suggests that there is little uncertainty in the values of w despite the small number of data points used for the Gaussian fits in Fig. 6.22. The uncertainty in the linear fit was quantified by taking the square roots of the diagonal terms in the covariance matrix.

The ambipolar diffusion coefficient was calculated to have a value of $1.21 \pm 0.07\ \text{cm}^2\ \text{s}^{-1}$. This is very close to the values found for MAPbBr₃ (1.59 to $2.41\ \text{cm}^2\ \text{s}^{-1}$)




Figure 6.23: Plot of the squared FWHM, w^2 , of the spatial PL distribution as a function of time. The values of w were extracted from Gaussian fits to the data plotted in Fig. 6.22. The orange line is a linear least-squares fit to the data points.

and MAPbI₃ (0.50 to 1.44 cm² s⁻¹) nanocrystals at room temperature [218]. This is also similar to the diffusion coefficient found for FAPbBr₃ bulk single crystals at room temperature (0.40 cm² s⁻¹) [219].

By applying the Einstein relation given in equation (6.3) with a temperature of 4.2 K we get a carrier mobility of $(3.3 \pm 0.2) \times 10^3 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$. Very few studies of carrier mobilities in hybrid perovskites have been carried out at such a low temperature, so comparison with other results is difficult. Extrapolating the T^{-1} dependence of mobility found by Zhi-Gang Yu *et al.* [220] in certain MAPbI₃ samples gives a value of around $3 \times 10^4 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$ at 4.2 K which is an order of magnitude higher than our result. Other groups report a $T^{-3/2}$ temperature dependence of mobility, and extrapolating from measurements of the mobility in MAPbI₃ [221] and FAPbI₃ [222] yields mobilities of $1 \times 10^4 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$ and $1.5 \times 10^3 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$ respectively. The latter of these is both more chemically similar to our own samples and is in much closer agreement with our result, but there is still a factor of 2 difference. For comparison, McMeekin *et al.* measured a mobility of $21 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$ in FA_{0.83}CS_{0.17}Pb(I_{0.6}Br_{0.4})₃ at room temperature [141].

There is an additional complication introduced by the fact that the excited carriers may not be at the same temperature as the rest of the sample. Not only would there be some small amount of local heating from the laser pulses, but the short delay between excitation and measurement (< 2 ns) may not give the carriers enough time to reach thermal equilibrium with the lattice. Because the temperature is so low, even an error of just a few Kelvin would produce a very large change in the calculated mobility. If present, this effect would cause us to underestimate the temperature (since 4.2 K is the minimum temperature achievable with the cryostat used in this experiment) and hence overestimate the mobility. Without a more accurate way to estimate the uncertainty in the carrier temperature, this result should be treated with a degree of caution.

6.4 Conclusions

Four different perovskite compositions were studied in this chapter, all based on the formula $\text{FA}_{0.83}\text{Cs}_{0.17}\text{Pb}(\text{I}_{1-x}\text{Br}_x)_3$ with $x \in \{0.1, 0.2, 0.3, 0.4\}$. Hybrid lead-halide perovskites in general are promising candidates for photovoltaic applications, and these particular materials show improved stability over a wider range of compositions and temperatures.

Micro-photoluminescence spectroscopy measurements were carried out over a range of temperatures from 4.2 K to 300 K. A positive correlation between band gap energy and temperature was identified, which is the opposite of the usual trend in most semiconductors but has been observed previously in other perovskites. The trend of increasing PL intensity with decreasing temperature reversed at temperatures below 50 K in the 10% and 20% bromine samples, suggesting the possibility of a phase change in this region.

Two types of μPL maps were performed both at cryogenic temperatures and at room temperature. The objective maps, in which excitation and collection spots were moved together, showed evidence of halide segregation (or possibly cation segregation). This segregation was more pronounced at 4.2 K than at 300 K and occurred over smaller length scales in samples with higher bromine fraction.

Mirror maps, in which the collection spot remained fixed while the excitation laser was scanned, allowed us to probe the carrier transport dynamics within these perovskites. In addition to the expected carrier diffusion process, photon recycling was also found to contribute to non-local photoluminescence. The carrier diffusion process was found to have a range of several microns, and this range increased with temperature in all four samples. No significant variation could be detected between the different sample compositions. A diffusion length on the order of microns would be very useful for efficiently extracting carriers from a photovoltaic cell.

Time-resolved micro-photoluminescence measurements were also performed on three of the samples (10%, 20%, and 40% Br) at cryogenic temperatures. A sophisticated theoretical model was developed which included contributions for monomolecular exciton decay as well as trap-assisted, bimolecular, and Auger recombination of free carriers. Distortions introduced by the detection apparatus were also modelled in the computer code. This model was fitted to the data and was found to accurately reproduce all the observed features.

The fitting algorithm was unable to find unique values for all of the parameters because the model was almost invariant under a particular parameter transformation, but some useful information could still be extracted. Trap-assisted recombination was found to play a negligible role with an effective lifetime of 1 ms or more across all samples and excitation powers. Excitonic decay played a more significant role, dominating emission from the 10% bromine sample at all excitation powers and from the other two samples at low powers. The exciton lifetime was found to be constant at around 100 ns in the 20% and 40% samples, and rose to 400 ns in the 10% sample at the lowest excitation power. Using excitation power as a proxy for temperature, the exciton lifetime was estimated to decrease linearly with increasing temperature. The Auger rate coefficient also showed signs of decreasing with increasing temperature, although more detailed measurements are needed to confirm that.

These experiments only scratched the proverbial surface, and there is a lot more to be done to fully understand these materials. Limited evidence of a phase change in the 4.2 K to 50 K region has been observed, so more detailed μ PL studies

over this temperature range would be worthwhile. We are also working on extending the TRPL measurements up to higher temperatures to gain a better understanding of the temperature dependence of these parameters.

7

Conclusions

7.1 PhC microcavities

In chapters 2 and 5 we discussed, respectively, FDTD simulations and experimental fabrication of photonic crystal microcavities composed of SU-8 photoresist patterned on top of GaAs PhC waveguides. The simulations yielded predictions for the frequency, Q factor, and mode volume of these cavities, as well as shedding light on the factors limiting the Q factor. More specifically, the breaking of the z -axis mirror symmetry by the addition of the SU-8 disk meant that the optical modes of the system could no longer be categorised into independent TE and TM modes. Light could leak from the TE mode into the TM mode, then radiate into the surrounding PhC due to the absence of a TM optical band gap.

A procedure for fabricating these cavities was developed, using a 405 nm continuous wave laser to expose the photoresist. 1D μ PL maps performed before and after exposure were used to verify the successful creation of spatially localised cavity modes. The success rate was 34 devices out of an attempted 40, which is quite reasonable for a newly developed process. Higher spectral resolution measurements were used to measure the Q factors of the cavities, and they were found to range from slightly over 2×10^3 up to a maximum of 7.4×10^3 . This maximum Q factor

is comparable to the Q factors of Notomi and L3 cavities fabricated via e-beam lithography on similar samples of GaAs.

Interestingly, the highest Q factor measured experimentally (7.4×10^3) was slightly higher than the Q factor predicted by the FDTD simulations (7.2×10^3). Normally the experimental Q factors are significantly lower than their theoretical counterparts due to imperfections in the fabrication. The explanation for this is two-fold: first, as mentioned earlier, the theoretical Q factor is limited by z -axis symmetry violation rather than by the finite size of the lattice. Secondly the size and shape of the SU-8 disks in the real devices varied slightly from the theoretical model. Many of the disk diameters were slightly larger than the $1\text{ }\mu\text{m}$ target, and disk diameter been shown to be positively correlated with Q factor.

These cavities were coupled to an ensemble of InGaAs quantum dots rather than single QDs, but the optical nature of the fabrication process has the potential to overcome one of the main difficulties in achieving on-chip dot-cavity coupling. When using e-beam lithography to fabricate PhC cavities it is very difficult to locate embedded QDs in order to achieve good spatial and spectral mode overlap. In contrast, an optical μPL system can locate embedded QDs to an accuracy of $\pm 25\text{ nm}$ using super-resolution microscopy techniques [187], then the same apparatus can immediately be used for cavity fabrication. As of the writing of this thesis we are still in the process of acquiring sufficiently low density QD samples in order to test coupling to a single QD.

The wavelength of these cavities was around $1.25\text{ }\mu\text{m}$ which is only very slightly short of the telecoms O-band (1.26 to $1.36\text{ }\mu\text{m}$) - a particularly low-loss region of the spectrum suitable for long-distance transmission through optical fibres. If coupling to single QDs is achieved, these devices could potentially be used as single photon sources for quantum key distribution systems. The small mode volume predicted by the FDTD simulations ($1.44(\lambda_0/n)^3$) is also very promising for achieving a high Purcell factor, enabling brighter single photon emission and reducing jitter.

