

Investigation of Photon Interactions with Semiconductor Quantum Dot Devices for Quantum Communication Applications

by

Jason Phoenix

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Science
in
Physics

Waterloo, Ontario, Canada, 2019

© Jason Phoenix 2019

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis was entirely written by the author, but all of the research was done in conjunction with employees of National Research Council Canada (NRC). In particular, all research presented in this thesis was conducted in a partnership between the author and Dr. Louis Gaudreau of the NRC. Through the use of terms such as ‘I’ and ‘we’ in the text, the author has attempted wherever possible to distinguish between collective actions and the actions of the author alone.

This research is all based upon the photon-to-spin project proposed and developed by Dr. Gaudreau and the NRC’s quantum physics group, which is headed by Dr. Andrew Sachrajda. All research presented here was conducted in NRC’s new SPIN Lab. The construction and design of this lab and its custom-built systems were primarily carried out by Dr. Gaudreau, Piotr Zawadzki, and Dr. Sergei Studenikin, with some very minor assistance from the author. Some of the high-frequency electronic systems were designed in collaboration with Professor Jan Kycia of the University of Waterloo.

The nanowire samples discussed in Chapter 4 were fabricated by Dr. Dan Dalacu of the NRC, who also assisted with the research described in that chapter. Dr. Dalacu provided access to optical equipment used for those experiments, such as a spectrometer, single-photon counters, and autocorrelation software for single-photon measurements.

The contents of Chapter 5 are to be published in a peer-reviewed journal in the near future, largely without further modification. While that chapter was written by the author, who will appear as first author on the associated paper, suggestions for the project and corrections to the paper were made by: Dr. Marek Korkusinski, Mr. Zawadzki, Dr. Alex Bogan, Dr. Studenikin, Dr. Robin Williams, and Dr. Sachrajda. All of these individuals will appear as co-authors on the paper, as will Dr. Gaudreau, with whom the research was conducted. As described in Chapter 5, Dr. Korkusinski created a genetic algorithm for the experiment, which is based on a mathematical model provided by the author and Dr. Gaudreau. Dr. Bogan devised the proof that the SPIN Lab’s cryogenic optical setup can uniquely identify any polarisation. This proof is mentioned in Chapter 5 and described in detail in Appendix F.

The author has attempted to acknowledge all contributions from collaborators at the appropriate points in each chapter.

Abstract

A major goal in the field of quantum communication is to achieve long-distance (>100 km) transmission of quantum information, which would allow for the formation of a global quantum network. Devices called ‘quantum repeaters’ will enable delicate quantum states to be transmitted over long distances without succumbing to the signal losses inherent in the use of optical fibres. This thesis presents my work on a new type of quantum repeater design, which will combine both the photonic and spin qubit platforms to achieve more robust and efficient quantum communication. This ‘photon-to-spin’ research was conducted at the National Research Council Canada (NRC) in Ottawa, in the labs of Dr. Andrew Sachrajda’s quantum physics group. This work has been aided by collaboration with Professor Jan Kycia’s low-temperature condensed matter physics group at the University of Waterloo, who helped design our cryogenic, high-frequency reflectometry circuit. In the following chapters, I will describe my work at the NRC’s newly-constructed SPIN Lab (or Spin-Photon Interactions in Nanostructures Laboratory). By combining quantum optics and quantum electronics technologies, the purpose of this lab is to enable the study of interactions between individual photons and electron spins in nanostructures. We will look at the setup of the SPIN lab in Chapter 2.

During the course of my degree, I have worked to develop the lab’s optical capabilities, helping to create a system for delivering light to nanostructure samples via optical fibre. As the samples are microscopic in scale and fully enclosed in a windowless dilution refrigerator, accurate light delivery requires the samples to be imaged in advance. As we shall see in Chapter 3, piezoelectric positioners are used to scan a lensed fibre over a sample, allowing imaging to be accomplished with sub-micron resolution. Nanostructure topography is observed by measuring photocurrent, reflected light, and photoluminescence emissions.

Nanowire quantum dots have been tested in the SPIN Lab’s dilution refrigerator for their potential to act as on-demand sources of fibre-coupled, indistinguishable single photons. In Chapter 4, I will review our techniques for direct coupling of nanowire emissions into a lensed fibre. Although the fibre’s collection efficiency is just 1–2%, this is sufficient to perform a variety of optical tests, including emission spectra analyses and an investigation of photon indistinguishability. The degree of second-order coherence is measured to be $g^{(2)}(0) = 0.095 \pm 0.027$ for fibre-coupled photons, surpassing the threshold required for a single photon source to have practical applications.

Finally, Chapter 5 summarises a method of fully controlling the polarisation of light sent through a fibre to a dilution refrigerator. As the birefringence of a fibre core alters polarisation, a newly-designed setup involving a mathematical fibre model and standard optical components is used to compensate for the fibre and deliver various polarisation states to the refrigerator. State fidelities of greater than 0.96 ± 0.007 are achieved while operating at cryogenic temperatures, and the successful delivery of randomly-selected elliptical states demonstrates that this system can transmit any arbitrary polarisation by fibre. These results will soon be submitted for publication, once the application to patent the polarisation system has been reviewed.

Acknowledgements

First and foremost, I would like to thank Dr. Louis Gaudreau. You have been a great friend and mentor throughout my time at NRC, and from you I have gained invaluable insight into the world of professional physics research. Our work on the SPIN Lab has been interesting, challenging, and enormously educative.

Thank you also to my supervisors, Dr. Sergei Studenikin and Professor Jan Kycia, for your support, assistance, suggestions, and tutelage. The courses I took with you as my professors (Professor Kycia during my undergrad and Dr. Studenikin during my master's) provided me with valuable knowledge of this field, knowledge which has been incredibly useful during my master's degree and which will no doubt contribute to any success I might have in the future.

Likewise, many thanks to Dr. Dan Dalacu for educating me about nanowires and providing significant assistance during our extensive experiments with your nanowire samples. Working with you has been a lot of fun.

Piotr Zawadzki, you have been a great friend, and I can't thank you enough for your support, for our fascinating and thought-provoking conversations, for educating me about cryogenics and the operation of dilution refrigerators, and for the beautiful pictures of the lab you have provided for this thesis. You have made my graduate experience both interesting and entertaining.

Thank you to Dr. Alex Bogan for being such an amiable office-mate, for our interesting conversations, and for the initial development of our experiment automation software. I would also like to thank both Dr. Bogan and Dr. Marek Korkusinski for helping to educate me about many of the more complicated details of our field of research, i.e. the theory.

Dr. Andrew Sachrajda and Dr. Robin Williams, I thank you both for use of your labs and for welcoming me into your quantum physics group at the National Research Council (NRC).

I would like to thank Dr. Melanie Campbell for your hard work and patience in helping me set up my unusual graduate position with the NRC. Thank you also to the NRC for providing me with the opportunity to perform my graduate research in one of your labs and for helping to fund my graduate studies. This experience has been fascinating and instructive, and it would not have been possible without the support of the NRC.

Finally, many thanks to my parents, Dean and Suzanne, for endowing me with an interest in understanding the fundamental nature of reality and for teaching me to always ask questions, especially the ones which seem most unanswerable.

Table of Contents

List of Figures	xi
List of Tables	xv
List of Abbreviations	xvi
List of Symbols	xviii
1 Introduction	1
1.1 Lateral, Gate-Defined Quantum Dots	4
1.1.1 Structure and Formation of a Lateral Quantum Dot	4
1.2 Nanowire Quantum Dots	7
1.2.1 Introduction to Photonic Nanowires	7
1.2.2 Fabrication of a Photonic Nanowire with Embedded Quantum Dot	9
1.2.3 Photoluminescence and Exciton Complexes in a Nanowire Quantum Dot	14
1.2.4 Confirming Single-Photon Emission with the Hanbury Brown-Twiss Autocorrelation Experiment	20
1.3 Quantum Dot Charge Sensing	25
1.3.1 Quantum Point Contact as a Charge Sensor	25
1.3.2 Radio Frequency Quantum Point Contact as a Charge Sensor	27
1.3.3 Theoretical Modelling of an LCR Tank Circuit	30
1.4 Light Polarisation	36
1.4.1 Classical vs. Quantum Polarisation	36
1.4.2 Waveplates and Polarisation Control	39

1.5	Mathematical Treatment of Light Polarisation	42
1.5.1	Mueller Calculus	42
1.5.2	Jones Calculus	46
1.5.3	Mathematical Description of an Optical Fibre	49
1.5.4	Fidelity of a Polarisation State	51
1.6	Fibre-Based Transmission of Polarised Light	54
1.6.1	Fibre Birefringence	54
1.6.2	Compensating for Fibre Birefringence	56
1.6.3	Long-Distance Transmission of Polarisation States	58
1.7	The Photon-to-Spin Qubit Transfer Protocol	60
1.7.1	Quantum Teleportation	60
1.7.2	Quantum Repeaters and Entanglement Distribution	65
1.7.3	The Photon-to-Spin Interface	67
1.7.4	Advantages of a Photon-to-Spin Quantum Repeater	71
2	The SPIN Lab	72
2.1	The Equipment and Systems of the SPIN Lab	73
2.1.1	The Dilution Refrigerator	73
2.1.2	The Cryogenic SPIN Setup	75
2.1.3	Equipment Control and Experiment Automation	78
2.2	Quantum Dot Samples	79
2.2.1	Mounting a Sample in the Refrigerator	80
2.2.2	Lateral Quantum Dots for Photon-to-Spin Qubit Transfer	81
2.3	Preliminary Testing of the RF-QPC Charge Sensor	83
2.3.1	Layout of the RF Reflectometry Circuit	83
2.3.2	Computer Modelling of the RF Reflectometry Circuit	87
2.3.3	In-Situ Tuning of the Resonance Frequency Using a Variable Capacitor	89
2.3.4	Testing the RF-QPC System on QPC Samples	91
2.4	Optical Setup for Fibre-Based Light Delivery and Polarisation Identification	95
2.4.1	Room-Temperature Optical Setup	96
2.4.2	Cryogenic Optical Setup	97
2.5	Conclusions	98

3	Distributed Imaging of Nanostructures	99
3.1	Operating Principles of Piezoelectric Nanopositioners	100
3.2	The SPIN Lab’s Piezoelectric Nanopositioners	103
3.2.1	Calibration of the Resistive Position Encoding System	104
3.2.2	Computer Control of the Positioning Stages	105
3.3	Imaging with Piezoelectric Nanopositioners	106
3.3.1	Development of a High-Resolution Imaging Procedure	106
3.3.2	Focusing the Fibre on a Sample	109
3.3.3	Heat Output of the Piezoelectric Positioners	112
3.3.4	Sample Damage Caused by Fibre During Imaging Scan	114
3.4	Options for Measuring Sample Topography	118
3.4.1	Photocurrent Imaging	118
3.4.2	Reflectance Imaging	121
3.4.3	Photoluminescence Imaging	123
3.5	Conclusions and Future Upgrades to the Imaging System	127
4	Fibre-Coupled Nanowire Quantum Dots as Sources of Single Photons	129
4.1	PL Measurements of Nanowire Quantum Dots	130
4.1.1	Structure and Layout of the Nanowire Samples	131
4.1.2	Experimental Setup used for Fibre-Based PL Collection	134
4.1.3	Acquiring and Processing Nanowire Emission Spectra	136
4.1.4	Analysis of Single-Dot Spectra	138
4.1.5	Analysis of Double-Dot Spectra	141
4.1.6	Pump Power Dependence and Lensed Fibre Collection Efficiency	144
4.2	Indistinguishable Photons	149
4.2.1	Lifetime-Limited Emission Linewidths	149
4.2.2	Emission Line Broadening due to Phonon Interactions	150
4.2.3	Spectral Wandering due to a Fluctuating Charge Environment	152
4.2.4	Quantifying Photon Indistinguishability	153
4.3	A Fibre-Coupled Source of Single Photons	155
4.3.1	Preparing for Hanbury Brown-Twiss Interferometry	155

4.3.2	Continuous Wave Hanbury Brown-Twiss Measurement	157
4.3.3	Pulsed Hanbury Brown-Twiss Measurement	160
4.3.4	Exciton Lifetime Measurement	163
4.4	Conclusions and Future Work	165
5	Full Polarisation Control of Fibre-Delivered Light in a Dilution Refrigerator	167
5.1	Details on the Polarisation Control Setup	168
5.2	Preliminary Characterisation of Optical Components	171
5.2.1	Determining Waveplate Retardances	172
5.2.2	The Depolarising Effects of a Gradient-Index Lens	175
5.3	Modelling the Optical Fibre's Effects on Polarisation	175
5.4	Delivery of Polarised Light	176
5.4.1	Step 1: Measure Fibre Effects	177
5.4.2	Step 2: Fit Fibre Parameters	177
5.4.3	Steps 3 & 4: Calculate Waveplate Orientations and Verify Delivery of States	179
5.5	Stability of Fibre Compensation System	184
5.6	Conclusions and Future Work	185
6	Conclusions and Outlook	188
	References	191
	APPENDICES	198
A	Quantum Teleportation via Bell State Measurements	199
B	Platypus Experiment Automation Software	202
B.1	Introduction to Platypus	202
B.1.1	Devices	203
B.1.2	Experiments	204
B.1.3	User Interface	205
B.2	Significance of Platypus	207

C	Interactive Data Plotting and Analysis Software	209
C.1	Introduction to the Interactive Plotter	209
C.1.1	Layout and Basic Features	209
C.1.2	Altering the Plot's Appearance	211
C.1.3	Comparing and Analysing Data	212
C.1.4	Trace Viewing and Spectrum Viewing	213
C.2	Significance of the Interactive Plotter	215
D	The SPIN Lab's Piezoelectric Nanopositioning System	216
D.1	Temperature-Dependent Positioner Behaviour	216
D.2	Positioner Calibration Curves	217
D.3	Hysteresis Effects	218
D.3.1	Extreme Hysteresis in the Vertical Positioner's Step Size	218
D.3.2	Positioner Voltage Probe Hysteresis	218
E	Details of the HBT Plot	220
E.1	'Negative' Time Delay in an HBT Plot	220
E.2	Peak Symmetry in a Pulsed HBT Plot	221
F	Uniquely Identifying Polarisation States	223

List of Figures

1.1	Lateral Double Quantum Dot Device	5
1.2	Nanowire Quantum Dot Samples	8
1.3	Selective-Area Patterning for Nanowire Growth	10
1.4	Nanowire Growth Stages	12
1.5	Mechanisms of Indium Transport to the Nanowire Catalyst	13
1.6	Exciton Trapping in a Quantum Dot	17
1.7	S-Shell Exciton Complexes	18
1.8	Hanbury Brown-Twiss Autocorrelation Experiment	22
1.9	Quantum Point Contact	25
1.10	Basic Operation of the RF-QPC	28
1.11	LCR Tank Circuit Schematic	31
1.12	Simulating a Simple Tank Circuit	35
1.13	Classical Electromagnetic Wave	37
1.14	Half-wave and Quarter-wave Plates Acting on Polarisation	40
1.15	Poincaré Sphere	46
1.16	Basic Quantum Repeater	67
1.17	Photon-to-Spin Hybrid Quantum Repeater	68
1.18	Photon-to-Spin Qubit Transfer	68
1.19	Optical Selection Rules for Photon-to-Spin Qubit Transfer	69
2.1	The SPIN Lab	72
2.2	SPIN Lab Dilution Refrigerator	74
2.3	Experimental Apparatus in the Dilution Fridge	76
2.4	The Cryogenic SPIN Setup	77

2.5	Sample Holder	81
2.6	Images of Operational Photon-to-Spin DQD Samples	82
2.7	Images of Broken Photon-to-Spin DQD Samples	82
2.8	The SPIN Lab RF Reflectometry Circuit	84
2.9	Simulating the SPIN Lab RF-QPC	88
2.10	Tuning the RF-QPC Resonance Frequency with a Variable Capacitor	90
2.11	Effect of QPC Resistance on RF-QPC Resonance	92
2.12	Effect of RF Signal Power on RF-QPC Resonance	93
2.13	Quality Factor of a Deep RF-QPC Resonance	94
2.14	Optical Setup for Fibre-Based Polarisation Delivery and Identification	95
2.15	Room-Temperature Optical Setup for Polarisation Delivery	97
3.1	Fibre Delivery of Light to DQD Sample	99
3.2	Slip-Stick Movement and Resistive Position Encoder of a Piezoelectric Nanopositioning Stage	101
3.3	Positioning of the Lensed Fibre Using XYZ Piezoelectric Stages	103
3.4	Position Calibration Curves for Piezoelectric Stages	105
3.5	Interlacing Distortion in Imaging Scans	107
3.6	Focusing Fibre on a Nanostructure	110
3.7	Focusing Fibre using Photoluminescence	112
3.8	Heating of Fridge due to Sample Imaging Procedure	113
3.9	Damage to Nanowire Sample Caused by Fibre During Scans	115
3.10	Damage to Lensed Fibre Caused by Contact with Sample Surface	117
3.11	Distributed Imaging Using Photocurrent	119
3.12	Imaging Mask Holes of Lateral Dot Samples Using Photocurrent	120
3.13	Imaging Mask Holes of Lateral Dot Samples Using Reflectance	121
3.14	Distributed Imaging Using Reflectance	122
3.15	Example PL Spectrum of a Nanowire	123
3.16	Distributed Imaging Using Photoluminescence at Three Different Wavelength Ranges	125
3.17	Photoluminescence Scan of a Row of Nanowires	126
3.18	High-Resolution Photoluminescence Map of a Nanowire Quantum Dot	127

4.1	Fibre-Coupled Nanowire	130
4.2	Nanowire Waveguide Shell and Gaussian Emission Profile	132
4.3	Layout of Nanowire Samples	133
4.4	Setup for Collection of Nanowire Photoluminescence Emissions	134
4.5	Spectrometer Dark Counts	137
4.6	Eliminating Cosmic Rays from Measured Spectra	138
4.7	Photoluminescence Spectrum of a Single-Dot Nanowire	139
4.8	Energy Band Diagram of Nanowire Defects	140
4.9	Photoluminescence Spectrum of a Double-Dot Nanowire	142
4.10	Pump Power Dependence of Nanowire Emissions	145
4.11	Phonon Sidebands in Quantum Dot Emission Spectra	151
4.12	Width of Zero Phonon Line versus Temperature	152
4.13	Setup for Hanbury Brown-Twiss Experiment	156
4.14	Isolating X^- Peak for Hanbury Brown-Twiss Experiment	157
4.15	Continuous Wave Hanbury Brown-Twiss Data and g_2 Fit	159
4.16	Pulsed Hanbury Brown-Twiss Data and g_2 Fit	162
4.17	Exciton Lifetime Measurement	164
5.1	Optical Setup for Fibre-Based Polarisation Delivery and Identification	169
5.2	Correcting for Noise in Polarisation Data	170
5.3	Fridge Heating During Polarisation Tests	171
5.4	Characterising Polariser and HWP	172
5.5	Characterising QWP Retardance and Offset	174
5.6	Polarisation Maps and Simulations	178
5.7	Simulated Signatures of the Six Standard Polarisations	180
5.8	Measured Signatures of the Six Standard Polarisations	181
5.9	Measured Signatures of Three Randomly-Selected Elliptical Polarisations	183
5.10	Stability of Polarisation Delivery System	185
B.1	Platypus: User Interface	206
B.2	Platypus: State Monitor	207

C.1	Interactive Plotting and Analysis Program	210
C.2	Plotting Program: Editing Plot Display	211
C.3	Plotting Program: Comparing Data Sets	212
C.4	Plotting Program: Trace Viewing	213
C.5	Plotting Program: Spectrum Viewing	214
D.1	Position Calibration Curves For Piezoelectric Stages	217
E.1	Explanation for the Symmetry of Peaks in a Pulsed HBT Experiment	221

List of Tables

1.1	Mathematical Representations of the Standard Polarisation States	44
1.2	Mathematical Representations of Standard Optical Elements	47
1.3	Polarisation Fidelity Equations	53
4.1	Sources of Peaks in a Nanowire Emission Spectrum	140
5.1	Measured Fidelities of Polarised Light Transmitted Through a Fibre	182
E.1	Explanation for the Symmetry of Peaks in a Pulsed HBT Experiment	222

List of Abbreviations

- 2DEG** 2-dimensional electron gas
- AOI** angle of incidence
- APD** avalanche photodiode
- BS** beam splitter
- BSM** Bell state measurement
- CNOT** controlled-NOT (gate)
- CR** cosmic ray
- CW** continuous wave (laser)
- D-A** donor-acceptor (impurities)
- DGD** differential group delay
- DQD** double quantum dot
- EM** electromagnetism
- EPC** electronic polarisation controller
- FWHM** full width at half maximum
- GRIN** gradient-index (lens)
- HBT** Hanbury Brown and Twiss (interferometry)
- HOM** Hong-Ou-Mandel (interferometry)
- HWP** half-wave plate
- MC** mixing chamber
- NA** numerical aperture

NRC National Research Council Canada
PCB printed circuit board
PL photoluminescence
PM polarisation maintaining
PMD polarisation mode dispersion
QD quantum dot
QKD quantum key distribution
QM quantum memory
QPC quantum point contact
QWP quarter-wave plate
RF radio frequency
RF-QPC radio frequency quantum point contact
RF-SET radio frequency single electron transistor
SEM scanning electron microscopy
SF (crystal layer) stacking faults
TEM transmission electron microscopy
TMI trimethylindium
WZ wurtzite (crystal structure)
ZB zincblende (crystal structure)
ZPL zero-phonon line

List of Symbols

$\lambda/2$ half-wave plate

$\lambda/4$ quarter-wave plate

A anti-diagonal polarisation

D diagonal polarisation

H horizontal polarisation

L left circular polarisation

R right circular polarisation

V vertical polarisation

X neutral exciton

X⁻ negatively-charged exciton

XX neutral biexciton

XX⁻ negatively-charged biexciton

Chapter 1

Introduction

As the related fields of quantum computing and quantum communication advance, an important step in their development will be the creation of a global quantum network. The ability to transmit quantum information over great distances, as well as the security ensured by various quantum communication protocols, will be key components of a future ‘quantum internet’. Photons have generally been considered as the ideal candidates for encoding and transmitting quantum information, given their speed, but we inevitably run up against the inherent signal attenuation experienced by light in standard telecom optical fibres.¹ A 1550 nm telecom signal, for example, will be attenuated by 0.2 dB/km in a standard fibre, meaning that 99% of the signal will be lost after it travels just 100 km through the fibre [1]. The limit on the fibre-based transmission of photons is therefore generally considered to be roughly 100 km or less, which is insufficient for the purpose of creating a global quantum network.

The solution to this problem has been to devise a method of using entangled states and quantum repeaters to reduce signal losses. Suppose we have a pair of entangled photons, labelled A and B, each of which is sent to a distant destination by optical fibre. At most, we now have roughly 200 km separating A and B. If we wish to send a quantum state, encoded on particle C, from A to B, we perform a special type of measurement (called a Bell state measurement) on particles A and C, which immediately projects particle B into the original quantum state of C, modulo a unitary transformation. A classical communication channel is then used to relay the result of the Bell state measurement (BSM), which is required to correct the unitary transformation of the transmitted state. This is the basic principle of quantum teleportation, which allows us to avoid physically transmitting delicate states by fibre. We are still limited to relatively short communication distances, however, until we begin using quantum repeater nodes. If there are two pairs of entangled photons, say A-B and C-D, a quantum repeater performs a BSM on B and C. The result is that A and

¹Free-space transmission of light results in less attenuation, even in atmosphere, but creating a global network of crisscrossing single-photon streams over the surface of the planet has obvious drawbacks. Optical fibres are really the only realistic option, in spite the issues they themselves present.

D, which could now be separated by up to 400 km, automatically become entangled. By using many repeater nodes and entangled pairs, the quantum communication distance can be extended far beyond the traditional 100 km limit.

In this thesis, I will introduce the concept of a new type of hybrid quantum repeater, which merges the photonic and spin qubit platforms [2]. A qubit, which is short for quantum bit, is the quantum analog of a digital bit. A digital bit can be either a 0 or a 1. Similarly, a qubit is a two-level system which can be in state $|0\rangle = (1, 0)^T$, state $|1\rangle = (0, 1)^T$, or some superposition of those two. In our hybrid system, photonic qubits (encoded in the photon polarisation state) are sent to a photon-to-spin conversion interface, where the qubit is transferred to the spin state of a charge carrier, such as an electron. Two spin qubits, each generated by a photon from a different entangled pair, can then be entangled, causing the two remaining photons to become entangled.

This hybrid system has several advantages over past quantum repeater designs, including non-destructive heralding of successful qubit transfer, long coherence times of spin qubits (which means longer quantum memory if one photon arrives before the other), and the potential to fully entangle the spin qubits. The results presented in this thesis summarise the progress which NRC has made towards creating a photon-to-spin conversion interface, where photonic qubits are translated into spin qubits. In particular, I will review my work on the optical components of this project, such as generating fibre-coupled single photons and controlling photon polarisation in fibres.

The creation of a hybrid photon-to-spin quantum repeater will require the combination of expertise and technologies in several fields, ranging from photonics and fibre optics to semiconductor quantum computing techniques. Although I have focused primarily on the optical aspects of this project during my master's degree, much of my research has had a significant overlap with other areas as well. In this chapter, I will attempt to provide a comprehensive review of the physics and technologies which underlie the experiments I have conducted, with the hope that the reader will better understand the contents of the remaining chapters.

To begin, Sections 1.1 and 1.2 will look at two types of quantum dots: lateral dots and dots embedded in photonic nanowires. In the photon-to-spin setup, the lateral dots will be used to confine the carrier spins and perform entangling operations. The overview of their properties will be relatively cursory, as the focus of my work has mostly been on the optical components of the project. We will look in far greater detail, however, at nanowire quantum dots, including the fabrication process and their photoluminescent properties. These devices have long been known to act as sources of indistinguishable single photons, which would be necessary for many quantum communication and optical quantum computing applications. In our experiments, we investigated the possibility of directly coupling nanowire emissions to an optical fibre, which might someday allow for

the creation of plug-and-play fibre-coupled sources of single photons. These fibre-coupled dots could provide the individual polarised photons needed for testing the photon-to-spin system, as well as the entangled photons to be used in the final version of the project.

Section 1.3 describes two devices used for charge sensing: the quantum point contact (QPC) and the radio frequency quantum point contact (RF-QPC). While both allow for detection of the carrier population in a quantum dot, the RF-QPC can give sub-microsecond resolution of charge events. In our photon-to-spin experiment, an RF-QPC will be used to detect the generation of a carrier in a quantum dot upon the arrival of a polarised photon.

The next three sections discuss the optics topics which are relevant to the photon-to-spin system. A brief review of polarisation, both classical and quantum, is provided in Section 1.4. We will also look at polarisation control using waveplates, as they are a major component of our system and will feature heavily in Chapter 5. Section 1.5 goes over the mathematics used for modelling polarisation, calculating the fidelity of one polarised state with respect to another, and describing the effects of an optical fibre on polarisation. These equations are an important part of our technique for compensating for the polarisation-altering effects of the fibres used in the photon-to-spin system. To give a broader understanding of polarisation in fibres, Section 1.6 contains a review of the literature in the area of fibre-based transmission of polarised light.

Finally, Section 1.7 will describe in detail our proposed photon-to-spin quantum repeater, including comparisons to other repeater designs and descriptions of challenges which must be overcome in order to successfully create such a system. The goal of my work has been to help the NRC progress towards the realisation of this system, with particular focus on the optical components. The remainder of this thesis will discuss the various portions of the photon-to-spin project upon which I have worked during my time at NRC.

1.1 Lateral, Gate-Defined Quantum Dots

Spin qubits, such as those used in the photon-to-spin system, are encoded in the spin states of charge carriers. In order to perform operations on them, we must be able to contain these carriers within a small, localised region, and the 3-dimensional confining potential used to accomplish this task is called a quantum dot (QD). A quantum dot is formed by manipulating the electric potential of a material so that charge carriers, such as electrons, may be confined to nanoscale regions, often considered 1D point wells (hence the term ‘dot’). The electric potential of the quantum dot is lower than in the surrounding material, either because the bulk material only permits higher-energy states or because the dot is confined by barriers of large electric potential. The trapped carriers are therefore prevented from escaping the well unless they receive additional energy. The dimensions and material properties of the quantum dot determine the number of carriers which may be trapped within it, although some dot devices are tunable, allowing one to manually alter their carrier populations. Whereas quantum dots historically confined a multitude of charge carriers, modern dot designs allow for the isolation of individual electrons or electron holes [3].

While there are many types of quantum dot structures, such as self-assembled dots based on crystal strain and artificially-grown crystal heterostructures, in this section we will focus on one particular design which may have many potential uses in the field of quantum computing: the lateral quantum dot. This will be merely a brief review of lateral dots; though they are an important component of the photon-to-spin project, my work with them was minimal. For an excellent and more-detailed introduction to these devices, I recommend the textbook *The Physics of Low-Dimensional Semiconductors: An Introduction* by Davies [4].

1.1.1 Structure and Formation of a Lateral Quantum Dot

The lateral quantum dot is based upon the concept of a two-dimensional electron gas (2DEG) and the manipulation of the local electric potential. This type of sample, an example of which is shown in Figure 1.1, is composed of two (or more) semiconductor layers. The material properties of this structure cause free electrons to become trapped at the intersection (called a heterojunction) of the two materials. The electrons move freely and with minimal interaction in the 2D plane of the heterojunction, but they are confined/quantized in the third dimension. This is the physical realization of the quantum well. In our samples, the 2DEG is formed between layers of gallium arsenide (GaAs) and aluminum gallium arsenide (AlGaAs).

Gold electrodes called gates are deposited on top of the wafer (shown in light blue in Fig. 1.1). By applying a sufficiently negative voltage to a gate, an electric field is generated which forces the electrons to vacate the region of the 2DEG directly beneath it. If multiple gates have voltage applied to them, large areas of the heterojunction can become depleted




Figure 1.1: Scanning electron microscope (SEM) image of a gate-defined, lateral double quantum dot. The dots, whose approximate locations are outlined by the dashed circles, are formed in the 2DEG by the proper application of voltage to the gates (marked in light blue). The 2DEG is sandwiched between layers of AlGaAs and GaAs in our samples. Its depth is exaggerated here for visualisation purposes; the actual 2DEG has no height.

of electrons. Through proper manipulation of the gates, a region containing just a few electrons may be cut off from the rest of the 2DEG, forming a finite potential well (i.e. a quantum dot). With the proper materials and fabrication techniques, lateral dots can instead be made to trap individual electron holes, with the heterojunction hosting a two-dimensional hole gas. As we shall see in Section 1.3, the number of electrons/holes in a lateral quantum dot can be determined using a quantum point contact, which can act as a sensitive, on-chip charge detector.

Unlike dots which are defined by the fixed, inert barriers of the sample’s material structure alone, gate-defined lateral dots are highly versatile. The sample shown in Figure 1.1 can form two quantum dots, the locations of which are approximately outlined by the dashed circles. This is called a double quantum dot (DQD). By altering the voltages applied to the gates, the sizes of the dots may be reduced to the point where they each contain just a single electron (or no electrons) [3]. The electrostatic barrier separating the dots can also be tuned, allowing interactions between their occupants. By creating large arrays of lateral dots, each uniquely initialized and the interactions between them carefully controlled, we would theoretically have a scalable quantum computing device.

As we shall see in Section 1.7, a major component of the photon-to-spin project is the use of lateral dots to trap photo-generated charge carriers, whose spins will correlate to

the polarisation states of photons incident on the sample.² While there is much to be said about the methodology and physics of electron trapping in a lateral dot, the main focus of this thesis will be on the optical components of the photon-to-spin system. We are therefore reviewing only the very basics of a lateral dot's operation, with the hope that the details provided here are sufficient to enable the reader to understand the most useful properties of these devices. In Chapter 2, we will briefly discuss some of the preliminary work which has been done with lateral dot samples as part of the photon-to-spin system.

²Section 1.2.3 describes in detail the principle of photo-excitation, where the absorption of a photon by the semiconductor material excites an electron to a higher energy state.

1.2 Nanowire Quantum Dots

Nanowires are long, thin structures which have diameters on the order of several nanometers and length-to-width ratios typically exceeding 20:1. They may be composed of semiconductor materials, metals, or superconductors, and they may be fabricated to include such features as quantum dots [5], waveguide shells [6], and/or branches [7]. The many optical, electrical, and mechanical properties of nanowires make them versatile devices with a broad range of applications, from use as high-gain photodiodes [8] to coatings for brain implants [9].

It is, however, their potential usefulness as reliable sources of single photons and entangled photon pairs which makes them particularly relevant to the field of quantum communication [10–14]. Quantum dots embedded in photonic nanowires are proven sources of highly-indistinguishable single photons, which could be used as flying qubits. In our study of these devices, we looked at the potential for integrating them into a fibre-based quantum network, as described in Chapter 4. Sources of such particles could be very useful in our tests of the photon-to-spin system, which will require individual photons to be delivered by fibre to a lateral quantum dot.

Over the next several pages, we will review the basic properties of photonic nanowires and nanowire quantum dots. In particular, we will focus on: the fabrication methods for our samples, the physics of photoluminescence (PL) through exciton generation/recombination, and a technique for verifying single-photon emission. An excellent overview of the state of research in the field of photonic nanowires can be found in the review article *Nanowire-based sources of non-classical light* by Dalacu,³ Poole and Williams [10]. Lodahl, Mahmoodian and Stobbe provide a detailed discussion of single-photon emission by semiconductor quantum dots in their article *Interfacing single photons and single quantum dots with photonic nanostructures* [13].

1.2.1 Introduction to Photonic Nanowires

Nanowires are tall, thin structures which exhibit many interesting electrical, mechanical and optical properties, making them incredibly useful for a wide variety of applications in quantum optics, nano-scale electronics, and even biology. The samples we will look at here are photonic nanowires with embedded quantum dots, which are grown for the purpose of directed single-photon emission [11]. Figure 1.2a) shows an example of one of these nanowires. The device is a tapered, cylindrical tube with a total height of about $5\ \mu\text{m}$ [6]. As we can see in Fig. 1.2b), nanowires are often grown in large arrays so that we may test many of them and determine which ones have the desired properties.

The main components of the photonic nanowire are shown schematically in Figure 1.2c). Nanowires usually have two main components: the core and the shell. The core is

³Dr. Dalacu is largely responsible for most of my knowledge on the subject of nanowires, so thank you, Dan, for our many conversations on the subject.




Figure 1.2: a) SEM image of a nanowire clad with a waveguide shell. b) Array of positioned, waveguide-clad nanowires with 10 μm pitch. (SEM images provided by Dan Dalacu.) c) Schematic cross-section of a nanowire, with quantum dot shown in purple, nanowire core in orange, and waveguide shell in blue. (Relative radii of core and shell are not to scale.)

typically quite narrow (about 20 nm in our samples) and is grown first. Once the core has reached a sufficient height, the fabrication conditions are altered such that radial growth is promoted instead, cladding the nanowire in a shell.⁴ Rather than just growing a wire with the width of the shell, there are several important reasons for growing the narrow core first, such as maintaining crystal structure purity. We shall cover these reasons in the next section, where we will look at the fabrication process in greater detail.