These results suggest that there's still a lot of scope for further refinement of this cavity design. Since the completion of this work, even higher Q factors

(up to 8.6×10^3) have been achieved and the reliability has also been improved significantly [223]. The obvious next step is to acquire a wafer with lower QD density ($\leq 10 \mu\text{m}[-2]$) and attempt to couple an optically-defined cavity to a single dot. The problem of imperfect spatial alignment between the probe and exposure beams can be overcome by using a beam splitter integrated into an optical fibre to merge the two lasers into a single beam. Temperature control could be used to fine-tune the spectral overlap between the cavity and the QD, but it would also be worth exploring other options for adjusting the cavity mode or QD emission wavelength on a per-device basis.

More ambitiously there are many interesting experiments which could be performed using multiple coupled dot-cavity systems. For example by coupling two cavities together, one of which is also coupled to a quantum dot, it is theoretically possible to achieve extremely strong photon anti-bunching with only a modest dot-cavity coupling strength via the unconventional photon blockade process [224]. It would also be interesting to investigate the possibility of creating networks of these cavities in which photons are guided from one device to the next via waveguides. If the cavity-to-waveguide coupling efficiencies are high enough then it could allow for the creation of on-chip quantum optical networks.

7.2 Hybrid Perovskites

Chapter 6 covered experimental studies of mixed-cation hybrid metal-halide perovskites. Four different polycrystalline samples were studied with compositions $\text{FA}_{0.83}\text{Cs}_{0.17}\text{Pb}(\text{Br}_x\text{I}_{1-x})_3$ where $x \in \{0.1, 0.2, 0.3, 0.4\}$. Two different measurement techniques were used: time-integrated μPL spectroscopy using a 532 nm CW laser, and time-resolved μPL using a pulsed 450 nm laser. Measurements were made both of the spatial structure of the PL emission and carrier diffusion, and of the temporal dynamics of carrier decay.

The diffusion profile was measured via μPL mapping, showing evidence of photon recycling in addition to conventional carrier diffusion. The diffusion length was estimated to be around $2 \mu\text{m}$ at 4.2 K increasing to $5 \mu\text{m}$ at room

temperature. This is similar to the values measured for other polycrystalline perovskites. No significant difference between the sample compositions could be measured. A measurement was also made of the ambipolar diffusion coefficient for the $x = 0.1$ sample ($D = 1.21 \pm 0.07 \text{ cm}^2 \text{ s}^{-1}$), and this value was within the range of reported diffusion coefficients for single crystal MA-based and FA-based lead-halide perovskites.

TRPL measurements were carried out on the $x = 0.1$, 0.2, and 0.4 samples at a temperature of 4.2 K. A theoretical model was developed to fit the data which included both excitons and free carriers as well as Auger recombination and trap-assisted recombination terms. Trap-assisted recombination was found to be negligible, having an effective lifetime on the order of milliseconds. At low excitation powers (around 25 nW) the excitonic decay dominated the luminescence, and excitons were found to contribute more to emission from the 10% bromine sample than the others, suggesting a higher exciton binding energy in that composition. At higher excitation powers the emission is dominated by free carriers in the other two samples.

These results, particularly the carrier diffusion measurements, reaffirm that FA-based hybrid metal-halide perovskites are promising candidates for photovoltaic applications. The long diffusion length makes it relatively simple to design devices to efficiently extract photogenerated carriers, while the dominance of free carrier emission bodes well for their performance at room temperature. The very low rate of trap-assisted recombination also hints at the potential for high efficiency. That said, the measurements shown here are only a small contribution to the body of knowledge on these materials and there is still much to be done in the nascent field of perovskite photovoltaics.

In the short term there are plans to extend the study of these materials to higher temperatures, allowing the carrier dynamics model and its parameters to be tested more thoroughly. There are also plans to study samples with larger grain sizes and added passivation layers to investigate what effect those parameters have on the optical and electronic properties. The μPL spectra recorded at different

temperatures show signs of a phase transition between 4 K and 50 K, so it would be worthwhile to measure how the electronic and optical properties vary over that transition. Performing more TRPL measurements at lower laser powers (< 100 pW) would also be useful, since that would more accurately reflect the illumination experienced by practical photovoltaic systems and would further reduce the magnitude of local heating.

Speaking more generally, the main hurdle which perovskite photovoltaics must overcome remains their relative instability compared to conventional semiconductors. For the low manufacturing costs to translate into low energy costs, perovskite-based solar cells must be able to survive for years in ambient conditions. There have been huge leaps forward in stability over the last few years, but there are still significant issues with temperature and moisture sensitivity which hurt the commercial viability of these materials. One possible alternative which has recently been receiving increasing research interest is the field of inorganic perovskites. Compositions such as CsPbX_3 exhibit many of the same electronic properties which make hybrid perovskites so appealing for photovoltaic applications, but also have greatly increased moisture and temperature stability [225, 226, 227].

7.3 Closing Thoughts

As a whole, the research undertaken in this thesis reaffirms the effectiveness and versatility of optical techniques, both for studying the fundamental properties of matter and as an essential part of cutting-edge technologies. The optical micro-cavity design which we have developed could be used to produce light with strictly quantum photon statistics for use in quantum cryptography or could even form part of an on-chip quantum computer. Hybrid perovskites such as those studied in chapter 6 have huge potential as low-cost photovoltaic materials and could soon be helping our species make the urgently needed transition away from fossil fuels. There is also a certain beauty to be found in the physics of light, and the symmetry between photonic crystals and electronic semiconductors is just one example of the deep mathematical elegance of the universe.

Appendices



Experiment Control Software

This appendix includes some additional details of the software which was developed to control the experimental apparatus and perform different types of measurements.

A.1 Spectrometer Control Software

Over the course of this research, bespoke software was developed to link together many of the electronic elements in the optical system. This allowed some repetitive tasks (eg. acquiring background spectra) to be reduced to a single button click, and provided fully automatic mapping capabilities in a number of different modes. The software was written using the LabView graphical programming language, so unfortunately it's not possible to provide full code listings here. Fig. A.1 does however show a screenshot of the front panel of this software.

size	encoding	description
8 bytes	ASCII	File signature: "BinspcBE"
4 bytes	unsigned integer	Length of metadata section, m
m bytes	UTF-8	Metadata in JSON format
4 bytes	unsigned integer	Number of detector pixels, n
$4n$ bytes	32 bit floating points	1D Array of wavelengths in nm
$4n$ bytes	32 bit signed integers	1D Array of counts per pixel

Table A.1: Binary file layout for the "Binspec" format designed to store a single spectrum. All binary numeric values are stored in big-endian (BE) byte order.




Figure A.1: A screenshot of the user interface of the software used to record μ PL spectra and perform maps.

size	encoding	description
8 bytes	ASCII	File signature: "BinmapBE"
4 bytes	unsigned integer	Length of metadata section, m
m bytes	UTF-8	Metadata in JSON format
4 bytes	unsigned integer	Number of detector pixels, n
$4n$ bytes	32 bit floating points	1D Array of wavelengths in nm
4 bytes	unsigned integer	Number of rows in the map, h
4 bytes	unsigned integer	Number of columns in the map, w
4 bytes	unsigned integer	Number of detector pixels, n
$4hwn$ bytes	32 bit signed integers	3D Array of counts per pixel

Table A.2: Binary file layout for the "Binmap" format designed to store hyperspectral maps. All binary numeric values are stored in big-endian (BE) byte order. In the 3D array of counts, the count for row y , column x , and pixel p would be found at relative offset $4(ywn + xn + p)$. In numpy-style indexing, this would be `array[y,x,p]`.

Custom binary file formats were also developed for efficient saving and loading of spectrometer data. By using a 32 bit signed binary representation instead of the ASCII decimal format which was previously in use, the file size could be reduced by roughly a factor of two. This didn't make much difference for single spectra which were only a few kilobytes in size, but it did have a significant impact on large maps which could routinely reach a size of several hundred megabytes. The other advantage was the inclusion of metadata about the spectrometer configuration and acquisition settings. Rather than manually recording the exposure time, grating, and other information in the filename or in a separate text file, this metadata could instead be automatically embedded into the file header itself.

The bitwise layout of these two formats is given in tables A.1 and A.2. JSON (javascript object notation) was chosen as a metadata format both because there are open source JSON libraries available for every major programming language and because, when viewed with a text editor, it is human-readable. Full documentation of the file format is included in the metadata section of every file. This provides future-proofing by guaranteeing that anyone who encounters a file in this format also has instructions detailing how to read it.

A.2 Position Optimisation Pseudocode

Below is pseudocode for the automatic position optimisation algorithm. This was used to automatically find the position corresponding to maximum μ PL intensity when looking at cavities and single quantum dots. It is essentially a hill-climbing algorithm which iterates through the four cardinal directions until no further improvements can be made. The stepsize was typically set to around $0.2\text{ }\mu\text{m}$.

```
direction = 0  # degrees
old_intensity = measurePL()
halt_count = 0
while halt_count < 8:
    move(direction, stepsize)
    new_intensity = measurePL()
    if old_intensity > new_intensity:
        move(direction + 180, stepsize)  # go back
        direction = direction + 90  # degrees
        halt_count = halt_count + 1
    if new_intensity > old_intensity:
        old_intensity = new_intensity
        halt_count = 0
```

B

FDTD Simulation Code

Here I provide code listings for some of the simulations performed in chapter 2. All of these are coded in a dialect of the programming language Scheme (which itself is a dialect of Lisp) and designed to work with MEEP version 1.1. To run these on a multi-core linux system, use a command such as:

```
mpirun -n 8 meep-openmpi code.ct1 | tee output.txt
```

Parts of this code make extensive use of tail recursion in Scheme. For a detailed explanation of tail recursion, I recommend reading the first chapter of *Structure and Interpretation of Computer Programs* by H. Abelson *et al.*[228].