Our wires are also equipped with quantum dots, which are made by briefly switching to a different growth material while making the core. In our samples, the substrate, core and shell are all composed of InP, while the quantum dot is nominally InAs (although residual phosphorus present during growth usually results in an $\text{InAs}_{0.25}\text{P}_{0.75}$ dot). The dot is therefore a small, flat disc of InAs, surrounded on all sides by uniform InP. Unlike a gate-defined lateral quantum dot, which electrostatically isolates electrons from the rest of the 2DEG, the quantum well of this dot arises from the difference in the electric potential between InP and InAs. Ambient charge carriers become trapped in the InAs layer as they propagate through the wire.

Our fabrication technique allows for the number of dots and their positioning within the wire to be engineered precisely. The nanowire experiments described in Chapter 4 involved both single- and double-dot samples. The core and dot have diameters of about 20 nm, with the dot(s) situated about 200 nm from the base of the nanowire [11]. The shell, on the other hand, has a diameter of about 200–250 nm at the base, although it tapers to a

⁴Given all the applications for which they are used, there is an incredible variety of nanowire designs. Some samples are grown with the cores left bare or only partially-clad, while in other samples, there are branches off of the main ‘trunk’ of the core [7]. Our samples, however, always have the design described here.

point near the top [15].

As mentioned earlier, one of the main purposes of including the quantum dot is to generate single photons. This is accomplished through the process of photoluminescence (PL), which is the emission of light via optical stimulation. Section 1.2.3 describes this phenomenon in more detail, but in essence, it begins with a laser of sufficiently high photon energy (usually > 1.5 eV in our experiments) generating electron-hole pairs within the InP of the wire. When an electron and hole fall into the quantum dot, they form an unstable bound state called an exciton, and the exciton's decay releases energy in the form of a photon. The exciton wavefunction, which is dictated by the dot's material composition and geometry, determines the wavelength of the emitted photon. Our samples typically emit at around 950 nm, with very little spread (~ 0.024 nm [11]) in the emission lines of a given dot. We will discuss excitons and photoluminescence in more detail in Section 1.2.3.

When an exciton decays, the photon may be emitted in any random direction, including radial emission towards the sidewalls of the nanowire. Without some method of collimating the dot emissions, most of the light would be lost, and it would not be possible to collect it all using an external optical system. The photonic nanowire therefore incorporates a shell, which acts as a waveguide for the PL photons. The difference in the index of refraction between the shell and the exterior environment (which is typically vacuum in our experiments) ensures that light is internally reflected, so the photons are directed to the nanowire tip.⁵ With the emitted photons coupled to the fundamental waveguide mode (HE_{11}), the tapering of the wire slowly causes their spatial mode to expand outside the wire as they propagate [6]. This evanescent coupling to the external mode reduces the likelihood of reflection once the light reaches the tip, with narrower tapering angles ($\sim 1^\circ$ – 2°) yielding better coupling to the external mode. The waveguide geometry also causes the nanowire emissions to have a Gaussian profile and a numerical aperture (NA) of 0.5 [16], implying a divergence angle of 60° . So by spreading out the wires, as in the array of Fig. 1.2b), the waveguide allows us to isolate and collect with great efficiency the individual photons emitted by a specific nanowire quantum dot.

1.2.2 Fabrication of a Photonic Nanowire with Embedded Quantum Dot

While there are many different types of nanowires, they can be broadly classified into two categories from a fabrication standpoint: top-down and bottom-up [10].⁶ Generally speaking, a top-down device is created by etching away a bulk material, leaving only a nanowire. This method might be used, for example, to excavate the substrate such that

⁵Half of the photons are actually directed downwards and are lost to the substrate, although this can be remedied by inserting a metallic mirror at the base of the nanowire [15].

⁶There are other fabrication options as well, such as nanowire branching [7] and template-assisted nanowire growth [17]. (Technically they are both bottom-up approaches, but they are fairly unique in their methods.)

the wire contains a self-assembled quantum dot which was originally in the bulk material. Unfortunately, top-down wires may have etching-induced damage [5] and tend to be non-uniform, both with respect to each other and in the shape of an individual wire. As they are usually constructed around self-assembled quantum dots or nitrogen-vacancy centres in diamond, top-down wires often contain several emitters which are randomly positioned throughout the structure, and they emit photons with varying efficiencies and wavelengths [10].

The bottom-up approach, on the other hand, involves growing a nanowire crystal from scratch on the substrate. This technique has several advantages, such as site-selective growth (i.e. predetermined wire positioning) [18], as well as highly uniform nanowires and dots. This technique allows for the quantum dot dimensions, composition, and position within the wire to be chosen with high precision, and the number of dots per wire can be selected as well [5, 10]. The nanowire samples used in our experiments are all grown from the bottom up using a combination of two growth methods: selective-area (SA) epitaxy and vapour-liquid-solid (VLS) epitaxy. With SA growth, a patterned mask is deposited on the substrate, which ensures nanowires only grow at specific locations. Unfortunately, using this technique alone usually results in wires which have indistinct, multi-layer quantum dot boundaries. VLS epitaxy, on the other hand, uses randomly-placed metallic seed particles to catalyse nanowire growth and promote sharp interfaces between the quantum dot and host semiconductor. Combining these two methods allows for controlled positioning of the nanowires and well-formed quantum dots.




Figure 1.3: Selective-area patterning of substrate for nanowire growth. a) The first step is to layer the SiO_2 mask (blue) and resist (orange) on top of the InP substrate (red). The area highlighted by the grey circle is the electron beam pattern. b) After etching in an HF solution, the resist has a small hole and the mask has a large one. The sample layers are separated in the diagram to show the larger hole in the SiO_2 mask. c) Placing a gold layer (yellow) on top of the resist causes a gold droplet to be deposited on the sample substrate. d) With the resist removed, we are left with a positioned gold drop at the centre of a hole in the mask.

The process for growing our nanowire samples begins with the preparation of a patterned growth substrate. The steps for this procedure are shown in Figure 1.3. First, a 20 nm layer of SiO_2 is deposited on the (111)B surface of a S-doped InP wafer [5, 18, 19]. A

resist material is placed on top and patterned using electron-beam lithography. A typical e-beam pattern for our samples is an array of circles, and a buffered HF solution is used to etch openings in the resist at the locations of these circles. The diameters of the holes in the resist ($\sim 20\text{--}40$ nm) will determine the widths of the nanowire cores. The resist openings expose the SiO_2 to the HF, so the mask is etched away at these points too. As the diameters of the SiO_2 openings (~ 250 nm) dictate the widths of the waveguide shells, the sample is left in the HF solution long enough to remove a larger area of the SiO_2 from beneath the resist. This ensures that the nanowire shell diameter is much larger than that of the core.

Once the sample has been removed from the HF, a layer of gold is deposited on top of the resist in a metal lift-off process. The gold falls through the holes in the resist, resulting in gold droplets placed directly on the InP substrate. The resist is then removed, leaving only the substrate, the SiO_2 mask, and a small drop of gold at the centre of each hole in the SiO_2 . To clean the bare InP surface in those regions where the SiO_2 has been removed, the sample is placed under a UV lamp, which generates ozone in the air [19]. The O_3 is highly reactive with the InP surface, coating it thoroughly. This ozone layer is then removed, etching away some of the InP and removing any remaining SiO_2 in the growth regions.

With the growth sites set by the gold lift-off process, the next step is to begin growing the nanowire cores using VLS epitaxy [5, 11, 18, 20]. The vapour stage is initiated with chemical beam epitaxy in a vacuum chamber, where precracked⁷ PH_3 provides the phosphorus and trimethylindium (TMI) provides the indium. The vaporous indium and phosphorus dissolve in the gold,⁸ which brings us to the liquid stage. Our samples are typically grown at temperatures between 420 and 450 °C, which is too low for growth on the bare surface of the substrate (this would be vapour-solid growth). The gold catalyses the reaction, however, and the absorbed indium and phosphorus are deposited in solid InP layers at the interface between the gold and the substrate, bringing us to the solid phase. The removal of vapour by the gold creates a gradient, which draws more material towards the wire and enables the process to continue. This is VLS epitaxy.

Bulk InP, as in the substrate, has a zincblende (ZB) crystal structure [10], where the unit cell is composed of two overlapping face-centred cubic cells, each composed of a different type of atom. Each atom is bonded to four others of the opposite species in a tetrahedral configuration. Sufficiently thin nanowires tend to grow with a pure wurtzite (WZ) structure, however, as long as the TMI levels are low enough [11]. The WZ unit cell is similarly composed of two interlaced cells, each of which has a hexagonal close-packed structure and is composed of one of the two atomic species. So using VLS epitaxy, the InP nanowire core grows axially through the successive deposition of atomic layers, each of which (ideally) has a WZ structure.

⁷Cracking: the carbon bonds of the organic precursor are broken. In this case, it occurs prior to entry into the vacuum chamber.

⁸The other elements of the III, V sources, such as hydrogen, are generally left behind as the In and P enter the gold.




Figure 1.4: The stages of nanowire growth. a) A gold drop (yellow) is deposited on the substrate (red) at the centre of the hole in the SiO₂ (blue). b) The gold drop catalyses the growth of the InP nanowire core (purple), with the core diameter equal to the diameter of the gold. c) Switching to InAs, the quantum dot (green) is grown near the nanowire base. d) Return to InP core growth. e) When the wire is sufficiently tall, alter the growing conditions to begin radial growth of the waveguide shell. The shell diameter will be equal to the diameter of the opening in the SiO₂.

The growth stages of our nanowire samples are shown in Figure 1.4. The WZ core, which is represented by the solid purple column in Fig. 1.4b)–e), necessarily has the same diameter as the gold seed (since the Au-InP interface is the only place where growth occurs at these temperatures). After 15–20 minutes, the core is roughly 200 nm tall, and the PH₃ is swapped for AsH₃ to create the InAs quantum dot [11, 15]. With an AsH₃ overpressure for 2–3 seconds, the quantum dot height will be about 5–7 nm tall. Residual phosphorus in the gold will be incorporated into the dot, roughly yielding InAs_{0.25}P_{0.75} layers, with the P concentration trailing off as the dot grows. InP core growth is then resumed for another 10–15 minutes, reaching a height of 4–5 μm by the end. (Some residual As will also penetrate the P layers just above the quantum dot.) Multiple dots can be grown in a single wire by briefly switching back to InAs growth at the appropriate intervals.

Comparing the growth times for the core before and after the dot (20 min vs. 15 min), it is evident that the growth rate is non-linear, with the wire growing exponentially faster as its height increases [18]. This process can be modelled by taking into account the methods by which the indium and phosphorus are absorbed by the catalyst. Although vaporous group III elements, such as the P and As, may occasionally bind directly to an InP surface

at these temperatures, they tend to be scattered instead [20]. They therefore contribute to nanowire growth only when absorbed directly by the gold particle at the top of the wire. Indium, on the other hand, will diffuse across an InP surface. Therefore, in addition to direct absorption by the catalyst, indium can also reach the gold by moving across the substrate and nanowire core. The indium diffusion length, which is highly dependent upon the growth conditions, is roughly $0.5\text{--}1\ \mu\text{m}$, but importantly, the indium will only diffuse across an InP surface. While a sample made without the use of SA epitaxy would allow for indium diffusion across the bare substrate, the SiO_2 mask on our samples reflects the TMI. We therefore have three processes by which indium reaches the gold particle:

- 1) **Direct Flux:** Indium is directly absorbed by the gold particle.
- 2) **Direct Sidewall Flux:** TMI makes contact with the sidewall of the nanowire core. The indium cracks and diffuses to the catalyst.
- 3) **Indirect Sidewall Flux:** The TMI is reflected off the SiO_2 mask, it hits the sidewall, and the indium moves along the surface to the gold.

Evidently these last two mechanisms become more relevant as the nanowire grows (especially indirect sidewall flux), since the sidewall area becomes much larger over time and far exceeds the surface area of the gold droplet. This is the reason for the non-linear growth rate. Figure 1.5 represents the three transport methods visually. As the samples are fabricated with an overpressure of PH_3 , the TMI is the limiting factor in the growth process, i.e. the instantaneous growth rate is dependent only upon the efficiency of the indium transport mechanisms for a given core height.




Figure 1.5: The three methods by which indium is transported to the gold catalyst: direct, direct sidewall, indirect sidewall.

One last thing to note about indium transport is that it binds to the surface after diffusion. If the diffusion length does not allow an indium atom to reach the gold, it will

settle on the surface of the nanowire core, resulting in radial growth [18]. Taller nanowire cores may therefore have a tapered appearance, as the particles closest to the gold are used for axial growth, while those at the base of the wire bind to the surface.

The final step in the fabrication process is to grow the waveguide shell. There are two methods which can be used to switch from axial to radial growth: increased temperature [5, 15] or an increase in group V flux [11]. By increasing the temperature to 500–515 °C, we move into the regime where vapour-solid growth is dominant, and the VLS process of the gold catalyst is suppressed. Alternatively, by increasing the PH₃ flow by a factor of 3, radial growth becomes more likely than axial growth catalysed by the gold. Under these conditions, the surface of the nanowire begins to grow outward, with the WZ structure of the core seeding the shell. The opening in the SiO₂ dictates the diameter of the shell, which swells to fill it at the base of the wire. Without the mask, the surrounding substrate would also be growing, which would use up resources and prevent proper shell growth. Figure 1.4e) shows the base of the nanowire when clad by the shell. Note that the shell fills the annulus surrounding the wire and has a WZ crystal structure (hence its purple colouring in the image instead of red, like the ZB substrate). The InAs(P) quantum dot is now completely surrounded by wurtzite InP.

The bare substrate surrounding the core contributes a lot to the shell growth, so the shell is initially conical in shape. As the shell fills out to the width of the opening in the mask, it becomes cylindrical near the base and tapered at the top [15]. When growth is stopped, the shell is roughly 200–250 nm in diameter at the level of the quantum dot, but just 20–30 nm wide at the tip. The shell growth time is roughly 94 min, and the growth conditions favour radial growth so heavily that the nanowire may only grow another 200 nm in height during this period [11]. The conditions for shell growth are calibrated to achieve very specific nanowire widths and tapering, which are necessary to achieve a Gaussian emission profile for PL photons and a numerical aperture of 0.5 [6, 16].

The combination of SA and VLS epitaxies therefore provides us with sharply-defined quantum dots embedded in positioned nanowires which are clad in waveguide shells. The gold catalyst of the VLS procedure allows core growth at lower temperatures while suppressing substrate growth. Importantly, the presence of the patterned mask, borrowed from the SA technique, ensures that shell growth takes precedence over substrate growth, allowing for complete and finely-tuned cladding of the nanowire core. The nanowire samples used in our experiments were grown by Dan Dalacu at the NRC.

1.2.3 Photoluminescence and Exciton Complexes in a Nanowire Quantum Dot

Photoluminescence (PL) is a property of certain materials whereby they emit light in response to stimulation by light of a lower wavelength. In particular, PL operates on the principle of photoexcitation, where an electron is promoted to a higher energy level by absorbing a photon of sufficiently high energy. Upon relaxation of the electron, a second

photon is emitted, albeit with a different energy from the first if there are multiple steps in the relaxation process (such as phonon emission or multi-photon emission).

Free Excitons

When an electron from the valence band of a semiconductor is excited to the conduction band, the vacancy left behind in the valence band is called an electron hole. We choose for simplicity to model the system in terms of the behaviour of this one hole, which acts as a positively-charged quasiparticle, rather than attempting to account for each of the millions of electrons present in the valence band. Both the conduction-band electron and valence-band hole are able to move about the material freely, as their bands are largely devoid of their respective particle species. Due to spin-orbit coupling, where the orbital angular momentum generates an electric field, which in turn interacts with the spin of the particle, semiconductor valence bands tend to have three sub-bands: light holes, heavy holes, and the split-off band [4]. In units of \hbar , both light and heavy holes have total angular momentum $\mathbf{J} = 3/2$, but the angular momentum projection in the z-direction is $\mathbf{J}_z = \pm 1/2$ for light holes, while $\mathbf{J}_z = \pm 3/2$ for heavy holes. These two types of holes are named for their relative effective masses. The split-off band, which has total angular momentum $\mathbf{J} = 1/2$, is named for the fact that it is usually much lower in energy than the other two. In our nanowire samples, the combination of cryogenic temperatures and confinement due to the quantum dot induces a significant separation between the heavy hole, light hole, and split-off sub-bands [21]. As we study only the ground-state emissions of the wires, we are therefore going to consider just the heavy holes in this section.

The ground state of a valence band (heavy) hole is the highest energy level (since all the valence electrons, which are real particles, are in the lowest energy states possible). A free electron is, of course, in its ground state when it occupies the lowest energy level of the conduction band. In a direct band gap semiconductor, the valence band energy is highest at the same crystal momentum as the lowest point in the conduction band. Since photons produce very little crystal momentum, a valence electron can be excited to the conduction band merely through the absorption of a photon with energy equal to the band gap. The materials from which we grow our nanowires, InP and InAs(P), are direct band gap semiconductors (as is GaAs, which forms part of our lateral quantum dot samples). Indirect band gap materials have a mismatch between the valence band maximum and conduction band minimum, so a change in crystal momentum is required to excite an electron from one to the other. Again, photons do not impart crystal momentum, so a phonon must either be emitted or absorbed in order for an electron to transition between the two band extrema. As the coincidence of simultaneous photon and phonon absorption is relatively rare, indirect band gap materials aren't often used for purposes which demand high photoexcitation efficiency.⁹

⁹At high temperatures, however, more phonons are present in the sample, increasing the chances of photoexcitation occurring. Alternatively, higher-energy photons will cause electrons in these materials to be promoted to much higher energy levels in the conduction band without changing the crystal momentum.

Carrier recombination occurs when an electron in the conduction band recombines with a hole by emitting energy. There are several types of recombination processes, although here we will mostly focus on radiative recombination, where the electron releases the energy in the form of light. Photoluminescence is therefore the process whereby the material is optically stimulated, by a laser for example, generating an electron-hole pair which eventually recombines by emitting a photon. The period between excitation and recombination is called the carrier lifetime, which varies depending on the material and the energy of the exciting photons. If the energy of the laser photons is much larger than the band gap, the electron will be promoted to a state far above the conduction band minimum. It may then release energy in the form of heat (i.e. phonons) as it relaxes to the bottom of the conduction band, before re-emitting a photon during recombination. In this process, called above-band excitation, the emitted photon will be lower in energy than the laser light. With resonant excitation, the energy of the laser photons is equal to the band gap.

As they have opposite charges, a free electron and a free hole may attract each other and form a bound state. The bound state of an electron and a hole is a hydrogenic quasiparticle called an exciton, and its properties are somewhat similar to positronium, the metastable bound state of an electron and a positron. In a solid, excitons may be formed when Coulomb forces attract the two particles and cause the surrounding electrons to push them together, creating a bound state. Like positronium, excitons are unstable and tend to decay over a period of a few nanoseconds. In a direct band gap material, exciton recombination results in the emission of a photon. The bound state exciton has a slightly lower energy than a free electron-hole pair, however, so it emits lower-energy photons during recombination.

Trapped Excitons in Quantum Dots

The presence of a quantum dot in the system introduces several new factors, starting with the fact that an electron-hole pair in a dot always forms an exciton [22]. The quantum dots in our nanowire samples are roughly 18–20 nm wide and 5–7 nm tall [11], putting us in the strong coupling regime where the dot dimensions are confining the spatial extent of the exciton. In this regime, the quantization energy of the dot far exceeds the Coulomb binding energy of the exciton.

The exciton trapped in a quantum dot may have four spin states quantized along the growth direction, which is the direction of greatest confinement: $|\uparrow\uparrow\rangle$, $|\uparrow\downarrow\rangle$, $|\downarrow\uparrow\rangle$, and $|\downarrow\downarrow\rangle$. Here we are using the conventional representation of a single arrow ($\uparrow\downarrow$) for the electron spin state: up for $+\hbar/2$, and down for $-\hbar/2$. A double arrow ($\uparrow\downarrow$) is used to represent the heavy hole angular momentum: up for $+3\hbar/2$, and down for $-3\hbar/2$.

With an electron spin of $1/2$ and heavy hole angular momentum of $3/2$, the states with anti-parallel spins have total angular momentum ± 1 . To conserve angular momentum, the emission of a photon (which is spin-1) can only occur when the total spin of the exciton is ± 1 . The $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$ states are therefore called the bright states, as each recombination results in the emission of a single photon. As we will discuss in Section 1.4.1, the spin

of a photon is equivalent to its polarisation. The most natural basis is typically circular polarisation, with a spin of $+1$ corresponding to right-circular polarisation (σ^+), and -1 spin corresponding to left-circular polarisation.

The parallel-spin states, with total spin ± 2 , are called the dark states [23]. They are optically inactive because two-photon emissions do not occur for the recombination of a single electron and hole. (Such a two-photon event would require an additional intermediate energy level within the band-gap, which would allow the electron to do a two-step recombination, but our devices do not have this additional level.) One of the spins may eventually flip, however, allowing recombination.

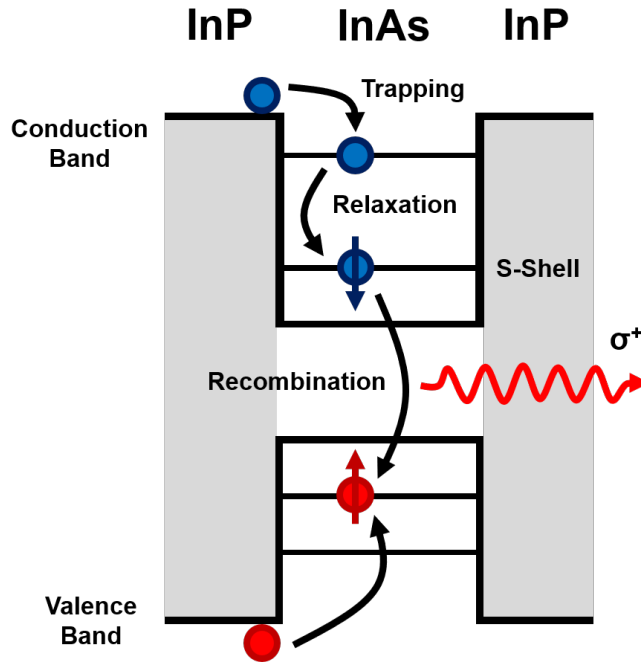


Figure 1.6: The stages of exciton trapping, relaxation, and recombination in a quantum dot for above-band excitation. In this case, the angular momentum of the ground-state exciton is $+1$, resulting in a photon with right circular polarisation.

While resonant excitation at the exciton binding energy generates a single ground-state exciton directly within the quantum dot,¹⁰ we tend to use above-band pumping of our samples, which causes many electron-hole pairs to be generated within the nanowire. Some number of them may recombine within the nanowire or the bulk InP of the sample substrate, but occasionally an electron and a hole will fall into the quantum dot, usually on a time scale of about 10–100 ps. Once trapped in the dot, they relax non-radiatively through phonon emission to the ground state over a period of tens of ps, forming an exciton.

¹⁰If the exciton is generated within the dot, but in a higher-energy state, it is called quasi-resonant excitation.

As the exciton lifetime within the dot is a few nanoseconds, which is much longer than the relaxation period, we typically only observe photons emitted by the recombination of ground-state excitons. Figure 1.6 shows this trapping-relaxation-recombination process of excitons in a quantum dot. Evidently, directly populating the ground state using resonant excitation results in somewhat shorter lifetimes and more consistent delays between the excitation pulse and recombination.

As would be expected from a particle trapped in a quantum well, the photon emission energy is highly dependent upon the dot geometry and composition. The small, InAs(P) dots in our samples, for example, tend to yield PL photons with wavelengths around 950 nm. When we examine a large array of dots, however, the unique properties of each dot give us a spread of ground state emissions between 900 nm and 1000 nm. It has been shown that increasing the size and arsenic content of our quantum dots produces higher-wavelength photons [24].

The charge environment of the quantum dot also affects the photon energies. In particular, the presence of additional carriers within or near the dot may shift the emission energy by several meV. The bound states formed by combinations of trapped excitons and carriers are called exciton complexes. The ground-state (s-shell) emissions of a quantum dot arise from four types of exciton complexes: the neutral exciton, the charged exciton, the biexciton, and the charged biexciton. These are depicted visually in Figure 1.7.




Figure 1.7: Examples of the four most common ground-state exciton complexes observed in our nanowire samples: the neutral exciton (X), the negatively-charged exciton (X^-), the biexciton (XX), and the negatively-charged biexciton (XX^-).

The neutral exciton, represented by the symbol X (or sometimes X^0), is merely a normal exciton composed of one hole and one electron. It may be formed by either a spin-up electron and spin-down hole, $|\uparrow\downarrow\rangle$, or vice versa, $|\downarrow\uparrow\rangle$. Fine structure splitting due to the presence of the single electron spin in a slightly asymmetric dot results in each of these states having a slightly different emission energy [23, 25]. This splitting is usually below

the resolution of our spectrometer, although it may be observed as a slight broadening of the emission line.

The charged exciton arises when another carrier is added to the dot. In our S-doped samples, this tends to be another electron, especially since the holes have a much larger effective mass and move more slowly through the nanowire. The symbol for the negatively-charged exciton is X^- , although sometimes the charged exciton is represented by X^* . The addition of the electron modifies the binding energy of the system, so when analysing the emission spectrum, one can easily distinguish between a photon produced by X and a photon produced by the recombination of the hole and one electron of X^- . Unlike the spectral doublet formed by X recombination, the singlet state of the electron pair in X^- eliminates fine structure splitting [25], so radiative recombination of the X^- always results in the same photon energy. When a negatively charged exciton recombines, an electron is left behind, trapped in the dot.

While two excitons in a bulk material may have minimal interaction, two excitons trapped in a quantum dot form a four-particle bound state called a biexciton, which has the symbol XX . The excitons usually attract one another, so the molecular binding energy causes the energy of a biexciton to be slightly below that of two free excitons [22]. In the strong confinement regime of our small quantum dots, each exciton in the pair recombines individually, although only the first one recombines with the $XX \rightarrow X$ energy, as the second one is just a trapped neutral exciton. As the system recombines to one of the two non-degenerate neutral exciton states, the biexciton emission line also exhibits some splitting, even though the paired holes and paired electrons ensure that every ground-state biexciton in a particular dot has the same energy.

Finally, we have the charged biexciton, XX^- . The presence of an additional electron in the p-shell modifies the binding energy of the biexciton, but we are still only considering the recombination of a hole and an s-shell electron (not the p-shell electron). The emissions of this exciton complex are far more complicated, as there are eight possible X^- states after recombination: $|\uparrow_s\uparrow_p\uparrow\rangle$, $|\downarrow_s\uparrow_p\uparrow\rangle$, $|\uparrow_s\downarrow_p\uparrow\rangle$, $|\downarrow_s\downarrow_p\uparrow\rangle$, $|\uparrow_s\uparrow_p\downarrow\rangle$, $|\downarrow_s\uparrow_p\downarrow\rangle$, $|\uparrow_s\downarrow_p\downarrow\rangle$, and $|\downarrow_s\downarrow_p\downarrow\rangle$. Those states with electron spins in opposite directions experience electron-electron exchange interactions, and all states experience some degree of electron-hole exchange. The different energies of the possible final states therefore result in different photon energies, so we observe multiple, often widely-separated, peaks in the emission spectrum due to XX^- recombination.

These are the only s-shell exciton complexes, excluding positively-charged states. Higher laser powers will produce more free carriers in the sample, however, and may generate higher charge states, such as X^{--} or XX^{---} for example. Carriers trapped near the dot may also serve to alter the emission energy, even though they do not form bound states with the excitons. An unstable electrostatic environment, due to the movement of carriers throughout the nanowire, will result in time-varying photon energies. This results in a broadening of the emission lines in the PL spectrum [13]. As the excitation power increases, these higher-order charge states generate minor satellite peaks about the main emission

lines. A higher pump power can also result in p-shell emissions due to continuous filling of the s-shell, and d-shell emissions are occasionally observed in emission spectra as well. In Chapter 4, we will look at the emission spectra of actual nanowire quantum dots and identify the sources of the various emission lines, such as exciton complexes.

1.2.4 Confirming Single-Photon Emission with the Hanbury Brown-Twiss Autocorrelation Experiment

The autocorrelation of a stream of photons, or the statistical correlation of a photon stream with a time-delayed version of itself, is typically measured by the degree of second-order coherence, given by the second-order correlation function $g^{(2)}(\tau)$. It is a statistical measure of the intensity fluctuations in an electromagnetic field over time, and it is analogous to the first-order correlation function, $g^{(1)}$, which deals with statistical changes in the electric field strength. Classically, the $g^{(2)}$ function is defined as:

$$\begin{aligned}
 g^{(2)}(\tau) &= \frac{\langle E^*(t)E^*(t+\tau)E(t+\tau)E(t) \rangle}{\langle E^*(t)E(t) \rangle \cdot \langle E^*(t+\tau)E(t+\tau) \rangle} \\
 &= \frac{\langle E^*(t)E^*(t+\tau)E(t+\tau)E(t) \rangle}{\langle E^*(t)E(t) \rangle^2} \\
 &= \frac{\langle I(t)I(t+\tau) \rangle}{\langle I(t) \rangle^2}
 \end{aligned} \tag{1.1}$$

where the angled brackets denote the time average, $E(t)$ is the electric field strength at a time t , $I(t)$ is the intensity, and τ is the time delay between the two versions of the signal. The equivalence of the first and second lines derives from the assumption that the average field strength is constant over time. In the quantum regime, where intensity can be associated with number of photons $n(t)$, we may rewrite this equation as:

$$g^{(2)}(\tau) = \frac{\langle n(t)n(t+\tau) \rangle}{\langle n(t) \rangle^2} \tag{1.2}$$

This equation provides the correlation between the number of photons at a given point in the stream at time t and the number of photons at time $t+\tau$. Suppose we are considering the case where, on average, there is only one photon present, so that if we consider a time t which coincides with one of these photons, then $g^{(2)}(\tau) = \langle n(t+\tau) \rangle$. The function now rises to $g^{(2)}(\tau) = 1$ only at those values of τ where another photon is statistically likely to be present, and it will drop to zero for times when it is unlikely that a photon is present.

If we now consider the value of $g^{(2)}(0)$, we run into an issue. We are presuming that we are observing these photons with a single-photon detector, but for $\tau = 0$, we would have to make two simultaneous measurements of the same detector. We are correlating the photon

at time t with itself. Even if the stream consists of just one photon, which would yield $g^{(2)}(\tau) = 0$ for all other values of τ , we will still obtain a value of $g^{(2)}(0) = 1$ because we are measuring the same value (i.e. 1) for both $n(t)$ and $n(t + \tau)$. This is confusing because we normally expect to see $g^{(2)}(\tau)$ rise to 1 only when another photon is present in the stream, whereas it should be zero if there are no additional photons.

To avoid this issue, we could instead consider a very small value of τ instead of zero; this would give us $g^{(2)}(\sim 0) = 0$, which is roughly what we want. Alternatively, we can just divide the stream of photons using a 50/50 beam splitter and send each stream to a separate detector. We will label these detectors D1 and D2. With this arrangement, we make the first measurement at time t with D1, while the delayed measurement at $t + \tau$ is always made by D2. Equation 1.2 then becomes:

$$g^{(2)}(\tau) = \frac{\langle n_{D1}(t)n_{D2}(t + \tau) \rangle}{\langle n_{D1}(t) \rangle^2} \quad (1.3)$$

where n_{Di} is the number of photons detected at detector Di (i.e. D1 or D2). We assume the average number of photons measured is the same for both detectors because the 50/50 beam splitter statistically divides the photon beam equally. With this version of the equation, the value of $g^{(2)}(0)$ refers to simultaneous measurements at the two detectors. If we only consider those instances where D1 observes a photon, then a $g^{(2)}(0)$ value of 0 indicates that for every detection of a photon at D1, there are never any photons detected at D2. This is the signature of photon antibunching, where the photons in a stream are separated from each other and approximately equally spaced.

The physical realisation of this autocorrelation measurement is the Hanbury Brown-Twiss (HBT) experiment,¹¹ the setup for which is depicted in Figure 1.8a). In this type of test, the photon stream (from a nanowire, in our case) is sent to a 50/50 beam splitter, which directs the light to two single-photon detectors, D1 and D2. The two detectors are equidistant from the splitter. In this experiment, we are measuring the time delay between a detection at D1 followed by a detection at D2. If photons are detected at both D1 and D2 simultaneously, corresponding to a count at $\tau = 0$, then the quantum dot must be emitting multiple photons simultaneously. If there is always a non-zero time delay between a detection at D1 and the next detection at D2, then the quantum dot must be emitting one photon at a time. Zero counts at zero time delay is therefore the signature of single-photon emission which we wish to observe. By observing many events, we build up a statistical average for $g^{(2)}(\tau)$ at all τ .

¹¹While it may be somewhat pedantic to make mention of this, HBT interferometry is named for the two gentlemen who devised it, Hanbury Brown and Twiss, hence the absence of a hyphen between the Hanbury and Brown in the experiment name.

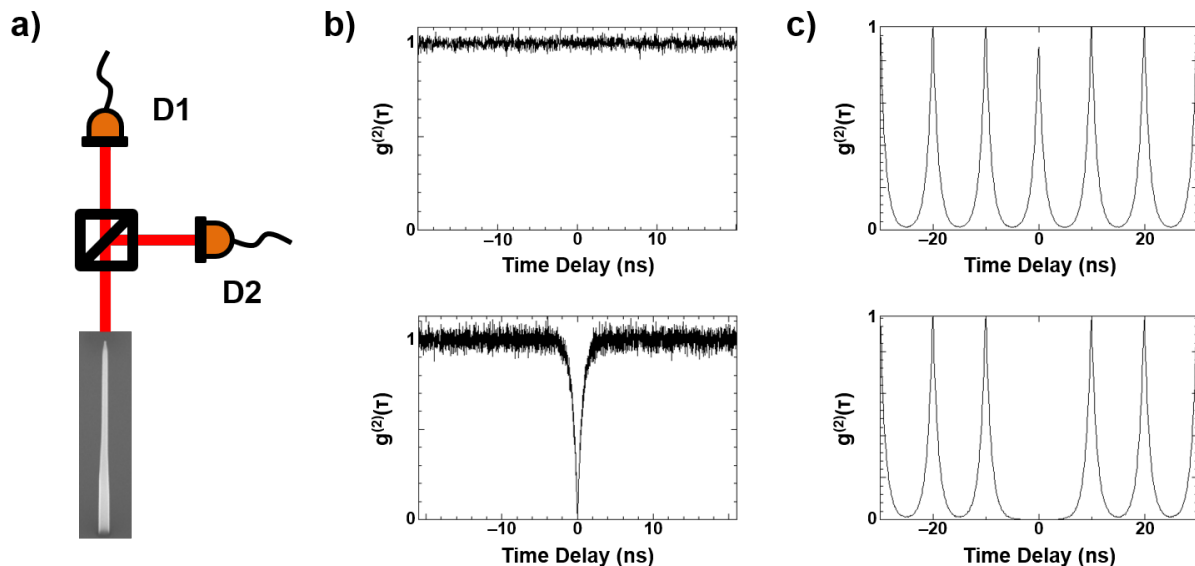


Figure 1.8: a) The setup used for a Hanbury Brown-Twiss autocorrelation experiment, with a nanowire as the single-photon source. b) Simulated $g^{(2)}(\tau)$ measurement with CW pumping. Top: Multi-photon emission. Bottom: Single-photon emission. c) Simulated $g^{(2)}(\tau)$ measurement with pulsed pumping. Top: Multi-photon emission. Bottom: Single-photon emission.