B.1 Band Structure of a hexagonal PhC

```

; this is in case we want to convert to SI units at any point
(define a 340e-9) ; (units of metres)

; All the following lengths are in units of a
(define struct_d (/ 200 340))
(define struct_r 0.27)
(define struct_pad 2)
(define struct_pml 1)

; PhC size
(define phc_size_x 0.5) ; width in units of 2a
(define phc_size_y 1) ; height in units of sqrt(3)a

; material definitions
(define Si (make dielectric (epsilon 12.14)))
(define SiO2 (make dielectric (epsilon 2.089)))
(define GaAs4K (make dielectric (epsilon 11.089)))

; source parameters
(define source_component Hz)
(define-param source_f 0.3) ; pulse center frequency
(define-param source_df 0.6) ; pulse width (in frequency)

; These are some constants which you shouldn't change
(define y_squash (/ (sqrt 3) 2))

; Important points in reciprocal space
; The following page gives a very good derivation of these:

```

```

; http://alphard.ethz.ch/Smajic/PCH000.htm
(define kvec_G (vector3 0 0)) ; Gamma
(define kvec_M (vector3 0 (/ 1 (sqrt 3)) 0)) ; M
(define kvec_K (vector3 (/ 1 3) (/ 1 (sqrt 3)) 0)) ; K

; set up the simulation volume
(define sim_w 1)
(define sim_h (* 2 y_squash))
(define sim_d (+ struct_d (* 2 (+ struct_pad struct_pml))))
(set! geometry-lattice (make lattice (size sim_w sim_h sim_d)))
(set! resolution 20)

; add PML layers
(set! pml-layers (list (make pml (thickness struct_pml) (direction Z))))

; symmetries
; Mirror phase:
; 1 => non-zero H field
; -1 => non-zero E field
; In this case we can't use XY mirror symmetry because the
; phase shift at the boundary will violate it for most k-points
(set! symmetries (list
  (make mirror-sym (direction Z) (phase 1))
))

; this is where we define our PhC
(define (phc_iter_x i j end geom_list)
  (cond ((> i end) geom_list)
        (else (phc_iter_x (+ i 1) j end
                          (cons

```

```

    (make cylinder
      (center (+ i (/ j 2)) (* j y_squash))
      (height struct_d)
      (radius struct_r)
      (material air))
    geom_list
  )
))
)
)

```

```

(define (phc_iter_y nx j end geom_list)
  (cond ((> j end) geom_list)
        (else (phc_iter_y nx (+ j 1) end
                          (phc_iter_x
                           (ceiling (- (- (/ j 2)) nx))
                           j
                           (floor (+ (- (/ j 2)) nx))
                           geom_list)
                          )
        ))
)
)

```

```

(define (phc nx ny geom_list)
  (phc_iter_y nx (- ny) ny geom_list)
)

```

; Here we use the above functions to generate our PhC holes
; and add other things such as the slab and substrate

```

(set! geometry (append (list
  ; things to add before the PhC holes go here
  (make block ; GaAs slab
    (center 0 0 0)
    (size infinity infinity struct_d)
    (material GaAs4K))
  )
  (phc phc_size_x phc_size_y (list))
  (list
    ; things to add after the PhC go here
  )))

; Set up sources
(define (rand-source)
  ; creates a random source with random phase somewhere in the sim volume
  (make source
    (src (make gaussian-src
      (frequency source_f)
      (fwidth source_df)
    ))
    (component source_component)
    (center
      (* sim_w (- (random 1.0) 0.5))
      (* sim_h (- (random 1.0) 0.5))
      0
    )
    (amplitude (exp (* 0+1i (random (* 2 pi)))))
  )
)

(define (rand-sources n slist)

```



```

; this takes a list of harminv points in the -y half and
; makes symmetrical copies in the +y half of the sim volume
; [a,b,c,...] => [(a1,a2),(b1,b2),(c1,c2),...]
(map (lambda (v) (cons
  v
  (vector3
    (cond
      ((< (vector3-x v) 0) (+ (vector3-x v) (/ sim_w 2)))
      (else (- (vector3-x v) (/ sim_w 2)))
    )
    (+ (vector3-y v) (/ sim_h 2))
    0
  )
)) hlist)
)
(define hmv-points (copy-hmvps (rand-hmvps 10 '())))

; we set the k-vector(s) we're interested in
(define k_points (interpolate 19 (listkvec_G kvec_M kvec_K kvec_G)))
(print "k points:" k_points "\n")

(begin (init-fields) (output-epsilon))
(map (lambda (k)
  (change-k-point! k)
  (restart-fields)
  ; output the expected phase shift between the two points
  (map (lambda (hmv)
    (print "phase shift: "
      (* 2 pi (vector3* k (vector3- (cdr hmv) (car hmv))))) "\n"
    )
  )
)

```

```
) hmv-points)
; The (map ...) returns a list of lists of harminv callables.
; The (apply append ...) turns that into a single list.
; The (apply after-sources ...) sets all of those harminv callables
; to run after the sources.
(run-sources+ 1000 (apply after-sources (apply append
  (map (lambda (hmv) (list
    (harminv source_component (car hmv) source_f source_df)
    (harminv source_component (cdr hmv) source_f source_df)
  )) hmv-points)
)))
) k_points)
```

B.2 SU-8 Cavity Q Factor

```

; this is in case we want to convert to SI units at any point
(define a 340e-9) ; (units of metres)

; These are some constants which you shouldn't change
(define y_squash (/ (sqrt 3) 2))

; some settings
(define-param subpixel-averaging? true)
(set! filename-prefix false)
(set! eps-averaging? subpixel-averaging?)
(set! Courant 0.5)

; All the following lengths are in units of a
(define struct_d (/ 200 340))
(define struct_r 0.27)
(define struct_pad_xy 0.5)
(define struct_pad_z 2)
(define struct_pml 1)
(define su8_thick (/ 100 340))
(define su8_r (/ 500 340))

; Parameters for the waveguide
(define use_wg true)
; Space between centres of holes = wg_width*sqrt(3)*a
(define wg_width (cond
  (use_wg 0.98)
  (else 1)
))

```

```

; PhC size
(define phc_size_x 23.5) ; number of units of a on each side of the middle
(define phc_size_y 12)   ; number of rows above/below the middle

; material definitions
(define Si (make dielectric (epsilon 12.14)))
(define SiO2 (make dielectric (epsilon 2.089)))
(define GaAs4K (make dielectric (epsilon 11.089))) ; n=3.33
(define SU8 (make dielectric (epsilon 2.465))) ; n=1.57

; source parameters
(define source_component Ey)
; When using xy mirror symmetry, the negative quadrant is the only real one.
; Sources in the other quadrants don't get simulated.
(define source_pos (vector3 -0.1 -0.1 0))
(define source_f 0.267) ; eps averaging on, res>5
(define source_df 0.001) ; pulse width (in frequency)
(define source_end (/ 10 source_df))

; set up the simulation volume
(define sim_w
  (* 2 (+ phc_size_x struct_r struct_pad_xy struct_pml))
)
(define sim_h
  (* 2 (+ (* y_squash phc_size_y) struct_r struct_pad_xy struct_pml))
)
(define sim_d
  (+ struct_d (* 2 (+ struct_pad_z struct_pml)))
)

```

```

(set! geometry-lattice (make lattice (size sim_w sim_h sim_d)))
(set! resolution 20)

;add PML layers
(set! pml-layers (list (make pml (thickness struct_pml))))

; symmetries
; Mirror phase:
;   1 in X => at x=0, Hx, Ey, and Ez components are non-zero
;  -1 in X => at x=0, Ex, Hy, and Hz components are non-zero
(set! symmetries (cond
  (subpixel-averaging? (list
    (make mirror-sym (direction X) (phase 1))
  ))
  (else (list
    (make mirror-sym (direction X) (phase 1))
    (make mirror-sym (direction Y) (phase -1))
  ))
))

; This is where we define our PhC using tail recursion.
(define (phc_iter_x i j end geom_list)
  (cond ((> i end) geom_list)
        ((and use_wg (= j 0)) (phc_iter_x (+ i 1) j end geom_list))
        (else (phc_iter_x (+ i 1) j end
          (cons
            (make cylinder
              (center
                (+ i (/ j 2)) ; x
                (* y_squash (+ j (cond

```

```

        ((= 0 j) 0)
        (else (* (- wg_width 1) (/ j (abs j)))))
    ))) ; y
    0 ; z
  )
  (height struct_d)
  (radius struct_r)
  (material air))
  geom_list
)
))
)
)

```

```

(define (phc_iter_y nx j end geom_list)
  (cond ((> j end) geom_list)
        (else (phc_iter_y nx (+ j 1) end
                          (phc_iter_x
                           (ceiling (- (- (/ j 2)) nx))
                           j
                           (floor (+ (- (/ j 2)) nx))
                           geom_list)
                          )
        ))
  )
)
)

```

```

(define (phc nx ny geom_list)
  (phc_iter_y nx (- ny) ny geom_list)
)

```

```

; Here we use the above functions to generate our PhC holes
; and add other things such as the slab and substrate
(set! geometry (append (list
  ; things to add before the PhC holes go here
  (make block ; GaAs slab
    (center 0 0 0)
    (size infinity infinity struct_d)
    (material GaAs4K))
  )
  (phc phc_size_x phc_size_y (list))
  (list
    ; things to add after the PhC go here
    (make cylinder ; SU-8 disk
      (center 0 0 (/ (+ struct_d su8_thick) 2))
      (height su8_thick)
      (radius su8_r)
      (material SU8))
    )))

; define some functions to mirror the sources correctly
(define (mirror-pos-x poslist)
  (apply append (map (lambda (pos)
    (cond ((= (vector3-x pos) 0) (list pos))
      ; a list of pos and its reflection in x
      (else (list pos (vector3- pos (vector3 (* (vector3-x pos) 2) 0 0))))
    )
  ) poslist))
)

(define (mirror-pos-y poslist)

```

```

(apply append (map (lambda (pos)
  (cond ((= (vector3-y pos) 0) (list pos))
        ; a list of pos and its reflection in y
        (else (list pos (vector3- pos (vector3 0 (* (vector3-y pos) 2) 0))))
  )
) poslist))
)

; define a function to turn a list of positions into a list of sources
(define (make-sources poslist)
  (map (lambda (pos)
    (make source
      (src (make gaussian-src
        (frequency source_f)
        (fwidth source_df)
      ))
      (component source_component)
      (center pos)
    )
  ) poslist)
)

; Set up sources
(set! sources (make-sources
  (mirror-pos-y (mirror-pos-x (list source_pos)))
))

; define the boundaries of the non-PML region
(define edge_x (+ phc_size_x struct_r struct_pad_xy))
(define edge_y (+ (* y_squash phc_size_y) struct_r struct_pad_xy))
(define edge_z (+ (/ struct_d 2) struct_pad_z))

```

```

; define some useful volumes/planes for analysis
(define vol-phc-plane (volume
  (center 0 0 0)
  (size sim_w sim_h 0)))
(define vol-su8-plane (volume
  (center 0 0 (/ (+ struct_d su8_thick) 2))
  (size sim_w sim_h 0)))
(define vol-cross-plane (volume
  (center 0 0 0)
  (size sim_w 0 sim_d)))
(define vol-integration (volume
  (center 0 0 0)
  (size (* 2 edge_x) (* 2 edge_y) (* 2 edge_z))))
(define vol-flux-x1 (volume
  (center edge_x 0 0)
  (size 0 (* 2 edge_y) (* 2 edge_z))))
(define vol-flux-y1 (volume
  (center 0 edge_y 0)
  (size (* 2 edge_x) 0 (* 2 edge_z))))
(define vol-flux-z1 (volume
  (center 0 0 edge_z)
  (size (* 2 edge_x) (* 2 edge_y) 0)))
(define vol-flux-z2 (volume
  (center 0 0 (- edge_z))
  (size (* 2 edge_x) (* 2 edge_y) 0)))

; Here we set up everything needed to find the Q factor and mode volume
(define Een 0)
(define Hen 0)
(define maxbox 0)

```

```

(define Py1 0)
(define Px1 0)
(define Pz1 0)
(define Pz2 0)
(define Qin 0)
(define Qout 0)
(define Qtotal 0)
(define modeV 0)
;  $\epsilon E^2$ 
(define (f1 r ex ey ez eps)
  (* eps (+ (* ez (conj ez)) (* ey (conj ey)) (* ex (conj ex)))))
)
(define (do-Een)
  (set! Een (+ Een (integrate-field-function
    (list Ex Ey Ez Dielectric)
    f1
    vol-integration
  )))
)
;  $H^2$ 
(define (f2 r hx hy hz)
  (+ (* hz (conj hz)) (* hy (conj hy)) (* hx (conj hx))))
)
(define (do-Hen)
  (set! Hen (+ Hen (integrate-field-function
    (list Hx Hy Hz)
    f2
    vol-integration
  )))
)

```

```

; max( $\epsilon$ * $E^2$ )
(define (do-maxbox)
  (set! maxbox (+ maxbox (max-abs-field-function
    (list Ez Ey Ex Dielectric)
    f1
    vol-integration
  )))
)

; Flux
(define (do-Px1) (set! Px1 (+ Px1 (flux-in-box X vol-flux-x1))))
(define (do-Py1) (set! Py1 (+ Py1 (flux-in-box Y vol-flux-y1))))
(define (do-Pz1) (set! Pz1 (+ Pz1 (flux-in-box Z vol-flux-z1))))
(define (do-Pz2) (set! Pz2 (+ Pz2 (flux-in-box Z vol-flux-z2))))

; Mode volume
(define (do-modeV)
  (set! modeV (meep-fields-modal-volume-in-box fields vol-integration))
)

; Output function
(define (output-integrals) (list
  (do-modeV)
  (set! Qin (/ (* pi source_f (+ Een Hen)) (+ (* 2 Py1) (* 2 Px1))))
  (set! Qout (/ (* pi source_f (+ Een Hen)) (- Pz1 Pz2)))
  (set! Qtotal (/ (* Qin Qout) (+ Qin Qout)))
  (print "=== START OF OUTPUT BLOCK ===\n")
  (print "Time = " (meep-time) "\n")
  (print "Een = " Een "\n")
  (print "Hen = " Hen "\n")
  (print "maxbox = " maxbox "\n")
  (print "Px1 = " Px1 "\n")
  (print "Py1 = " Py1 "\n")

```

```

(print "Pz1 = " Pz1 "\n")
(print "Pz2 = " (- Pz2) "\n")
(print "Qin = " Qin "\n")
(print "Qout = " Qout "\n")
(print "Qtotal = " Qtotal "\n")
(print "ModeV_meep = " modeV "\n")
(print "ModeV_calc = " (/ Een maxbox) "\n")
(print "=== END OF OUTPUT BLOCK ===\n")
))

; Run the simulation
(run-sources+ (/ 5 source_f)
  (at-beginning
    (with-prefix "phc" (in-volume vol-phc-plane output-epsilon))
    (with-prefix "su8" (in-volume vol-su8-plane output-epsilon))
    (with-prefix "cross" (in-volume vol-cross-plane output-epsilon))
  )
  (after-sources
    (harminv source_component source_pos source_f source_df)
    (at-every (/ 0.1 source_f)
      do-Een do-Hen do-maxbox
      do-Px1 do-Py1 do-Pz1 do-Pz2
    )
  )
  (after-sources+ (/ 3 source_f)
    (at-every (/ 0.05 source_f)
      (to-appended "phc-ey" (in-volume vol-phc-plane output-efield-y))
      (to-appended "cross-ey" (in-volume vol-cross-plane output-efield-y))
      (to-appended "phc-ez" (in-volume vol-phc-plane output-efield-z))
      (to-appended "cross-ez" (in-volume vol-cross-plane output-efield-z))
    )
  )
)

```

```
    )  
  )  
  (at-end output-integrals)  
)
```

C

Perovskite TRPL Fitting Code

This appendix contains listings of some of the code used to model carrier decay dynamics in the perovskite samples and to fit the model parameters to the experimental time-resolved PL data. In order to save space, most of the code related to parsing the TRPL data files has been omitted. Comments have been added where necessary to indicate the meaning of variables.

C.1 Carrier Decay Simulation Code

```

from typing import Tuple

import scipy.interpolate as interpolate
import scipy.integrate as integrate
import numpy as np
import trpl_files

edge_threshold = 0.6

def dn_by_dt(t, n, r_xt1, r_ft1, r_ft2, r_ft3):
    n_f, n_x = n
    # r_xt1 = rate of excitonic decay ( $\sim n$ )
    # r_ft1 = rate of free carrier trapping ( $\sim n$ )
    # r_ft2 = rate of free carrier bimolecular decay ( $\sim n^2$ )
    # r_ft3 = rate of free carrier auger recombination ( $\sim n^3$ )
    df_dt = - r_ft1 * n_f - r_ft2 * (n_f ** 2) - r_ft3 * (n_f ** 3)
    dx_dt = - r_xt1 * n_x
    return [df_dt, dx_dt]

def jacobian(t, n, r_xt1, r_ft1, r_tot2, r_tot3):
    n_f, n_x = n
    return np.array([
        [-r_ft1 - 2*r_tot2*n_f - 3*r_tot3*(n_f**2), 0],
        [0, -r_xt1]])

def emission(n, r_x1, r_f1, r_f2, r_f3):
    n_f, n_x = n
    return r_x1*n_x + r_f1*n_f + r_f2*(n_f**2) + r_f3*(n_f**3)