Figure 1.8b) shows Monte Carlo simulations of the $g^{(2)}(\tau)$ data acquired using the HBT setup for a continuous-wave (CW) laser source [26].¹² These plots can be understood as the likelihood of a detection by D2 at some time τ after a detection at D1. When the laser is in CW mode, excitons continuously form and recombine in the quantum dot, so the randomized photon emission times result in $g^{(2)}(\tau) = 1$ at large τ . If multiple photons are emitted simultaneously by the nanowire, then Equation 1.3 dictates that $g^{(2)}(0) = 1$, as can be seen in the upper image of Fig. 1.8b). If there is single-photon emission only, then $g^{(2)}(0) = 0$, as shown in the lower panel of Fig. 1.8b). A combination of multi-photon and single-photon events over time will result in $0 < g^{(2)}(0) < 1$.

In the case of a pulsed laser, the nanowire can only emit photons after it has been excited by a laser pulse. Ideally, each pulse should be short enough that there is only enough time for a single photon to be emitted per pulse. If a photon is detected by D1 at $t = 0$, then the next photon can only be generated and detected when the next laser pulse occurs. This produces regular peaks in the autocorrelation plot, as we see in Figure

¹²Note that although the HBT experiment describes τ as the time between a detection at D1, followed by a D2 detection, these plots extend to negative values of τ . While this seems to suggest a reversal of the D1-D2 order, this is not in fact the case. A delay line on D2 is used to observe these negative values of τ . Appendix E explains this issue in greater detail, but for now, simply consider the fact that the $g^{(2)}(\tau)$ function is symmetric about $\tau = 0$.

1.8c). The beam splitter divides the photon stream with a 50/50 probability, so there are some instances where photons are only directed to D1 for several pulses before a D2 detection occurs. As we are only looking at the delay between the first D1 detection and a D2 detection, this type of event is the reason why we observe multiple peaks in the $g^{(2)}(\tau)$ function, separated by the pulse period of the laser.¹³ Figure 1.8c) contains simulations of HBT measurements for a pulsed laser. If there is some multi-photon emission, as in the upper panel, an additional peak will appear at $\tau = 0$. We again get $g^{(2)}(0) = 0$ if there is only single-photon emission, as in the lower panel.

We now ask the question: what is the threshold value of $g^{(2)}(0)$ which indicates that single-photon events are occurring with statistically greater probability than multi-photon events? Returning briefly to Equation 1.1, we can rewrite the electric field in terms of creation (\hat{a}^\dagger) and annihilation (\hat{a}) operators [10, 13]:

$$g^{(2)}(\tau) = \frac{\langle \hat{a}^\dagger(t)\hat{a}^\dagger(t+\tau)\hat{a}(t+\tau)\hat{a}(t) \rangle}{\langle \hat{a}^\dagger(t)\hat{a}(t) \rangle^2} \quad (1.4)$$

Now let's re-examine the case where $\tau = 0$, the indicator of single-photon emission. We will label the n -photon Fock state at time t as $|n\rangle$. Noting that the time-average is equivalent to the expectation value in the quantum regime, the above equation gives us the following:

$$\begin{aligned} g^{(2)}(0) &= \frac{\langle n | \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} | n \rangle}{(\langle n | \hat{a}^\dagger \hat{a} | n \rangle)^2} \\ &= \frac{\sqrt{n}\sqrt{n-1} \langle n | \hat{a}^\dagger \hat{a}^\dagger | n-2 \rangle}{(\sqrt{n} \langle n | \hat{a}^\dagger | n-1 \rangle)^2} \\ &= \frac{n(n-1) \langle n | n \rangle}{(n \langle n | n \rangle)^2} \\ &= \frac{n(n-1)}{n^2} \\ &= 1 - \frac{1}{n} \end{aligned} \quad (1.5)$$

where we have dropped the argument (t) everywhere for simplicity. This result indicates that if $n > 1$, then $g^{(2)}(0) \geq 0.5$. The statistical boundary for single-photon emission could therefore be defined as $g^{(2)}(0) < 0.5$, as such a value would indicate that, on average, there are less than 2 photons present at a time [10, 27]. As an average photon number state of $n < 1.5$ would seem to indicate a more realistic definition of 'statistically more single-photon events', one could instead consider $g^{(2)}(0) < 1 - 1/1.5 = 0.3$ to be a better

¹³The amplitude of these peaks should decay for large τ , as the probability of only D1 detections decreases.

single-photon threshold. For most quantum applications, however, it is expected that a single-photon source should have $g^{(2)}(0) < 0.1$ in order to be more useful than a highly-attenuated pulsed laser [28]. In Chapter 4, I provide a review of our HBT experiments with fibre-coupled nanowires, during which we attempted to reach this threshold.

1.3 Quantum Dot Charge Sensing

We have now looked at two types of quantum dots and the methods by which carriers may become trapped within them, but the question remains: how do we detect the presence of carriers within a quantum dot? In the case of the nanowire quantum dot, excitons within the dot recombine, releasing detectable photons. With lateral dots, however, individual electrons or holes remain stably trapped, with no photon emissions to indicate their presence. Fortunately, we can detect them by measuring the changes in the electrostatic environment induced by the addition or removal of carriers from the dot. There are several systems capable of making such a measurement, and here we look at two of them: the quantum point contact (QPC) and the radio-frequency QPC.

1.3.1 Quantum Point Contact as a Charge Sensor

A QPC is essentially just a one-dimensional passage through which a quantized electrical current flows. Being very sensitive to their electrostatic environment, QPCs are often used to measure changes in the carrier populations of nearby quantum dots. Figure 1.9a) shows an example of a QPC next to a lateral DQD.

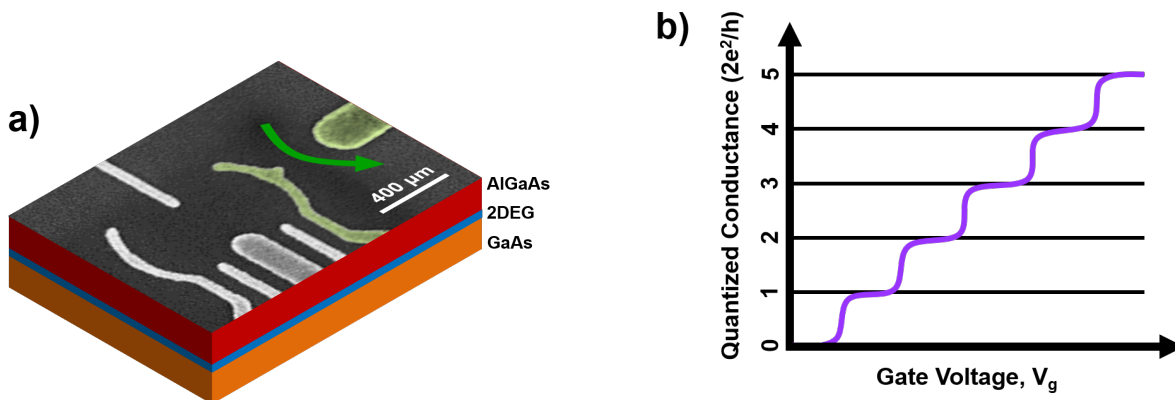


Figure 1.9: a) Double quantum dot sample with a 2DEG current (green arrow) flowing through the adjacent QPC. The green gates constrict the QPC, causing the conductance to become quantized. The 2DEG (blue) is formed by an AlGaAs-GaAs heterojunction. (2DEG is over-emphasized in this image; a real 2DEG is 2D and has no depth.) b) Example of what the quantized conductance of a QPC might look like. (This is a drawing, not actual data.)

The physics behind the operation of a QPC is fairly straightforward. Suppose, as in Fig. 1.9a), that metal gates are deposited on top of a 2DEG.¹⁴ The 2DEG regions directly

¹⁴The details of this type of device (a lateral DQD) are provided in Section 1.1. Also, note that we are only looking at electron-trapping devices here, although QPCs can be made to detect holes instead.

below the gates are depleted by the application of sufficiently negative gate voltages. While a quantum dot is formed by using several gates to cut off a few electrons from the rest of the 2DEG, a QPC is formed by two adjacent gates creating a narrow passage in the 2DEG. In Fig. 1.9a), the QPC gates are highlighted in green, and a bias applied across the QPC causes an electrical current (green arrow) to flow between these gates. By reducing the voltage applied to the gates, the QPC passage narrows further, blocking the majority of the current. The QPC is formed when the transverse confinement creates a 1D passage through which only electrons with specific energy states may propagate (much like the allowable states are quantized by a quantum well or a waveguide). As the QPC narrows, the current is reduced in sudden steps as each electron mode stops being transmitted. Thus, the conductance of a QPC is quantized according to the number of electron modes N which can pass through the barrier. The QPC conductance is given by the equation [4]:

$$G = \frac{2e^2}{h}N \quad (1.6)$$

where e is the elementary charge, h is Planck's constant, and the factor of 2 accounts for the two spin directions available in each mode. The conductance quantum $G_Q = 2e^2/h$ is therefore the conductance contributed by each transverse mode.¹⁵ The total conductance G decreases in a step-like fashion, in integer multiples of G_Q , as the QPC resistance increases. An example of the conductance as a function of QPC gate voltage, V_g , is shown in Figure 1.9c), where we can see G dropping in steps as V_g becomes more negative. The steps in this graph are smooth curves, rather than sharp steps, due to the tunnelling of electrons through the QPC [4].

If the gate voltages are adjusted such that the conductance is halfway between two of the steps,¹⁶ then the current is highly dependent on the QPC resistance; even a very small change in the electrostatic environment will result in a large change in conductance. If there is a nearby quantum dot, as in Fig. 1.9a), then the addition of a single electron to the dot will serve to decrease the local electrostatic potential and further narrow the QPC. This results in a measurable decrease in the current. Similarly, the removal of an electron will cause a sudden increase in the current. With this setup, the QPC therefore acts as a charge detector, allowing accurate detection of changes in the quantum dot's carrier population. Evidently, the QPC is more sensitive to nearby changes in the electrostatic environment than to distant ones, since the strength of the electric potential falls off as $1/r$.

¹⁵This equation for quantized conductance originates from the calculation of electrical current propagating across a barrier in one dimension, wherein one integrates the Fermi functions of states on either side of the barrier. Whereas we classically consider more current to be flowing if the electrons are moving faster, in this quantum derivation, the faster (i.e. higher-energy) states are fewer in number. Since the decrease in the density of states matches the increase in velocity as a function of energy, we are left with an equation which is independent of the exact energies or wave numbers of the electron states. For a full derivation of Equation 1.6, see [4, pp. 163-165,184-187].

¹⁶This is only possible thanks to the fact that the steps are not discontinuous, but merely exhibit steep slopes.

Importantly, the QPC is able to detect the presence of a nearby charge carrier without affecting its spin state. This is due to the fact that the charge and spin of an electron are completely decoupled from each other [2]. As we shall see in Section 1.7, this feature of the QPC will play a major role in heralding the arrival and successful absorption of polarised photons in a photon-to-spin qubit transfer system.

The only downside to this detection method is that it can be rather slow, with a bandwidth of <1 MHz [29]. It is primarily limited by the RC time constant of the QPC resistance and the wiring capacitance. A faster charge detection technique (bandwidth >1 MHz) uses RF signals in conjunction with a QPC, allowing for sub-microsecond resolution of changes in electron populations. This is the RF-QPC, which we will explore next.

1.3.2 Radio Frequency Quantum Point Contact as a Charge Sensor

A radio frequency quantum point contact (RF-QPC) is a circuit designed to measure changes in the resistance of a QPC on short time scales. Unlike the DC-QPC, which directly measures the DC current passing through a quantum point contact, the RF-QPC uses a high-frequency AC signal and a reflectometry circuit to detect changes in the QPC. As the QPC resistance affects the reflection coefficient of the circuit, changes in the QPC's electrostatic environment are measurable as changes in the amplitude of the reflected AC signal. Building on the techniques developed a decade earlier for the RF single electron transistor (RF-SET) [30], the RF-QPC has been shown to act as an extremely sensitive charge sensor, with its high bandwidth allowing for fast detection of changes in nearby electron/hole populations [29, 31]. In the SPIN Lab, we have constructed a system capable of RF-QPC charge detection, as we will see in Chapter 2.

Figure 1.10a) shows a basic schematic of the RF-QPC, which we will use to understand its operation.¹⁷ The main components of the reflectometry circuit are the source (1) and readout (2) ports for the RF signal, the circulator which directs the incident and reflected signals to their respective destinations, and the tank circuit. A tank circuit, which is any inductor-capacitor (LC) circuit, alternately stores energy in the electric field of its capacitor and the magnetic field of its inductor. A charge build-up on a capacitor drives a current, causing the inductor to generate a magnetic field. According Faraday's law of induction, the changing magnetic flux through the inductor creates a voltage across it, which then drives the current in the same direction. This again builds up a charge on the capacitor, now with opposite polarity. The sequence of events then reverses, with current flowing in the opposite direction, and repeats until the energy is fully dissipated by the circuit resistance.

Reminiscent of water splashing back and forth against the walls of a tank (hence the name 'tank' circuit), this harmonic oscillation of energy between the two components occurs

¹⁷Components such as amplifiers, attenuators, and bias tees have been omitted from the diagram for clarity.

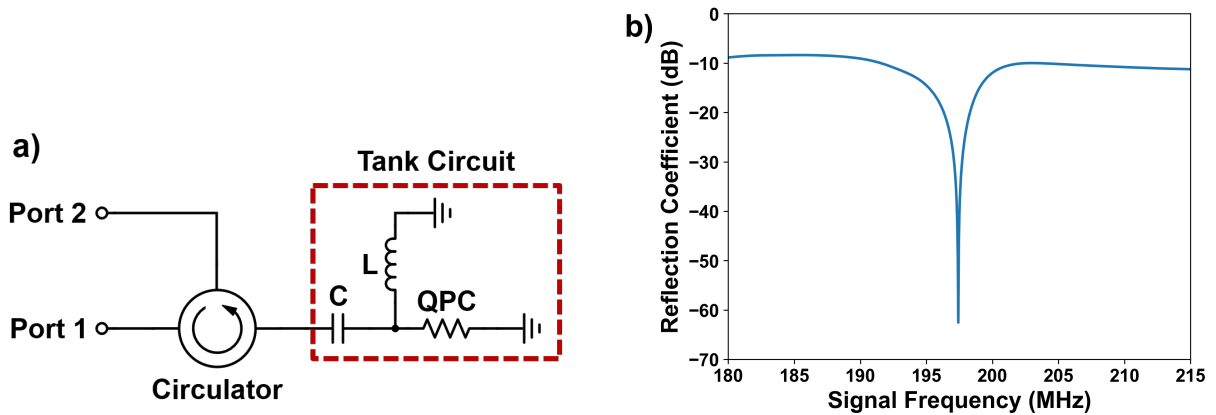


Figure 1.10: a) Basic RF-QPC setup. RF signals are sent from Port 1, through the circulator, to the tank circuit. The amplitude of the reflected signal, which is detected at Port 2, depends on the QPC resistance. b) Experimental data showing an example of the frequency response of the RF-QPC. In this case, the resonance occurs at a frequency of 197.4 MHz.

at a specific resonance frequency f_0 . For a tank circuit composed of just an inductor and capacitor, the resonance frequency is given by:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \quad (1.7)$$

$$\omega_0 = \frac{1}{\sqrt{LC}}$$

where L is the inductance of the inductor, C is the capacitance of the capacitor, and ω_0 is the angular resonance frequency. As our tank circuit incorporates a QPC, which we model as a resistor (R), it will have a very slightly modified resonance frequency. (We will construct a theoretical model of such a tank circuit in the next section.) Nevertheless, if we drive a tank circuit with an AC signal of frequency f_0 , a greater proportion of the signal will be transmitted, whereas more of the signal will be reflected by the tank circuit at other frequencies. As we will discuss in the next section, the amplitude of the reflected signal on resonance is almost entirely determined by the resistance of the QPC.

To use the RF-QPC for charge detection, we begin by sending an RF signal from Port 1 to the circulator, which directs the signal to the tank circuit. Depending on the signal frequency and QPC resistance, some portion of the signal may be transmitted to ground, while the rest is reflected and measured at Port 2. By monitoring the amplitude of the reflected signal at the resonance frequency, changes in QPC resistance can be detected. Since the resonance frequency is typically several hundred MHz, this means that the RF-QPC can be used as a fast charge detector.

Figure 1.10b) shows experimental data of the frequency response spectrum of an RF-QPC in which the resonance frequency is 197 MHz. The vertical axis is the tank circuit reflection coefficient,¹⁸ which is the ratio of the reflected signal’s amplitude to the amplitude of the input signal. When on resonance, most of the signal is transmitted through the circuit, so the reflected signal is small. The depth of the resonance is an important measure of the quality of RF-QPC operation. It is important for the resonance signal to be as deep as possible so that subtle changes in QPC resistance can be easily detected.

The quality factor is a preliminary measure of the RF-QPC’s sensitivity. The quality factor is proportional to the ratio of the energy stored in a circuit to the energy it dissipates:

$$\begin{aligned}
 Q &\equiv 2\pi \frac{\text{energy stored}}{\text{energy dissipated}} \\
 &= 2\pi f_0 \frac{\text{energy stored}}{\text{power dissipated}}
 \end{aligned}
 \tag{1.8}$$

In this equation, f_0 is the resonance frequency of the circuit. For higher quality factors (>3), we may make the following approximation:

$$Q = \frac{f_0}{\text{FWHM}}
 \tag{1.9}$$

where FWHM is the full width of the resonance when it is at half its maximum depth.¹⁹ We can therefore see that Q is a measure of the spread of frequencies contributing to the resonance; we are scaling the resonance frequency by the range of frequencies in the resonance. A higher quality factor indicates a sharper, and likely deeper, resonance. The earliest RF-QPC systems weren’t able to surpass quality factors of about 15 [32, 33]. We find from Equation 1.9 that $Q \approx 32$ for the data in Figure 1.10b), which was obtained on our system. This is decent and will allow for detection of single-electron events, but there remains room for improvement.

Although we have yet to test our RF-QPC on an actual quantum dot, we should mention here for the sake of completeness that the standard measure of an RF-QPC’s charge sensitivity is given by the following equation [29, 31, 33]:

$$\delta q = \frac{e}{\sqrt{\text{BW}} \times \text{SNR}}
 \tag{1.10}$$

¹⁸Although not shown in the simple diagram of Fig. 1.10a), we are actually attenuating our signal in line 1 before it gets to the circulator (to avoid damaging the sample), then amplifying the returning signal in line 2. The degrees of attenuation and amplification in the system do not match perfectly, hence the -10 dB vertical offset in Fig. 1.10a).

¹⁹The FWHM is measured according to the depth of the resonance when written as a function of power, not in decibels. For Q measured in dB, we use the width at the 3dB point, i.e. 3dB down from the top, which corresponds to half the maximum power.

where e is the charge of an electron, BW is the detection bandwidth, and SNR is the signal-to-noise ratio. With the SNR in the denominator, δq can be thought of as proportional to the noise-to-signal ratio, or the fraction of the signal which is noise. The charge sensitivity is effectively a measure of the RMS charge noise per $\sqrt{\text{Hz}}$ [34]. A smaller value for δq indicates a more sensitive detector, i.e. one which has less noise and is capable of measuring smaller signals. A similar measure is used for RF-SET and dispersive RF-gate measurements. Currently, most of these devices can achieve charge sensitivities on the order of 10^{-4} to $10^{-5} e/\sqrt{\text{Hz}}$ [29, 31, 33].

1.3.3 Theoretical Modelling of an LCR Tank Circuit

To understand the basic operation of a reflectometry tank circuit, we will model it using the impedances of its components. We begin by examining the reflection coefficient, which is the ratio of the reflected amplitude of an EM wave to the incident amplitude. In a circuit where an AC signal propagates along a transmission line of impedance Z_0 to a load with impedance Z_L , the current (and voltage) will be the same in both the transmission line and load if $Z_L = Z_0$. If the impedances do not match, then, by Ohm's law, the currents will not match. This implies that in the case of non-matching impedances, there must be a second signal to account for the sudden change in current, and we posit a backwards-propagating wave which is produced by the reflection of some portion of the incident wave off the load.²⁰

If the transmission line's forward-propagating signal has a voltage V_+ and the reflected signal has a voltage V_- (with associated currents I_+ and I_- respectively), then the definition of the reflection coefficient is the ratio of the two:

$$\Gamma \equiv \frac{V_-}{V_+} = \frac{I_-}{I_+} \quad (1.11)$$

Note that the second equivalence comes from Ohm's law, as both signals experience the same transmission line impedance, Z_0 .

The voltage across the load, V_L , must be the sum of the two counter-propagating voltages ($V_L = V_+ + V_-$), while the current on the load, I_L , is the difference between the currents of the two waves: $I_L = I_+ - I_-$, where I_+ and I_- are considered to have positive values.²¹ We can now follow a standard derivation of Γ in terms of impedances. It begins by writing V_+ as a function of the other variables.

²⁰Recall that the current must be the same everywhere in a series circuit, so if there is a discontinuity in the amplitude of the current, then we may hypothesize that there are actually two waves and we are observing the sum of the two. If they are counter-propagating signals, the discontinuity in the current can be explained by the proposition that one of the waves does not exist beyond that point.

²¹The total current in the transmission line must also be the difference of these two currents, and it is therefore equivalent to I_L , as expected.

$$\begin{aligned}
V_+ &= V_L - V_- \\
&= Z_L I_L - V_- \\
&= Z_L (I_+ - I_-) - V_- \\
&= Z_L \left(\frac{V_+}{Z_0} - \frac{V_-}{Z_0} \right) - V_- \\
V_+ \left(\frac{Z_L}{Z_0} - 1 \right) &= V_- \left(\frac{Z_L}{Z_0} + 1 \right) \\
\frac{V_-}{V_+} &= \frac{\frac{Z_L}{Z_0} - 1}{\frac{Z_L}{Z_0} + 1}
\end{aligned} \tag{1.12}$$

The end result is therefore:

$$\Gamma = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{1.13}$$

From this definition of the reflection coefficient, it is evident that there is no reflection when $Z_L = Z_0$, i.e. when the load impedance matches that of the transmission line. For all other load impedances, however, a portion of the signal will always be reflected. In the case of our reflectometry circuit, the load is the tank circuit. Therefore, by calculating the tank circuit impedance, we can explore the behaviour of the reflected signal as a function of tank circuit resistance, capacitance and inductance.




Figure 1.11: A simplified LCR tank circuit schematic, one which is commonly used when analysing an RF-QPC circuit [31]. We are modelling the QPC as a resistor and have included parasitic resistance and capacitance.

To calculate the resonance properties of the tank circuit, we will use the circuit diagram in Figure 1.11. This image is a reasonable representation of the real circuit, including the capacitor (C), inductor (L), QPC (which we model as a resistor, R), and parasitic capacitance to ground (C_p). An additional resistive element (r) is placed in series with the inductor to represent parasitic resistance within the circuit.²² In addition to the tank circuit, our actual reflectometry setup includes several bias tees, so our calculations here are merely to provide a basic understanding of this system.²³

For reference while we perform calculations and approximations here, let's look at some typical values for the various circuit elements and properties. The tank circuit in our system was designed to achieve a resonance frequency of $f = 200$ MHz. The capacitor C in an RF-QPC tank circuit will usually have a capacitance of a few pF, while L may be 0.1 - 1 μ H. Typical values of the QPC resistance R range from a few tens of kilo-ohms to infinite resistance (i.e. QPC is pinched off), although we generally expect impedance matching for values of about 10 - 150 k Ω . The parasitic capacitance in a system like this might be on the order of a few pF, while the parasitic resistance r could be somewhere between 1 and 10 Ω [31].

With these values in mind, we can start to mathematically model the system. Using the simple schematic of Figure 1.11, we can easily write out the tank circuit impedance:

$$\begin{aligned}
Z &= \frac{-i}{\omega C} + \left(\frac{1}{R} + i\omega C_p + \frac{1}{i\omega L + r} \right)^{-1} \\
&= \frac{-i}{\omega C} + \frac{R}{1 + iR \left(\omega C_p - \frac{1}{\omega L + r} \right)} \frac{1 - iR \left(\omega C_p - \frac{1}{\omega L + r} \right)}{1 - iR \left(\omega C_p - \frac{1}{\omega L + r} \right)} \\
&= \frac{R}{1 + R^2 \left(\omega C_p - \frac{1}{\omega L + r} \right)^2} - i \left[\frac{1}{\omega C} + \frac{R^2 \left(\omega C_p - \frac{1}{\omega L + r} \right)}{1 + R^2 \left(\omega C_p - \frac{1}{\omega L + r} \right)^2} \right]
\end{aligned} \tag{1.14}$$

Electrical resonance is defined to be the frequency at which the circuit impedance is real [35], so the impedance equation above has been separated into real and imaginary components. The reactance (i.e. the imaginary component) therefore determines the resonance frequency, while the resistance (i.e. the real component) determines the amplitude of the reflected signal on resonance. Assuming the values of all circuit components are known, we can calculate the resonance frequency by setting the imaginary component of Equation 1.14 to zero. Solving the exact equation for the resonance frequency is overly complicated for this brief investigation of the RF-QPC, so to simplify things, we will assume the QPC resistance R is sufficiently large that we may do the following:

²²As the purpose of this resistor is merely to represent other losses in the system, its exact placement within the circuit is largely irrelevant. Changing its location merely requires it to have a different resistance in order to match the real system, but it should not affect the main results of the simulation.

²³For further details on our actual reflectometry circuit, see Section 2.3.

$$\begin{aligned}
0 &= -i \left[\frac{1}{\omega C} + \frac{R^2 \left(\omega C_p - \frac{1}{\omega L+r} \right)}{1 + R^2 \left(\omega C_p - \frac{1}{\omega L+r} \right)^2} \right] \\
&\approx -i \left[\frac{1}{\omega C} + \frac{R^2 \left(\omega C_p - \frac{1}{\omega L+r} \right)}{R^2 \left(\omega C_p - \frac{1}{\omega L+r} \right)^2} \right] \\
&= \frac{1}{\omega C} + \frac{1}{\omega C_p - \frac{1}{\omega L+r}} \\
&= \omega C_p - \frac{1}{\omega L+r} + \omega C \\
&= \omega (C_p + C) (\omega L + r) - 1 \\
&= L(C_p + C)\omega^2 + r(C_p + C)\omega - 1
\end{aligned} \tag{1.15}$$

The assumption made in the second line is quite reasonable, even assuming some extreme (though still realistic) values for the variables. This is a quadratic equation, so if we let $x = C_p + C$, the resonance frequency is:

$$\begin{aligned}
\omega &= \frac{-rx \pm \sqrt{r^2x^2 + 4Lx}}{2Lx} \\
&\approx \frac{-rx \pm 2\sqrt{Lx}}{2Lx} \\
&= \frac{-r}{2L} \pm \frac{1}{\sqrt{Lx}} \quad \rightarrow \text{negative solution is inadmissible} \\
&\rightarrow \frac{-r}{2L} + \frac{1}{\sqrt{Lx}} \\
&= \frac{1}{\sqrt{L(C_p + C)}} - \frac{r}{2L}
\end{aligned} \tag{1.16}$$

The approximation in the second line arises from the consideration of typical values for the circuit components. This result shows that the resonance frequency depends almost entirely upon the values of the inductor, the capacitor, and the parasitic capacitance. Depending on the value of the parasitic resistance r , the resonance frequency may be shifted by a few MHz.²⁴ Based upon the inherent parasitic capacitance and resistance of the system, the inductor and capacitor may be chosen so as to achieve the desired resonance frequency. Note that if we set C_p and r to zero, as in an ideal circuit, we retrieve the ideal tank circuit resonance frequency of Equation 1.7.²⁵

²⁴A downshift of up to about 50MHz is possible if r is sufficiently large and L sufficiently low.

²⁵The approximation we made in the second line removed the effect of the QPC from the equation, but in reality it does have a slight influence on the resonance frequency.

With the resonance determined by the circuit reactance, Equation 1.11 indicates that the real component of the impedance must dictate the amplitude of the reflected wave on resonance. As the resonance frequency, capacitance, and inductance are predetermined during construction of the circuit, and r and C_p are set by the system itself, the only free variable to control the tank circuit impedance is the QPC resistance R . On resonance (i.e. with the reactance equal to zero), Equation 1.14 gives a tank circuit impedance of:

$$Z = \frac{R}{1 + R^2 \left(\omega C_p - \frac{1}{\omega L + r} \right)^2} \approx \frac{1}{R \left(\omega C_p - \frac{1}{\omega L + r} \right)^2} \quad (1.17)$$

Equation 1.11 tells us that the amplitude of the reflected wave is only zero when the system reaches an impedance-matched state, where $Z = Z_0$. Setting the tank circuit impedance to 50Ω , we can therefore calculate the QPC resistance which yields 100% signal transmission at the resonance frequency:

$$R = \frac{1}{50 \left(\omega C_p - \frac{1}{\omega L + r} \right)^2} \quad (1.18)$$

Assuming $C_p = 2 \text{ pF}$, $r = 1 \Omega$, $C = 1.5 \text{ pF}$, $L = 180 \text{ nH}$, and a resonance frequency of 200 MHz (these are the values used in our system), this equation suggests the impedance-matching condition is met when the QPC resistance is $5.6 \text{ k}\Omega$.²⁶

Putting all of these ideas together, Figure 1.12 shows the frequency response of the system as a function of capacitance, inductance and QPC resistance.²⁷ These simulations, which were produced from Equations 1.14 and 1.16, show that the resonance frequency decreases as the inductance and capacitance increase. The QPC resistance affects only the amplitude of the reflected signal, which is also influenced by inductance and capacitance. The QPC resistance would have a slight effect on the resonance, but for the assumption we made in Equation 1.15. Impedance matching can be seen to occur in Fig. 1.12c), with settings of $L = 200 \text{ nH}$, $C = 0.4 \text{ pF}$, and $R = 50 \text{ k}\Omega$.

This type of simulation can be used to help determine which inductor and capacitor should be used to construct the circuit so as to achieve the desired resonance frequency. As we will discuss in Section 2.3.2, however, we employed circuit modelling software to simulate our system, as the impedance equation becomes far more complicated when one includes the three bias tees in our circuit.

²⁶Impedance matching generally occurs at around $100 \text{ k}\Omega$ in our actual system, as its bias tees, which weren't considered here, alter its behaviour.

²⁷A parasitic capacitance of 2 pF and a parasitic resistance of 1Ω were used to make these plots.

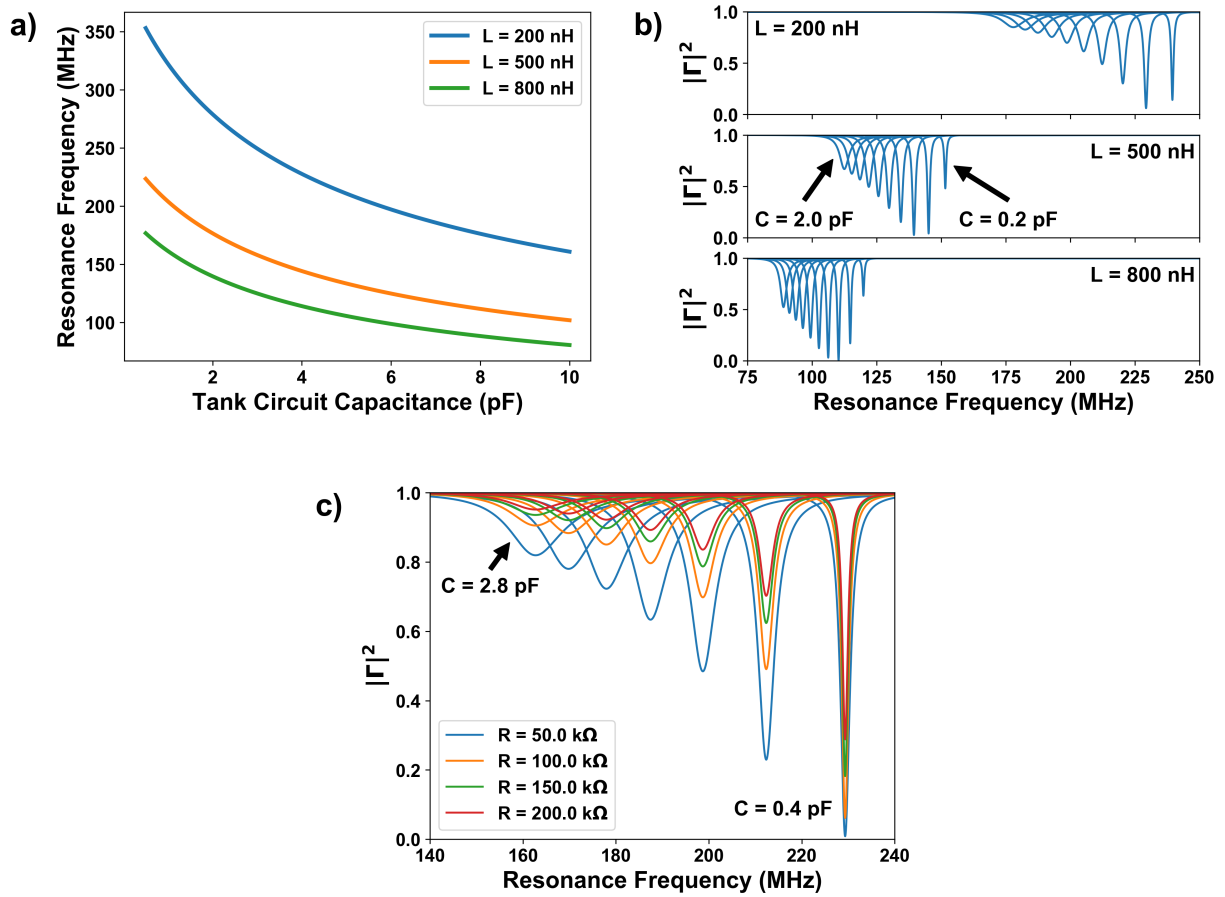


Figure 1.12: Simulating the properties of a tank circuit, using constant parasitic values of $C_p = 2$ pF and $r = 1$ Ω . a) Resonance frequency as a function of tank circuit capacitance and inductance, with constant $R = 100$ k Ω . b) Frequency response for various capacitor and inductor values, with $R = 100$ k Ω . c) Frequency response as a function of tank circuit capacitance and QPC resistance, with $L = 200$ nH.

1.4 Light Polarisation

The best option for transmitting quantum information over great distances is to encode it in light. The speed and easy manipulability of photons, not to mention the simplicity of their long-distance transmission via optical fibre, has made them an important focus of quantum communication research. Photons also offer several options for encoding quantum information, such as polarisation encoding, energy/wavelength encoding, and time-bin encoding.

In our attempts to develop a photon-to-spin quantum repeater, we have chosen to use photon polarisation states as our qubits because they offer the simplest option for the direct transfer of quantum information to carrier spin states. In the context of long-distance transmission by fibre, however, we run into the issue that standard optical fibres generally do not preserve the polarisation state of the light they transmit. As we will discuss in Section 1.6 and Chapter 5, this unfortunate fibre optic property demands that some form of polarisation compensation technique be employed in order to accurately transmit the desired states. Our own system accounts for the fibre's effects through the use of waveplates, which are special crystals designed to alter polarisation in a fully controllable manner.