```

```

def solve_ode(n0: Tuple[float, float], r_xt1, r_ft1,
              r_ft2, r_ft3, r_xr1, r_fr1, r_fr2, r_fr3, dt):
    # x = excitons
    # f = free carriers
    # t = total rate (radiative + non-radiative)
    # r = radiative fraction of total rate (between 0 and 1)
    # 1, 2, 3 = number of excitons/carriers involved in the process
    solver = integrate.ode(dn_by_dt, jacobian)
    solver.set_integrator('dopri5')
    solver.set_initial_value(n0, 0)
    solver.set_f_params(r_xt1, r_ft1, r_ft2, r_ft3)
    solver.set_jac_params(r_xt1, r_ft1, r_ft2, r_ft3)
    ts, ys = [], []

    def solout(t, y):
        ts.append(t)
        ys.append(emission(y, r_xt1*r_xr1, r_ft1*r_fr1,
                           r_ft2*r_fr2, r_ft3*r_fr3))

    duration = 1000
    solver.set_solout(solout)
    solver.integrate(duration)

    out_times = np.arange(0, duration, dt)
    spline = interpolate.InterpolatedUnivariateSpline(ts, ys)
    counts = [spline.integral(t, t+dt) for t in out_times]
    return np.array(counts)

def convolve(func, irf):
    return np.convolve(func, irf, mode="full")

```

```

def wrap(ys, ts, offset, period):
    dt = ts[1] - ts[0]
    data = np.zeros(len(ts))
    for i, y in enumerate(ys):
        t = (offset + (i * dt)) % period
        new_i = int(0.5 + t / dt)
        if 0 <= new_i < len(data):
            data[new_i] += ys[i]
    return data

def rising_edge(data):
    peak = np.max(data.counts)
    return data.times[data.counts > peak * edge_threshold][0]

def calculate_decay(data: trpl_files.TRPL, power, irf: trpl_files.TRPL,
                    nscale_f, nscale_x, r_xt1, r_ft1, r_ft2, r_ft3):
    n_f = power * nscale_f
    n_x = power * nscale_x
    dt = data.times[1] - data.times[0]
    rep_period = data.get_meta("rep_period")
    curve = solve_ode((n_f, n_x), r_xt1, r_ft1, r_ft2, r_ft3,
                      1, 0, 1, 0, dt)
    convolved = convolve(curve, irf)
    t = data.times
    t0_guess = t[int(len(t) / 2)]
    wrapped = wrap(convolved, t, t0_guess, rep_period)
    mid_height = (np.max(wrapped) * edge_threshold) \
        + (np.min(wrapped) * (1-edge_threshold))
    for iii in range(len(wrapped)):

```

```
    if wrapped[iii - 1] <= mid_height <= wrapped[iii]:
        t_displace = t0_guess - t[iii]
        break
else:
    print("Failed to find rising edge.")
    t_displace = t0_guess
t0 = rising_edge(data) + t_displace
wrapped = wrap(convolved, t, t0, rep_period)
return wrapped
```

C.2 Parameter Fitting Code

```

from typing import List

import fit_files
import trpl_files
import decay_model
import numpy as np
from numpy import random
import scipy.optimize as optimize

def error_func(params, scale_func, bounds, data: trpl_files.TRPL,
               power: float, irf: trpl_files.TRPL):
    # regions are given as a list of (start, end) tuples in units of ns
    # if start > end, wrapping is implied
    region_bounds = ((48, 51), (51, 75), (75, 48))
    region_weights = (1, 2, 2)
    regions = []
    for tmin, tmax in region_bounds:
        if tmin >= tmax:
            region = np.nonzero(np.logical_or(data.times < tmax,
                                              data.times >= tmin))[0]
        else:
            region = np.nonzero(np.logical_and(data.times >= tmin,
                                              data.times < tmax))[0]
        regions.append(region)
    params = scale_func(params)
    for arg, lim in zip(params, bounds):
        if lim[0] is not None and arg < lim[0]:
            return np.inf
        if lim[1] is not None and arg > lim[1]:

```

```

        return np.inf

    curve = decay_model.calculate_decay(data, power, irf, *params)
    error = 0

    for region, weight in zip(regions, region_weights):
        diffs = np.log10(np.clip(curve[region], 1e-8, None)) \
            - np.log10(data.counts[region])
        error += np.mean(diffs ** 2) * weight

    return error

bounds = [(0, None), (0, None), (0, None),
          (0, None), (0, None), (0, None)]
factors = np.array([1e3, 10, 1e-2, 1e-6, 1e-6, 1e-9])

def scale_func(params):
    return params * factors

def scale_func_inv(params):
    return params / factors

def accept_test(f_new, x_new, f_old, x_old):
    for arg, lim in zip(x_new, bounds):
        if lim[0] is not None and arg < lim[0]:
            return False
        if lim[1] is not None and arg > lim[1]:
            return False
    return True

def main():
    # File loading code has been omitted for brevity.
    # "asc" is the filename of the PL spectrum data

```

```

# "path" is the filename of the TRPL data
# "background" is the TRPL data without excitation
# "irf_norm" is the normalised IRF data

wavelengths, counts = read_asc(asc)
count_total = np.sum(counts)
count_range = np.sum(counts[np.nonzero(
    np.logical_and(wavelengths >= wlmin, wavelengths <= wlmax)
)])
wl_scale_factor = count_total/count_range

data = trpl_files.load_dat(path) * wl_scale_factor / duration
# round the rep_period to the nearest multiple of 12.5 ns
# valid for any pulse-picked laser with a base rate of 80 MHz
rep_period = np.round(data.meta['rep_period'] / 12.5) * 12.5
data.meta['rep_period'] = rep_period
# crop the trailing zeros off the data
data = data.crop(0, rep_period)
# subtract the average background
data = data.sub_mean(background)

# our starting points are taken from a log distribution
starting_points = np.power(10, (random.rand(5, len(bounds))-0.5)*2)
fit_results = []
min_args = {"args": (scale_func, bounds, data,
                    power, irf_norm.counts)}
for sp in starting_points:
    print("Starting from:", sp)
    # this is the main fitting function we use
    result = optimize.basinhopping(error_func, sp,

```

```
        minimizer_kwargs=min_args,  
        accept_test=accept_test,  
        callback=found_min,  
        stepsize=1,  
        niter=100)  
  
    fit_results.append(scale_func(result.x))  
  
# The code for storing and plotting fit_results  
# has been omitted for brevity.
```

Bibliography

- [1] C. Kittel, in *Introduction to Solid State Physics*, 8th edition ed. (John Wiley & Sons, 2005), Chap. 7, p. 173. 7
- [2] C. Kittel, in *Introduction to Solid State Physics*, 8th edition ed. (John Wiley & Sons, 2005), Chap. 2, pp. 33–34. 7
- [3] H. Takagi *et al.*, Opt. Express **20**, 28292 (2012). 9
- [4] A. Benmerkhi, M. Bouchemat, and T. Bouchemat, Optik - International Journal for Light and Electron Optics **124**, 5719 (2013). 9
- [5] A. Thom and C. J. Apelt, in *Field Computations in Engineering and Physics* (D. van Nostrand Company Ltd, London, 1961), p. v. 10
- [6] K. Yee, IEEE Transactions on Antennas and Propagation **14**, 302 (1966). 10
- [7] A. Taflov, IEEE Transactions on Electromagnetic Compatibility **EMC-22**, 191 (1980). 10
- [8] A. F. Oskooi *et al.*, Computer Physics Communications **181**, 687 (2010). 10, 37
- [9] A. Taflov, *Computational Electrodynamics: The Finite-Difference Time-Domain Method* (Artech House, Boston, 1995). 10
- [10] M. N. Sadiku, in *Numerical Techniques in Electromagnetics*, second edition ed. (CRC Press, London, 2000), pp. 163–164. 12
- [11] J.-P. Berenger, Journal of Computational Physics **114**, 185 (1994). 13
- [12] S. G. Johnson and J. D. Joannopoulos, *Introduction to Photonic Crystals: Bloch's Theorem, Band Diagrams, and Gaps (But No Defects)*, (2003) URL: <http://ab-initio.mit.edu/photons/tutorial/photonic-intro.pdf>, . 15
- [13] D. C. Reynolds *et al.*, Journal of Applied Physics **61**, 342 (1987). 15
- [14] B. Shew, C. Kuo, Y. Huang, and Y. Tsai, Sensors and Actuators A: Physical **120**, 383 (2005). 15, 86
- [15] V. A. Mandelshtam and H. S. Taylor, The Journal of Chemical Physics **107**, 6756 (1997). 18
- [16] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, in *Photonic Crystals: Molding the Flow of Light*, 2nd edition ed. (Princeton University Press, Princeton, New Jersey, 2008), Chap. 5, p. 76. 18

- [17] S. Boscolo, C. Conti, M. Midrio, and C. G. Someda, *J. Lightwave Technol.* **20**, 304 (2002). 30
- [18] R. Biswas, Z. Y. Li, and K. M. Ho, *Applied Physics Letters* **84**, 1254 (2004). 30
- [19] B.-S. Song, S. Noda, T. Asano, and Y. Akahane, *Nature materials* **4**, 207 (2005). 32, 105
- [20] S. Tomljenovic-Hanic, C. M. de Sterke, and M. J. Steel, *Opt. Express* **14**, 12451 (2006). 32, 105
- [21] S.-H. Kwon, T. Süner, M. Kamp, and A. Forchel, *Opt. Express* **16**, 11709 (2008). 32, 105
- [22] P. W. Milonni and J. H. Eberly, in *Laser Physics* (John Wiley & Sons, Hoboken, New Jersey, 2010), Chap. 5, p. 198. 39
- [23] D. Englund and J. Vučković, *Opt. Express* **14**, 3472 (2006). 39
- [24] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, *Opt. Express* **13**, 1202 (2005). 39, 43
- [25] Lumerical, *Quality factor calculations*, URL: https://kb.lumerical.com/en/diffractive_optics_cavity_q_calculation.html, retrieved 2018, (Archived at <https://archive.is/5HMMt>). 39
- [26] A. Wolski, arXiv **1111.4354**, v2 (2011). 40
- [27] D. Englund, I. Fushman, and J. Vuckovic, *Opt. Express* **13**, 5961 (2005). 41
- [28] R. Coccioli *et al.*, *IEE Proceedings - Optoelectronics* **145**, 391 (1998). 43
- [29] J. E. Geusic, H. M. Marcos, and L. G. V. Uitert, *Applied Physics Letters* **4**, 182 (1964). 46
- [30] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich, *Phys. Rev. Lett.* **7**, 118 (1961). 46
- [31] P. F. Moulton, *J. Opt. Soc. Am. B* **3**, 125 (1986). 47
- [32] R. Stolen and A. Ashkin, *Applied Physics Letters* **22**, 294 (1973). 47
- [33] D. E. Spence, P. N. Kean, and W. Sibbett, *Opt. Lett.* **16**, 42 (1991). 47
- [34] P. Ho, L. A. Glasser, E. P. Ippen, and H. A. Haus, *Applied Physics Letters* **33**, 241 (1978). 48
- [35] W. Neumann, in *Fundamentals of dispersive optical spectroscopy systems* (SPIE Press, Bellingham, Washington USA, 2014), Chap. 2, p. 15. 50
- [36] R. Storn and K. Price, *Journal of Global Optimization* **11**, 341 (1997s). 57
- [37] A. A. Michelson and E. W. Morley, *American Journal of Science* **Series 3 Vol. 34**, 333 (1887). 63

- [38] R. L. Powell and K. A. Stetson, J. Opt. Soc. Am. **55**, 1593 (1965). 63
- [39] F. J. Duarte *et al.*, Journal of Optics **13**, 035710 (2011). 63
- [40] V. Vali and R. W. Shorthill, Appl. Opt. **15**, 1099 (1976). 63
- [41] J. E. Baldwin and C. A. Haniff, Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **360**, 969 (2002). 63
- [42] T. C. Sum *et al.*, Applied Physics Letters **83**, 1707 (2003). 63
- [43] D. E. Aspnes, S. M. Kelso, R. A. Logan, and R. Bhat, Journal of Applied Physics **60**, 754 (1986). 63
- [44] A. Perot and C. Fabry, Astrophysical Journal **9**, 87 (1899). 63
- [45] L. Vegard, Zeitschrift für Physik **5**, 17 (1921). 70
- [46] A. R. Denton and N. W. Ashcroft, Phys. Rev. A **43**, 3161 (1991). 70
- [47] K. A. G. Jr. and G. H. Vineyard, Journal of Applied Physics **33**, 3444 (1962). 70
- [48] D. Bimberg *et al.*, in *Semiconductors: Physics of Group IV Elements and III-V Compounds*, edited by O. Madelung (Springer-Verlag, Berlin, 1982), Vol. 17a, Chap. 2.6, 2.10, 2.14, pp. 164, 166, 218, 234, 297, 300. 70
- [49] A. Cho, Journal of Applied Physics **41**, 2780 (1970). 70
- [50] A. Cho, Journal of Crystal Growth **150**, 1 (1995). 70
- [51] J. H. Neave, B. A. Joyce, P. J. Dobson, and N. Norton, Applied Physics A **31**, 1 (1983). 71
- [52] J. D. Kim, X. Chen, and J. J. Coleman, in *Handbook of Crystal Growth: Thin Films and Epitaxy*, 2nd edition ed., edited by T. F. Kuech (Elsevier, Amsterdam, 2015), Vol. volume 3, part A, Chap. 10, pp. 445–446. 72
- [53] Y. Nakata, Y. Sugiyama, and M. Sugawara, in *Self-Assembled InGaAs/GaAs Quantum Dots*, edited by M. Sugawara (Academic Press, London, 1999), Chap. 2, pp. 119–132. 72
- [54] A. Pimpinelli and J. Villain, in *Physics of crystal growth* (Cambridge university press, Cambridge, 1998), Vol. 19, Chap. 5, pp. 73–75. 73
- [55] T. Michely and J. Krug, in *Islands, Mounds and Atoms* (Springer, Berlin, 2004), Chap. 3, p. 61. 73
- [56] A. Sasaki, Journal of Crystal Growth **160**, 27 (1996). 73
- [57] L. Tinkler *et al.*, Applied Physics Letters **106**, 021109 (2015). 73
- [58] P. Kumar, S. Kanakaraju, and D. DeVoe, Applied Physics A **88**, 711 (2007). 73
- [59] X. Sun *et al.*, Solid-State Electronics **53**, 1032 (2009). 73

- [60] S. Sioncke *et al.*, ECS Transactions **16**, 451 (2008). 73
- [61] J.-N. Aqua and T. Frisch, Comptes Rendus Physique **16**, 741 (2015). 74
- [62] Z. Alferov and R. Kazarinov, Authors certificate 28448, U.S.S.R., Sov. Phys. Solid State **9**, 208 (1963). 75
- [63] H. Kroemer, Proceedings of the IEEE **51**, 1782 (1963). 75
- [64] J. Rosenberg *et al.*, IEEE electron device letters **6**, 491 (1985). 75
- [65] H. Schneider *et al.*, Infrared physics & technology **42**, 283 (2001). 75
- [66] M. Behet, J. Bekaert, J. D. Boeck, and G. Borghs, Sensors and Actuators A: Physical **81**, 13 (2000). 75
- [67] K. Barnham *et al.*, Applied Surface Science **113-114**, 722 (1997), proceedings of the Eighth International Conference on Solid Films and Surfaces. 75
- [68] D. A. B. Miller, Opt. Photon. News **1**, 7 (1990). 75
- [69] A. Ekimov, A. Efros, and A. Onushchenko, Solid State Communications **56**, 921 (1985). 76
- [70] R. Rossetti, S. Nakahara, and L. E. Brus, The Journal of Chemical Physics **79**, 1086 (1983). 76
- [71] A. Nozik, Physica E: Low-dimensional Systems and Nanostructures **14**, 115 (2002). 76
- [72] Q. Sun *et al.*, Nature photonics **1**, 717 (2007). 76
- [73] D. Loss and D. P. DiVincenzo, Phys. Rev. A **57**, 120 (1998). 76
- [74] X. Michalet *et al.*, Science **307**, 538 (2005). 76
- [75] E. Jang *et al.*, Advanced Materials **22**, 3076 (2010). 76
- [76] M. Sugawara, in *Self-Assembled InGaAs/GaAs Quantum Dots*, edited by M. Sugawara (Academic Press, London, 1999), Chap. 1, pp. 9–11. 77
- [77] B. P. Reid *et al.*, Japanese Journal of Applied Physics **52**, 08JE01 (2013). 77
- [78] A. Zora, C. Simserides, and G. P. Triberis, Journal of Physics: Conference Series **245**, 012037 (2010). 77
- [79] M. Čalić, Ph.D. thesis, École Polytechnique Fédérale de Lausanne, Switzerland, 2013. 77
- [80] P. M. Petroff, in *Single Quantum Dots: Fundamentals, Applications and New Concepts*, edited by P. Michler (Springer-Verlag, Berlin, 2003), Chap. 1, pp. 1–21. 78
- [81] C. Schneider, S. Hofling, and A. Forchel, in *Quantum Dots: Optics, Electron Transport and Future Applications*, edited by A. Tartakovskii (Cambridge University Press, Cambridge, 2012), Chap. 1, pp. 1–18. 79