In this section, we will review the basics of polarisation, including the differences between its classical and quantum descriptions. This is followed by a brief overview of waveplates and their use in controlling the polarisation state of light.

1.4.1 Classical vs. Quantum Polarisation

The connection between the polarisation of classical electromagnetic (EM) waves and the spin states of photons in the quantum regime is an important aspect of our work on the photon-to-spin project. To arrive at the concept of photon polarisation, we will review classical polarisation and photon spin states, and we will then explore the connection between them.

Classical Polarisation

Classically, polarisation is defined by the orientation of the electric field in an electromagnetic wave. Figure 1.13 shows a light wave with the E-field drawn in red. As the electric field is aligned along the y-direction, this light is defined to be vertically-polarised. In instances where the electric field oscillates in a plane formed by the z-axis and another line (e.g. the y-axis in Fig. 1.13), the light is said to be linearly-polarised. The four standard linear polarisations are horizontal (x-axis), vertical (y-axis), diagonal ($+45^\circ$ from x-axis), and anti-diagonal (-45° from x-axis). The corresponding symbols for these four states are H, V, D, and A respectively.

Diagonally-polarised light has an E-field with x and y components which are equal in amplitude and oscillating in phase with each other. In some cases, however, the x and y




Figure 1.13: A classical electromagnetic wave (vertically-polarised).

components of the electric field may be oscillating out of phase. This will result in the E-field direction evolving over time, and it will appear to orbit the axis of propagation (i.e. the z-axis). If the two field components are 90° out of phase and equal in amplitude, the E-field vector will rotate at a constant rate over time, tracing out a circle. This is called circular polarisation. If the direction of E-field oscillation obeys the right-hand rule, the polarisation state is called right circular (R), while it is called left circular (L) if it rotates in the opposite direction. Together, the horizontal, vertical, diagonal, anti-diagonal, right circular and left circular states are the six standard polarisations which are typically used in experiments. All other non-linear states are called elliptical polarisations, as their E-fields trace out an ellipse as the wave propagates.

Quantum Polarisation

In quantum mechanics, the quantized particle of light is the photon, which is a massless, spin-1 particle. Quantum electrodynamics allows us to form a link between the spin states of these light quanta and the polarisation of a classical EM wave [36, 37]. We will review this connection through a quick derivation of the allowable states of the electromagnetic four-potential.²⁸

The electrodynamics Lagrangian, i.e. the Lagrangian for a massless, spin-1 particle, is given by [36, 37]:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}^2 \quad \text{where} \quad F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \quad (1.19)$$

Here, $F_{\mu\nu}$ is the electromagnetic tensor, which describes an EM field in spacetime, and $A^\mu = (\phi/c, \mathbf{A})$ is the electromagnetic four-potential. The Lagrangian is gauge-invariant for

²⁸I do not go into great detail here, but see [36, pp. 118-120] and [37, pp. 223-227] for more information on this derivation. The quantum electrodynamics equations in this section are largely compiled from these sources, although these are well-known results.

the transformation $A_\mu \rightarrow A_\mu + \partial_\mu \lambda$. According to the Euler-Lagrange equations, the four equations of motion for this Lagrangian are therefore [36]:

$$\square A_\mu - \partial_\mu(\partial_\nu A_\nu) = 0 \quad (1.20)$$

Under the constraints of the source-free Maxwell equations ($\partial_\mu \partial^\mu A_\nu = 0$) and the Lorentz condition ($\partial^\mu A_\mu = 0$), the solution to Equation 1.20 is [37]:

$$A_\mu(x) = a_0 \varepsilon_\mu(p) e^{-ip \cdot x} \quad (1.21)$$

where a_0 is a constant, p is the momentum, x is the position, and $\varepsilon_\nu(p)$ is the photon spin vector. In the Coulomb gauge, the constraints for this solution become $\varepsilon^\mu = (0, \vec{\varepsilon}')$, with $\varepsilon^\mu \varepsilon_\mu = 1$ and $p^\mu \varepsilon_\mu = 0$. For a massless particle travelling in the z-direction, the four-momentum is given by $p^\mu = (E, 0, 0, E)$, so $p^\mu p_\mu = 0$. With these requirements eliminating two degrees of freedom, we are left with two solutions to the equations of motion [36, 37]:

$$\varepsilon^{\mu(x)} = (0, 1, 0, 0) \quad \text{and} \quad \varepsilon^{\mu(y)} = (0, 0, 1, 0) \quad (1.22)$$

Equations 1.21 and 1.22 therefore dictate that the only possible values for the electromagnetic four-vector are [37]:

$$A_{\mu(x)} = a_0 e^{-ip \cdot x} \hat{x} \quad \text{and} \quad A_{\mu(y)} = a_0 e^{-ip \cdot x} \hat{y} \quad (1.23)$$

which are vectors in the x- and y-directions.

To link these results to the classical regime, recall that the classical electric field can be written as:

$$\mathbf{E} = -\nabla \phi - \frac{\partial \mathbf{A}}{\partial t} \quad (1.24)$$

Taking this into account, the solutions in Equation 1.23 imply that the E-field of the photon wave is oriented along the x- and y-axes, corresponding to the exact definitions of horizontal ($\varepsilon^{\mu(x)}$) and vertical ($\varepsilon^{\mu(y)}$) polarisations. The normalised photon spin vector, $\varepsilon_\mu(p)$, is therefore also the photon polarisation vector, as it dictates the orientation of the electric field. A mathematical treatment of polarisation will be given in Section 1.5, where we shall see that these two solutions are in fact the exact equations for horizontal and vertical polarisations written as Jones vectors.

An alternative basis for A_μ would be:

$$\varepsilon^{\mu(L)} = \frac{1}{\sqrt{2}}(0, 1, i, 0) \quad \text{and} \quad \varepsilon^{\mu(R)} = \frac{1}{\sqrt{2}}(0, 1, -i, 0) \quad (1.25)$$

These solutions, which correspond to the right and left spin states (or the helicity) of the spin-1 photon, are also the right and left circular polarisation states. All other polarisations are superpositions of these two states. The photon is limited to two spin states by its masslessness, which demands it travel at the speed of light. An observer cannot move faster than a photon, which would cause a reversal of its spin direction with respect to its direction of propagation. We therefore always observe the spin direction to be fixed either parallel or anti-parallel to the propagation direction.

Importantly, a beam of light is composed of photons. A polarised beam has 100% of its photons identically-polarised, while an unpolarised beam has all of its photons randomly-polarised, with no statistically significant portion of them all polarised in a specific direction. In a partially-polarised beam, some portion of its constituent photons are polarised in random directions, but the sum of their polarisations gives an overall polarisation for the beam

Although photons in a light beam may interact with their environment, they do not interact with each other. Consequently, if the polarisation of the beam is changed, by a waveplate for example, it means that the polarisation of each individual photon has been changed in the same way (assuming no depolarisation of the beam). We may therefore use some quantum mechanical methods to analyse and describe the polarisation state of the beam, such as when we consider the ‘state fidelity’ of a beam.

1.4.2 Waveplates and Polarisation Control

The polarisation of light is often changed when it passes through matter, as the photons interact with other particles. The result is often a randomized polarisation state when the light reaches the other side of the substance. It is possible, however, to engineer materials which can controllably and coherently alter the polarisation of light. These devices are typically birefringent crystal plates called ‘waveplates’. Birefringence is a property of certain crystalline solids which have two different indices of refraction perpendicular to both the direction of light propagation and each other[38]. When light passes through such a material, the x- and y-components of the electric field travel at different speeds through the crystal, so the two crystal axes are called the ‘fast’ and ‘slow’ axes based on the relative phase velocities with which the electric field components travel along them. While the fast axis has little effect on the speed of the light, the slow axis reduces the speed of light, which is why these materials are often called ‘retarders’. The difference in the rates at which the E-field’s x- and y-components propagate causes the phase and relative amplitudes between them to be altered, which of course changes the direction of polarisation.

As the effect on the polarisation depends on how the electric field components are oriented with respect to the fast and slow-axes, waveplates can have different effects on different types of incident polarisation. For example, linear polarisation aligned with either the fast or slow axis will pass through the waveplate unchanged, but linear polarisation which is not aligned with either axis will experience a phase between its x- and y-components.

Rotating a waveplate will also change its effect on the light, depending on the alignment of the incident polarisation with the new orientation of the waveplate axes.

Retardance, written as a fraction of wavelength, is the relative phase difference between the x- and y-components of photons travelling along the two axis modes of a birefringent material [38]. Waveplates are typically engineered to produce very specific retardances in the light which passes through them: 0.5 for the half-wave plate (HWP or $\lambda/2$) and 0.25 for the quarter-wave plate (QWP or $\lambda/4$). The retardance experienced by light travelling through a birefringent material depends on the angle of incidence and the wavelength, although some materials are designed to have a large optical bandwidth. Light with a wavelength outside the appropriate range may therefore experience an imperfect waveplate retardance, leading to slightly different changes in the polarisation than would be produced by an ideal waveplate.




Figure 1.14: Waveplates acting on polarised light. The black line on each waveplate marks the orientation of its fast axis. a) The HWP ($\lambda/2$) acts to rotate linear (and elliptical) polarisation. b) The QWP ($\lambda/4$) induces a phase shift between the x- and y-components of the E-field. The effect is a transformation of linear polarisation into circular, and vice versa.

The result of a 0.5 retardance is that half-wave plates rotate linearly-polarised light, keeping the x- and y-components of the E-field in phase.²⁹ The relative angle between the fast axis of the HWP and the polarisation orientation will determine the degree of polarisation rotation. For example, in the case depicted in Figure 1.14a), vertically-polarised light is rotated as it passes through a rotated HWP, resulting in the diagonal state after the waveplate. A second rotated HWP produces horizontal polarisation. As a rule, to rotate linear polarisation by an angle θ , the light must pass through a HWP which is offset by $\theta/2$ from the orientation of the incident polarisation. (So in the example given in the figure,

²⁹A HWP can also rotate elliptically-polarised light, but it will have no effect on circularly-polarised light.

the first waveplate has a fast axis rotated by 22° from vertical, while the second is rotated by 45° from the vertical.)

A quarter-wave plate induces a phase shift between the x and y components of the E-field, thereby transforming linearly-polarised light into circularly-polarised light and vice versa. Elliptical states can also be generated using a QWP, depending on the waveplate orientation. Figure 1.14b) demonstrates linear light being transformed into elliptical and circular states after passing through two QWPs at different angles. To generate right circular polarisation from linear, the waveplate should be rotated by $+45^\circ$ from the incident polarisation direction, while a QWP angle of -45° will produce left circular polarisation.

The combination of a rotating HWP followed by a rotating QWP therefore allows linear light to be transformed into any polarisation. In our attempts to compensate for the polarisation-altering effects of an optical fibre, we use a HWP+QWP setup to transform the polarisation before it enters the fibre, allowing us to control the output state.

1.5 Mathematical Treatment of Light Polarisation

Mathematical modelling is a major component of our procedure for dealing with the effects of an optical fibre on light polarisation. It not only allows us to characterise the behaviour of our wavelength-dependent waveplates, but it also enables us to exactly determine the influence of the fibre on the polarisation. With this knowledge of the fibre, we can then calculate the waveplate orientations which ensure the desired polarisation states can be transmitted.

Mathematical descriptions of polarisation states and of the optical elements affecting polarisation are given by Mueller calculus and Jones calculus. While Mueller calculus is a more general representation than Jones calculus, allowing one to account for partial polarisation, both are important tools when modelling polarisation states. In Sections 1.5.1 and 1.5.2, we will review both of these methods of describing polarisation, as I will use both when explaining the results of our polarisation experiments in Chapter 5.³⁰

Through the mathematical manipulation of matrices representing retarding materials, we can derive a valid and practical equation for describing an optical fibre's impacts on polarisation. This derivation is presented in Section 1.5.3. A quantitative measure of our success in transmitting a polarised state through a fibre is given by the fidelity of the delivered state with respect to the ideal, so we will review the mathematics of this measurement in Section 1.5.4.

1.5.1 Mueller Calculus

With Mueller calculus, Stokes vectors are used to represent the state of either coherent or incoherent light which is polarised, partially polarised, or unpolarised. The Stokes vector is written as:

$$S_\mu = \begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix} \quad (1.26)$$

I , Q , U , and V are the traditional labels for the four Stokes parameters, but we will attempt to stick to the more modern labels S_i .

The first Stokes parameter, S_0 , is the intensity of the light, while the other three parameters describe the polarisation state of the light. S_1 describes the proportion of the light which is horizontally- or vertically-polarised. It can be understood for the moment as the

³⁰When describing states or retarders in the text, I will generally stick to Jones calculus, as its mathematical representations are simpler. Mueller calculus, however, is necessary for describing certain depolarising effects.

amount of light in a polarised beam which can pass through a horizontally-oriented linear polariser. (We'll look at a more accurate description shortly.) The S_2 and S_3 parameters represent the degree of polarisation in the diagonal/anti-diagonal and right/left circular bases respectively.

To define the i^{th} element of a Stokes vector, we imagine measuring the intensity I_i of light passing through a particular filter. This filter permits the transmission of *half* the incident light and deflects the rest. For S_0 , all polarisations are transmitted, so S_0 is defined as twice the measured intensity of light passing through the filter,³¹ i.e. $S_0 = 2I_0$. For Stokes parameters 1, 2, and 3 respectively, their associated filters are a horizontally-oriented linear polariser, a linear polariser oriented at an angle of $+45^\circ$, and a right-circular polariser (one which does not permit transmission of linear or left-polarised light). The Stokes parameters are then defined to be [38]:

$$\begin{aligned} S_0 &= 2I_0 \\ S_1 &= 2I_1 - 2I_0 \\ S_2 &= 2I_2 - 2I_0 \\ S_3 &= 2I_3 - 2I_0 \end{aligned} \tag{1.27}$$

From this definition, we can see that when the light is completely horizontally-polarised, $S_1 = 1$ (assuming a normalized light intensity), while $S_1 = -1$ if the light is completely vertically-polarised. If there is no polarisation in the H/V basis, then $S_1 = 0$. The final two Stokes parameters behave similarly in their respective bases: $S_2 = 1$ (-1) for state D (A), and $S_3 = 1$ (-1) for state R (L). The various possible elliptical states are represented by combinations of these three parameters. Table 1.1 lists the Stokes vectors for the most commonly-used polarisation states.

When the light is fully polarised, the final three Stokes parameters are related to the intensity by the equation:

$$S_0^2 = S_1^2 + S_2^2 + S_3^2 \tag{1.28}$$

It may be, however, that the light is partially depolarised (i.e. some of the photons are randomly-polarised, while the remainder are all polarised in the same direction). In this case, the total intensity will be greater than the sum of the squares of the final three Stokes parameters:

$$S_0^2 \geq S_1^2 + S_2^2 + S_3^2 \tag{1.29}$$

³¹Since the filter only lets half the light through, this definition of S_0 is equivalent to the description above, where we stated it to be the total intensity of the light.

State	Stokes Vector	Jones Vector
$ H\rangle$	$\begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$
$ V\rangle$	$\begin{pmatrix} 1 \\ -1 \\ 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$
$ D\rangle$	$\begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \end{pmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$
$ A\rangle$	$\begin{pmatrix} 1 \\ 0 \\ -1 \\ 0 \end{pmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$
$ R\rangle$	$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix}$
$ L\rangle$	$\begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}$
Unpolarised	$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$	N/A

Table 1.1: The Jones and Stokes representations of the standard polarisation states

So S_1 , S_2 , and S_3 only describe that portion of the light which is polarised. It is therefore useful to consider what fraction of total light intensity is polarised. From Equation 1.29, it is evident that the degree of polarisation is:

$$p = \frac{\sqrt{S_1^2 + S_2^2 + S_3^2}}{S_0} \quad (1.30)$$

Therefore, to account for partial polarisation in the definition given by Equation 1.27, we must multiply S_1 , S_2 , and S_3 by p .

On the other hand, an alternative, but equivalent, definition of the Stokes vector is [39]:

$$S_\mu = \begin{pmatrix} S_0 \\ S_1 \\ S_2 \\ S_3 \end{pmatrix} = \begin{pmatrix} I \\ Ip \cos(2\psi) \cos(2\chi) \\ Ip \sin(2\psi) \cos(2\chi) \\ Ip \sin(2\chi) \end{pmatrix} \quad (1.31)$$

These equations for the parameters S_1 , S_2 and S_3 are reminiscent of the translation equations between Cartesian and spherical coordinates, with 2χ as the polar angle and 2ψ as the azimuthal angle.³² The sphere formed by these coordinates is called the Poincaré sphere, and it represents the space of all possible polarisation states.³³ The Stokes vector represents a polarisation state on this sphere. Figure 1.15a) shows the Poincaré sphere drawn on the axes of S_1 , S_2 and S_3 . The first Stokes parameter, $S_0 = I$, dictates the radius of the sphere, so Ip represents the length of the state vector within it [39, 40]. We will assume throughout this document, unless otherwise stated, that the intensity is normalized, so the Poincaré sphere is a unit sphere. In an equivalent representation, Figure 1.15b) shows the Poincaré sphere with the axes labelled according to the six standard polarisations.

While we now have a mathematical description of the polarisation state of light, we need some method of describing changes in the state. Objects which affect polarisation, such as waveplates and polarisers, are represented using 4×4 matrices called Mueller matrices. For example, the Mueller matrix describing a horizontally-oriented linear polariser is:

$$P_H = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (1.32)$$

and a generic linear retarding material with a horizontally-oriented fast axis is represented by:

$$D_H(\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\phi) & \sin(\phi) \\ 0 & 0 & -\sin(\phi) & \cos(\phi) \end{pmatrix} \quad (1.33)$$

³²While not covered here, a full derivation of these equations for the Stokes parameters, using the polarisation ellipse, can be found in *Optics: Principles and Applications* by K. K. Sharma [39].