- [82] A. Huggenberger *et al.*, Applied Physics Letters **98**, 131104 (2011). 79
- [83] A. V. Kuhlmann *et al.*, Nature communications **6**, 8204 (2015). 79
- [84] K. Mukai and M. Sugawara, in *Self-Assembled InGaAs/GaAs Quantum Dots*, edited by M. Sugawara (Academic Press, London, 1999), Chap. 5, pp. 209–238. 79
- [85] G. La Rocca, *Electronic Excitations in Organic Nanostructures*, Vol. 31 of *Thin Films and Nanostructures* (Academic Press, Amsterdam, 2003), pp. 97 – 128. 79
- [86] J. M. Elward and A. Chakraborty, Journal of Chemical Theory and Computation **9**, 4351 (2013), pMID: 26589152. 80
- [87] D. Gammon and D. G. Steel, Technical report, NAVAL RESEARCH LAB WASHINGTON DC (unpublished). 80
- [88] M. Sugawara, in *Self-Assembled InGaAs/GaAs Quantum Dots*, edited by M. Sugawara (Academic Press, London, 1999), Chap. 1, pp. 97–107. 80
- [89] O. Akimoto and E. Hanamura, Journal of the Physical Society of Japan **33**, 1537 (1972). 80
- [90] R. J. Young *et al.*, in *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies* (Optical Society of America, 2006), p. JTuC2. 81
- [91] C. Santori *et al.*, Physical Review B **66**, 045308 (2002). 81
- [92] N. Akopian *et al.*, Journal of Applied Physics **101**, 081712 (2007). 81
- [93] P. Michler *et al.*, Science **290**, 2282 (2000). 81
- [94] D. W. Minsek and E. L. Alemy, *Epoxy photoresist composition with improved cracking resistance* (2004) US6716568 B1. 82
- [95] J. Liu *et al.*, Microsystem Technologies **10**, 265 (2004). 82
- [96] H. Lorenz *et al.*, Sensors and Actuators A: Physical **64**, 33 (1998), tenth IEEE International Workshop on Micro Electro Mechanical Systems. 82
- [97] E. H. Conradie and D. F. Moore, Journal of Micromechanics and Microengineering **12**, 368 (2002). 82
- [98] N. Chronis and L. P. Lee, Journal of Microelectromechanical Systems **14**, 857 (2005). 82
- [99] C.-H. Lin, G.-B. Lee, B.-W. Chang, and G.-L. Chang, Journal of Micromechanics and Microengineering **12**, 590 (2002). 82
- [100] H. Sato, H. Matsumura, S. Keino, and S. Shoji, Journal of Micromechanics and Microengineering **16**, 2318 (2006). 82
- [101] P. Abgrall *et al.*, Electrophoresis **28**, 4539 (2007). 82

- [102] S.-Q. Xie *et al.*, Microelectronic Engineering **85**, 914 (2008), proceedings of the Micro- and Nano-Engineering 2007 Conference. 82
- [103] J. Greener *et al.*, Lab Chip **10**, 522 (2010). 82
- [104] S. Cho *et al.*, Thin Solid Films **475**, 68 (2005). 82
- [105] L. Pang, W. Nakagawa, and Y. Fainman, Optical Engineering **42**, 2912 (2003). 82
- [106] R. Du, S. Ssenyange, M. Aktary, and M. T. McDermott, Small **5**, 1162 (2009). 82
- [107] B. Riedel *et al.*, Opt. Express **18**, A631 (2010). 82
- [108] O. P. Parida and N. Bhat, in *Int. Conf. on Opt. and Photon.(CSIO)* (Central Scientific Instruments Organization, New Delhi, 2009), pp. 4–7. 82, 85, 86
- [109] M. Aktary *et al.*, Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena **21**, L5 (2003). 82
- [110] S. Juodkazis *et al.*, Nanotechnology **16**, 846 (2005). 82
- [111] J. V. Crivello and J. H. W. Lam, Journal of Polymer Science: Polymer Chemistry Edition **17**, 977 (1979). 82
- [112] J. M. Shaw *et al.*, IBM Journal of Research and Development **41**, 81 (1997). 82
- [113] A. del Campo and C. Greiner, Journal of Micromechanics and Microengineering **17**, R81 (2007). 84
- [114] N. C. LaBianca and J. D. Gelorme, Proc.SPIE **2438**, 2438 (1995). 84, 85
- [115] H. Lorenz *et al.*, Journal of Micromechanics and Microengineering **7**, 121 (1997). 84
- [116] H. Lorenz, M. Laudon, and P. Renaud, Microelectronic Engineering **41-42**, 371 (1998), international Conference on Micro- and Nanofabrication. 84
- [117] MicroChem, *SU-8 2000 Permanent Epoxy Negative Photoresist Processing Guidelines*, URL: http://www.microchem.com/pdf/SU-82000DataSheet2000_5thru2015Ver4.pdf, retrieved 2018, (Archived at <http://www.webcitation.org/6x77CaVAc>). 84
- [118] L. E. Scriven, MRS Proceedings **121**, 717 (1988). 84
- [119] R. Martinez-Duarte and M. J. Madou, in *Microfluidics and Nanofluidics Handbook: Fabrication, Implementation, and Applications*, edited by S. K. Mitra and S. Chakraborty (CRC Press, New York, 2012), Chap. SU-8 Photolithography and Its Impact on Microfluidics, p. 249. 84, 86
- [120] T. A. Anhoj, A. M. Jorgensen, D. A. Zauner, and J. Hübner, Journal of Micromechanics and Microengineering **16**, 1819 (2006). 84
- [121] J. Serbin, A. Ovsianikov, and B. Chichkov, Opt. Express **12**, 5221 (2004). 85

- [122] W. Teh *et al.*, Journal of applied physics **97**, 054907 (2005). 85
- [123] MicroChem, *SU-8 Developer Safety Data Sheet*, (2008) URL: http://nano.pse.umass.edu/sites/default/files/msds_ghs_su8developer_rev0.pdf, (Archived at <http://www.webcitation.org/6x8U7qN3I>). 86
- [124] M. C. Peterman, P. Huie, D. M. Bloom, and H. A. Fishman, Journal of Micromechanics and Microengineering **13**, 380 (2003). 86
- [125] T. Kamiya *et al.*, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **267**, 2317 (2009), proceedings of the 11th International Conference on Nuclear Microprobe Technology and Applications and the 3rd International Workshop on Proton Beam Writing. 86
- [126] B. Bêche, N. Pelletier, E. Gaviot, and J. Zyss, Optics Communications **230**, 91 (2004). 86
- [127] National Renewable Energy Laboratory, *Photovoltaic Research Cell Efficiency Chart*, (2018) retrieved 2018, (Archived at <https://web.archive.org/web/20181016032645/https://www.nrel.gov/pv/assets/pdfs/pv-efficiencies-07-17-2018.pdf>). 87, 136, 137
- [128] D. B. Mitzi, Chemistry of Materials **13**, 3283 (2001). 87
- [129] M. I. Saidaminov *et al.*, Nature communications **6**, 7586 (2015). 87
- [130] J. Borchert *et al.*, ACS Energy Letters **2**, 2799 (2017). 87
- [131] Q. Chen *et al.*, Nano Today **10**, 355 (2015). 87, 89, 91
- [132] C. Li *et al.*, Acta Crystallographica Section B: Structural Science **64**, 702 (2008). 88
- [133] Z. Li *et al.*, Chemistry of Materials **28**, 284 (2016). 88
- [134] O. A. Syzgantseva, M. Saliba, M. Grätzel, and U. Rothlisberger, The Journal of Physical Chemistry Letters **8**, 1191 (2017), pMID: 28229595. 88
- [135] G. Niu, X. Guo, and L. Wang, J. Mater. Chem. A **3**, 8970 (2015). 89
- [136] Y. Hu *et al.*, J. Mater. Chem. A **5**, 25258 (2017). 89
- [137] G. Abdelmageed *et al.*, Applied Physics Letters **109**, 233905 (2016). 89
- [138] J. Zhao *et al.*, Scientific reports **6**, 21976 (2016). 89
- [139] J. Lee *et al.*, Advanced Energy Materials **5**, 1501310 (2015). 89
- [140] T. Baikie *et al.*, J. Mater. Chem. A **1**, 5628 (2013). 89
- [141] D. P. McMeekin *et al.*, Science **351**, 151 (2016). 90, 137, 157, 161, 178
- [142] G. E. Eperon *et al.*, Energy Environ. Sci. **7**, 982 (2014). 90, 140
- [143] Y. Ogomi *et al.*, The journal of physical chemistry letters **5**, 1004 (2014). 90

- [144] J. H. Noh *et al.*, Nano Letters **13**, 1764 (2013), pMID: 23517331. 90
- [145] B.-w. Park *et al.*, J. Mater. Chem. A **3**, 21760 (2015). 90, 139
- [146] T. Jesper Jacobsson *et al.*, Energy Environ. Sci. **9**, 1706 (2016). 90, 139
- [147] A. Buin *et al.*, Chemistry of Materials **27**, 4405 (2015). 90, 139
- [148] W. Shockley and H. J. Queisser, Journal of applied physics **32**, 510 (1961). 90
- [149] S. Rühle, Solar Energy **130**, 139 (2016). 90
- [150] W. Guter *et al.*, Applied Physics Letters **94**, 223504 (2009). 90
- [151] I. Borriello, G. Cantele, and D. Ninno, Phys. Rev. B **77**, 235214 (2008). 91
- [152] Y. Chang, C. Park, and K. Matsuishi, Journal-Korean Physical Society **44**, 889 (2004). 91
- [153] K. X. Steirer *et al.*, ACS Energy Letters **1**, 360 (2016). 91
- [154] M. V. Kovalenko, L. Protesescu, and M. I. Bodnarchuk, Science **358**, 745 (2017). 91
- [155] G. Xing *et al.*, Science **342**, 344 (2013). 91
- [156] W. Rehman *et al.*, Advanced Materials **27**, 7938 (2015). 91, 157
- [157] A. A. Zhumekenov *et al.*, ACS Energy Letters **1**, 32 (2016). 91, 148
- [158] Q. Dong *et al.*, Science **347**, 967 (2015). 91, 148
- [159] D. H. Kobe and V. C. Aguilera-Navarro, Phys. Rev. A **50**, 933 (1994). 94
- [160] M. Fox, in *Quantum Optics – An Introduction, Oxford Master Series In Physics* (Oxford University Press, Oxford, 2006), Chap. 5, pp. 75–101. 94
- [161] M. Fox, in *Quantum Optics – An Introduction, Oxford Master Series In Physics* (Oxford University Press, Oxford, 2006), Chap. 7, pp. 126–146. 95
- [162] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, Rev. Mod. Phys. **74**, 145 (2002). 96, 102
- [163] B. Schneier, *Applied Cryptography* (John Wiley & Sons, 1996). 96
- [164] P. Shor, SIAM Review **41**, 303 (1999). 96
- [165] R. A. Perlner and D. A. Cooper, in *Proceedings of the 8th Symposium on Identity and Trust on the Internet, IDtrust '09* (ACM, New York, NY, USA, 2009), pp. 85–93. 96
- [166] M. Grassl, B. Langenberg, M. Roetteler, and R. Steinwandt, in *Post-Quantum Cryptography*, edited by T. Takagi (Springer International Publishing, Cham, 2016), pp. 29–43. 96
- [167] W. Wootters and W. Zurek, Nature **299**, 802 (1982). 97

- [168] D. Yang, *Physics Letters A* **360**, 249 (2006). 97
- [169] C. H. Bennett and G. Brassard, *Theor. Comput. Sci.* **560**, 7 (2014). 97
- [170] G. Brassard, N. Lütkenhaus, T. Mor, and B. C. Sanders, *Phys. Rev. Lett.* **85**, 1330 (2000). 98
- [171] T. Schmitt-Manderbach *et al.*, *Phys. Rev. Lett.* **98**, 010504 (2007). 99
- [172] D. Stucki *et al.*, *New Journal of Physics* **11**, 075003 (2009). 99
- [173] V. Scarani, A. Acín, G. Ribordy, and N. Gisin, *Phys. Rev. Lett.* **92**, 057901 (2004). 99
- [174] A. Acín, N. Gisin, and V. Scarani, *Phys. Rev. A* **69**, 012309 (2004). 99
- [175] X.-B. Wang, *Phys. Rev. Lett.* **94**, 230503 (2005). 99
- [176] K. Inoue and T. Honjo, *Phys. Rev. A* **71**, 042305 (2005). 99
- [177] M. Schiavon, G. Vallone, F. Ticozzi, and P. Villoresi, *Phys. Rev. A* **93**, 012331 (2016). 99
- [178] E. Knill, R. Laflamme, and G. J. Milburn, *nature* **409**, 46 (2001). 99
- [179] M. Fox, in *Quantum Optics – An Introduction, Oxford Master Series In Physics* (Oxford University Press, Oxford, 2006), Chap. 10, pp. 194–211. 99, 100, 101
- [180] P. Grangier and I. Abram, *Physics World* **16**, 31 (2003). 102
- [181] B. W. Shore and P. L. Knight, *Journal of Modern Optics* **40**, 1195 (1993). 103
- [182] F. S. F. Brossard *et al.*, *Applied Physics Letters* **97**, 111101 (2010). 104
- [183] H. Kim *et al.*, *Nature Photonics* **7**, 373 (2013). 104
- [184] C. Simon *et al.*, *The European Physical Journal D* **58**, 1 (2010). 104
- [185] Q. A. Turchette *et al.*, *Phys. Rev. Lett.* **75**, 4710 (1995). 104
- [186] M. Gschrey *et al.*, *Applied Physics Letters* **102**, 251113 (2013). 105
- [187] K. H. Lee *et al.*, *Applied physics letters* **88**, 193106 (2006). 105, 183
- [188] S. T. Thornton and A. Rex, in *Modern Physics for Scientists and Engineers*, fourth edition ed. (Brooks/Cole, Boston, 2013), Chap. 11, p. 398. 110
- [189] L. Feng *et al.*, *Advanced materials* **14**, 1857 (2002). 116
- [190] L. Jiang, Y. Zhao, and J. Zhai, *Angewandte Chemie* **116**, 4438 (2004). 116
- [191] A. Vorobyev and C. Guo, *Journal of Applied Physics* **117**, 033103 (2015). 116
- [192] W. Barthlott, M. Mail, and C. Neinhuis, *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **374**, (2016). 116

- [193] A. B. D. Cassie and S. Baxter, *Trans. Faraday Soc.* **40**, 546 (1944). 116
- [194] J. Schanda, in *Colorimetry* (Wiley-Blackwell, 2007), Chap. 3, pp. 43–44. 124
- [195] A. Kojima, K. Teshima, Y. Shirai, and T. Miyasaka, *Journal of the American Chemical Society* **131**, 6050 (2009), pMID: 19366264. 136
- [196] N.-G. Park *et al.*, *Nature Energy* **1**, 16152 (2016). 137
- [197] G. Grancini *et al.*, *Nature communications* **8**, 15684 (2017). 137
- [198] T. Dittrich *et al.*, *The Journal of Physical Chemistry C* **119**, 23968 (2015). 139
- [199] K. Wu *et al.*, *Phys. Chem. Chem. Phys.* **16**, 22476 (2014). 139
- [200] C. Yu *et al.*, *Journal of Applied Physics* **110**, 063526 (2011). 139
- [201] P. Gratia *et al.*, *Journal of the American Chemical Society* **138**, 15821 (2016), pMID: 27960332. 145
- [202] C.-Z. Li *et al.*, *Advanced Materials* **26**, 6262 (2014). 148
- [203] O. A. Abdelraouf and N. K. Allam, *Solar Energy* **137**, 364 (2016). 148
- [204] S. D. Stranks *et al.*, *Science* **342**, 341 (2013). 148
- [205] J. Philibert, *Diffusion Fundamentals* **4**, 6 (2006). 148
- [206] J. Balenzategui and A. Martí, *Solar Energy Materials and Solar Cells* **90**, 1068 (2006). 154
- [207] L. M. Pazos-Outón *et al.*, *Science* **351**, 1430 (2016). 154
- [208] A. Tejada *et al.*, *Journal of Applied Physics* **123**, 175302 (2018). 157
- [209] R. A. Street, M. Schoendorf, A. Roy, and J. H. Lee, *Phys. Rev. B* **81**, 205307 (2010). 162
- [210] L. Tzabari and N. Tessler, *Journal of Applied Physics* **109**, 064501 (2011). 162
- [211] M. B. Johnston and L. M. Herz, *Accounts of Chemical Research* **49**, 146 (2016), pMID: 26653572. 162
- [212] E. Hairer, S. P. Norsett, and G. Wanner, *Solving ordinary differential equations I, Nonstiff Problems*, Vol. Volume 8 of *Springer Series in Computational Mathematics* (Springer-Verlag, Berlin, 1993). 163
- [213] Z. Li and H. A. Scheraga, *Proceedings of the National Academy of Sciences* **84**, 6611 (1987). 167
- [214] B. Olson, I. Hashmi, K. Molloy, and A. Shehu, *Advances in Artificial Intelligence* **2012**, 3 (2012). 167
- [215] L. M. Smith *et al.*, *Phys. Rev. B* **38**, 5788 (1988). 176

- [216] B. A. Ruzicka, L. K. Werake, H. Samassekou, and H. Zhao, *Applied Physics Letters* **97**, 262119 (2010). 176
- [217] B. A. Ruzicka *et al.*, *Phys. Rev. B* **82**, 195414 (2010). 176
- [218] W. Tian *et al.*, *Journal of the American Chemical Society* **137**, 12458 (2015), pMID: 26390276. 176, 178
- [219] H. Zhu *et al.*, *Advanced Materials* **29**, 1603072 (2017). 178
- [220] Z.-G. Yu, *The Journal of Physical Chemistry Letters* **7**, 3078 (2016), pMID: 27459897. 178
- [221] R. L. Milot *et al.*, *Advanced Functional Materials* **25**, 6218 (2015). 178
- [222] M. C. Gélvez-Rueda, N. Renaud, and F. C. Grozema, *The Journal of Physical Chemistry C* **121**, 23392 (2017). 178
- [223] S. A. Lennon *et al.*, *Opt. Express* **26**, 32332 (2018). 184
- [224] X. Cheng, H. Ye, and Z. Yu, *Superlattices and Microstructures* **105**, 81 (2017). 184
- [225] J. Liang *et al.*, *Journal of the American Chemical Society* **138**, 15829 (2016), pMID: 27960305. 186
- [226] J. Liang, J. Liu, and Z. Jin, *Solar RRL* **1**, 1700086 (2017). 186
- [227] P. Becker *et al.*, *Advanced Energy Materials* **9**, 1900555 (2019). 186
- [228] H. Abelson, G. J. Sussman, and J. Sussman, *Structure and Interpretation of Computer Programs*, second edition ed. (MIT Press, London, 1996). 192