³³The Poincaré sphere is analogous to the Bloch sphere, which represents the space of possible states for a two-level quantum system.

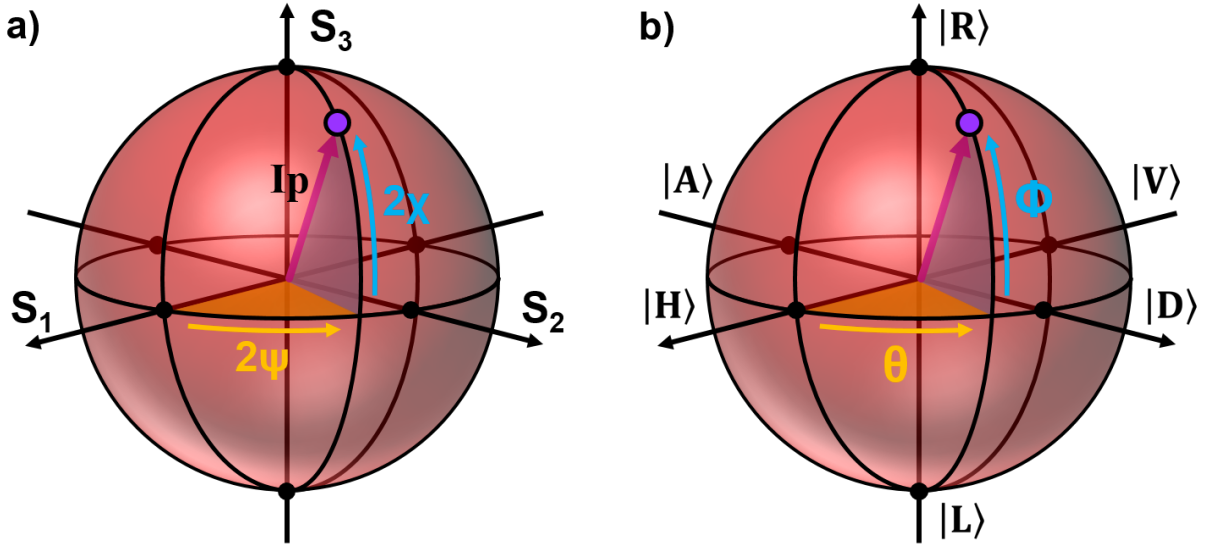


Figure 1.15: a) Poincaré sphere and polarisation vector depicted in terms of Stokes parameters. b) Poincaré sphere and polarisation vector depicted in terms of polarisation states.

where ϕ is the phase shift induced by the waveplate. The Mueller matrices describing the most common optical elements are given in Table 1.2. (The table also includes the rotation matrix, which is described in the following section to take advantage of the simplicity of Jones matrices.)

To calculate the Stokes vector of the light after it has passed through an optical element, we multiply the appropriate Muller matrix by the input Stokes vector. For example, for an input polarisation S_i^μ , the output polarisation S_f^μ after passing through a linear polariser P is: $S_f^\mu = P \cdot S_i^\mu$. The matrix representing each successive element in the optical is multiplied from the left. Mueller calculus thereby allows one to determine the polarisation of any light propagating through a system.

1.5.2 Jones Calculus

When the light source is coherent, polarisation states can be modelled using the much simpler Jones calculus. Jones vectors have only two elements, representing the components of the electric field in the x- and y-directions. For coherent, polarised light propagating along the z-axis, the Jones vector is:

$$\mathbf{J} = \begin{pmatrix} E_x \\ E_y \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\phi_x} \\ E_{0y}e^{i\phi_y} \end{pmatrix} = e^{i\phi_x} \begin{pmatrix} E_{0x} \\ E_{0y}e^{i\phi} \end{pmatrix} \quad (1.34)$$

In this equation, E_{0j} represents the wave amplitude in the j^{th} direction, while the value

Optical Element	Mueller Calculus	Jones Calculus
Linear Polariser	$\frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$
Half-Wave Plate	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$	$\begin{pmatrix} e^{-i\pi/2} & 0 \\ 0 & e^{i\pi/2} \end{pmatrix}$
Quarter-Wave Plate	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix}$	$\begin{pmatrix} e^{-i\pi/4} & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$
Rotation Matrix, $R(\theta)$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(2\theta) & \sin(2\theta) & 0 \\ 0 & -\sin(2\theta) & \cos(2\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}$

Table 1.2: Representations of the most common optical elements using Jones and Mueller calculus. All elements are oriented horizontally, but in accordance with Equation 1.40, each can be rotated using its respective rotation matrix.

$\phi = \phi_y - \phi_x$ is the relative phase between the x- and y-components of the electric field.³⁴ Light which is horizontally-polarised along the x-axis has the state $|H\rangle = (1, 0)^T$, while vertical polarisation is $|V\rangle = (0, 1)^T$. The Jones vector representations of the standard states are displayed in Table 1.1.

The Jones vector also gives the total light intensity:

$$I = |\mathbf{J}|^2 = |E_x|^2 + |E_y|^2 \quad (1.35)$$

A Jones vector is always considered to be normalized, $I = 1$, except after passage through a polariser (which we will discuss below). This normalisation requirement allows the state to be placed on the Poincaré sphere if it is rewritten as:

³⁴The global phase $e^{i\phi_x}$ may be discarded because it does not affect the intensity, which is the only measurable property.

$$\begin{aligned}
\mathbf{J} &= E_{0x} |H\rangle + E_{0y} e^{i\phi} |V\rangle \\
&= \cos(\theta/2) |H\rangle + \sin(\theta/2) e^{i\phi} |V\rangle
\end{aligned}
\tag{1.36}$$

where θ dictates the degree of polarisation in H and V, and the phase shift ϕ gives the coordinates in the D/A-R/L plane.³⁵ Figure 1.15b) depicts a Jones vector on the Poincaré sphere, which is shown in its standard representation, with R at the top and L at the bottom.

Unlike Mueller calculus, Jones calculus maintains information about the phase between the two field components, but it can only deal with a fully coherent beam of light. Incoherent or partially-polarised light must be described by Stokes vectors. Jones vectors, which are pure states, must therefore always extend to the full radius of the Poincaré sphere, hence the normalised intensity.

As mentioned above, however, the square of a Jones vector (i.e. the light intensity) is not necessarily equal to 1 if the light has passed through a projector, such as a polariser. A polariser will filter out some component of the electric field, yielding a reduced intensity. For example, a linear polariser oriented along the horizontal axis is modelled by the following 2×2 Jones matrix:

$$P_H = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}
\tag{1.37}$$

As with Mueller calculus, the output state of the light passing through this polariser is obtained by multiplying its matrix by the input state:

$$\begin{aligned}
\mathbf{J}_f &= P_H \cdot \mathbf{J}_i \\
&= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} E_{0x} \\ E_{0y} e^{i\phi} \end{pmatrix} \\
&= \begin{pmatrix} E_{0x} \\ 0 \end{pmatrix}
\end{aligned}
\tag{1.38}$$

This shows mathematically that the polariser matrix extracts the component of the electric field which is aligned with it, as we would expect.

To represent a rotated polariser, or indeed any rotated optical element, we employ rotation matrices. The standard 2×2 rotation matrix is:

$$R(\theta) = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}
\tag{1.39}$$

³⁵This is the same as the standard equation for a state on the Bloch sphere.

To rotate our entire coordinate frame, we multiply our input vector by the rotation matrix: $\mathbf{J}'_f = R(\theta) \cdot \mathbf{J}_i$. In this rotated frame, the representation of our rotated optical element (we'll call it A') has the standard form from Table 1.2. We can therefore easily apply it to our input vector: $\mathbf{J}'_f = A' \cdot R(\theta) \cdot \mathbf{J}_i$.

We typically wish to perform all of our calculations in the lab frame, however, so we must convert this result back to our original reference frame. It so happens that the inverse of the rotation matrix, $R^{-1}(\theta)$, is $R(-\theta)$. Applying the inverse rotation matrix to the right-hand side, our output is again in the lab frame: $\mathbf{J}_f = R(-\theta) \cdot A' \cdot R(\theta) \cdot \mathbf{J}_i$. From this, we can clearly see that the representation of an optical element A in the lab frame, after it has been rotated by an angle θ from the horizontal, is given by:

$$\begin{aligned} A &= R(-\theta) \cdot A' \cdot R(\theta) \\ &= \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} A' \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \end{aligned} \quad (1.40)$$

This rotation matrix can be applied to any Jones representation of a horizontally-oriented optical element. Mueller calculus also has a rotation matrix, although it is a 4×4 matrix to match the dimensionality of other Mueller matrices. Both of these representations are included in Table 1.2.

Finally, in addition to the polariser, the other major optical element is the waveplate, which changes the relative phase between the two components of the E-field. The mathematical representation of a waveplate D with its fast axis oriented horizontally is:

$$D(\phi) = \begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix} \quad (1.41)$$

This waveplate induces a phase change of ϕ between the two components of the electric field. The retardance of the waveplate is $\phi/2\pi$ in radians, or $\phi/360^\circ$ in degrees. As discussed in Section 1.4.2, a half-wave plate, which induces a phase change of π between the two components, therefore has a retardance of 0.5, while a quarter-wave plate, which is designed with $\phi = \pi/2$, has a retardance of 0.25. As with the polariser, waveplates can be rotated by application of the rotation matrix, $R(\theta)$, in accordance with Equation 1.40. Table 1.2 provides the Jones representations of the three most common optical elements.

1.5.3 Mathematical Description of an Optical Fibre

The development of a mathematical description of an optical fibre's effect on polarisation is a relatively recent achievement, having been demonstrated for the first time in 2015 [41]. In their paper on the application of Mueller calculus to tissue analysis using a fibre-based endoscope, Manhas et al. reveal their technique for characterising a fibre: the fibre

is considered to be composed of a long series of n retarding plates. Each plate M_i has a unique retardance ϕ_i , and a fast axis rotated by an angle θ_i from the horizontal. The fibre F may then be described by:

$$\begin{aligned}
F &= M_{n-1}M_{n-2} \dots M_2M_1M_0 \\
&= R(-\theta_{n-1})D(\phi_{n-1})R(\theta_{n-1})R(-\theta_{n-2})D(\phi_{n-2})R(\theta_{n-2}) \dots \\
&= \prod_{i=1}^n R(-\theta_{n-i})D(\phi_{n-i})R(\theta_{n-i})
\end{aligned} \tag{1.42}$$

where $R(\theta_i)$ is a rotation matrix, and $D(\phi_i)$ represents the i^{th} waveplate.³⁶ When this equation is written using Mueller matrices, it can be shown that the application of the pull-through lemma to this equation reduces it to the following simplified result [41]:

$$F = R(-\theta_a)D(\phi)R(\theta_b) \tag{1.43}$$

Calculating the polarisation of light at the fibre output is now merely a matter of solving for the three fibre parameters: θ_a , θ_b , and ϕ . As we shall see in Chapter 5, this is achieved by sending known states into the fibre and comparing them to the output states. Note that although this equation is reminiscent of the usual rotation equation (Eq. 1.40), which has only one angle θ , the angles θ_a and θ_b are separate variables in this equation because we are concatenating many retarding plates.

In general, we tend to write this equation using Jones matrices, even though Jones calculus does not account for depolarisation and photon losses. Given the relatively short fibres we use, however, these properties have negligible effects on our experiments. (If necessary, these particular properties can be determined anyway by careful characterisation of the output states themselves.) On the other hand, a Mueller matrix representation of the fibre should be capable of accounting for all scenarios and may therefore be more useful for characterising the effects of long fibres.

Stepping back from the description of a fibre (or any retarding material) as a series of rotated waveplates, a more general option for representing an optical element is the matrix:

$$X = \begin{pmatrix} ae^{i\theta_1} & be^{i\theta_2} \\ -be^{-i\theta_2} & ae^{-i\theta_1} \end{pmatrix} \tag{1.44}$$

where a , b , θ_1 , and θ_2 are all real values. Provided that $a^2 + b^2 = 1$, the determinant of this matrix will be 1. This requirement also reduces the number of independent variables to 3,

³⁶See the previous two sections for the explicit Jones and Mueller representations of these two matrices.

just as there are in Equation 1.43. The ‘series of waveplates’ description of the fibre is likely a specific case of this generic matrix. In both instances, the three unknown parameters can be determined through a comparison of the input and output states of the system, and it is expected that both options should yield equally-valid results.

Finally, a (more complicated) alternative to Mueller or Jones Calculus for describing the behaviour of light inside a fibre is to use the non-linear Schrödinger equation [42]. Fortunately, the simple matrix representations of Jones calculus are largely sufficient for our purposes.

1.5.4 Fidelity of a Polarisation State

In performing tests of our ability to transmit polarisation states via fibre, we need some measure of our success. To quantify the quality of the fibre-delivered state with respect to the ideal, we can calculate the fidelity. In quantum mechanics, the fidelity of a pure measured state $|\psi\rangle$ with respect to the expected state $|\phi\rangle$ is defined as:

$$F = |\langle\phi|\psi\rangle|^2 \tag{1.45}$$

If we are considering mixed states, we apply a generalized definition of fidelity. With respect to the ideal state ξ , the fidelity of a given polarisation with density matrix ρ is defined as follows:

$$F = \left(\text{Tr}\sqrt{\sqrt{\xi}\rho\sqrt{\xi}}\right)^2 \tag{1.46}$$

While this equation³⁷ for the fidelity of polarised light is excessively long when fully written out in terms of Stokes/Jones parameters, one can easily derive the fidelity equations for the six standard polarisation states. While the derivation is fairly straightforward, we will work through the fidelity equation for the vertical state V as an example. Using Stokes parameters, the standard equation for the density matrix ρ of a state on or within the Poincaré sphere³⁸ is given by [43]:

$$\rho = \frac{1}{2}S^\mu\sigma_\mu = \frac{1}{2}\begin{pmatrix} 1 + S_1 & S_2 - iS_3 \\ S_2 + iS_3 & 1 - S_1 \end{pmatrix} \tag{1.47}$$

where $\sigma^\mu = (\mathbb{I}, \sigma_x, \sigma_y, \sigma_z)$ is the Pauli four-vector, and a normalised intensity ($S_0 = 1$) is assumed. According to Table 1.1, the density matrix for perfect vertical polarisation must therefore be:

³⁷In this equation, ρ and ξ can actually be interchanged without altering the fidelity. It is best to use whichever form is easiest to solve.

³⁸The same equation also applies to the Bloch sphere when written in terms of the Bloch vector components.

$$\xi_V = \frac{1}{2} \begin{pmatrix} 1 + (-1) & 0 \\ 0 & 1 - (-1) \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \quad (1.48)$$

We must now take the square root of this matrix. The root of a 2×2 matrix X is defined as:

$$\sqrt{X} = \sqrt{\begin{pmatrix} X_{00} & X_{01} \\ X_{10} & X_{11} \end{pmatrix}} = \frac{1}{t} \begin{pmatrix} X_{00} + s & X_{01} \\ X_{10} & X_{11} + s \end{pmatrix} \quad (1.49)$$

$$\begin{aligned} \text{where } s &= \sqrt{\delta} \\ t &= \sqrt{\tau + 2s}, \\ \text{and } \delta &= X_{00}X_{11} - X_{01}X_{10} \\ \tau &= X_{00} + X_{11} \end{aligned}$$

From this, the root of the vertical polarisation density matrix is easily calculated to be $\sqrt{\xi_V} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$, which is the same as the original density matrix, ξ_V . In fact, if a state ξ is fully polarised, it is easy to show that $\sqrt{\xi} = \xi$ is always true. (It is never true in the case of partially-polarised states though.)

Applying the formula for a generic density matrix to represent the measured state ρ (Eq. 1.47), we now compute the following portion of the fidelity equation:

$$\begin{aligned} \sqrt{\xi_V} \rho \sqrt{\xi_V} &= \frac{1}{2} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 + S_1 & S_2 - iS_3 \\ S_2 + iS_3 & 1 - S_1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & S_2 - iS_3 \\ 0 & 1 - S_1 \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} 0 & 0 \\ 0 & 1 - S_1 \end{pmatrix} \end{aligned} \quad (1.50)$$

Using Equation 1.49 to take the square root of this result, we get:

$$\sqrt{\sqrt{\xi_V} \rho \sqrt{\xi_V}} = \frac{1}{\sqrt{2(1 - S_1)}} \begin{pmatrix} 0 & 0 \\ 0 & 1 - S_1 \end{pmatrix} \quad (1.51)$$

and so the fidelity of a measured state ρ with respect to ξ_V is:

$$F_V = \left(\text{Tr} \sqrt{\sqrt{\xi_V} \rho \sqrt{\xi_V}} \right)^2 = \left(\frac{1 - S_1}{\sqrt{2(1 - S_1)}} \right)^2 = \frac{1 - S_1}{2} \quad (1.52)$$

The fidelity of any state ρ with respect to perfectly vertical polarisation is therefore dependent only upon the S_1 parameter of ρ . From Table 1.1, we can see that a measurement of pure vertical polarisation gives $F_V = 1$, while the other standard states have a fidelity of zero with respect to the vertical state. Elliptical states will have $0 \leq F_V < 1$.

The results are similar for the other five standard polarisations, whose fidelity equations are summarised in Table 1.3.

State	Fidelity Equation
Horizontal	$\frac{1 + S_1}{2}$
Vertical	$\frac{1 - S_1}{2}$
Diagonal	$\frac{1 + S_2}{2}$
Anti-Diagonal	$\frac{1 - S_2}{2}$
Right Circular	$\frac{1 + S_3}{2}$
Left Circular	$\frac{1 - S_3}{2}$

Table 1.3: Fidelity equations for the six standard polarisation states. The Stokes parameters refer to the state vector describing the measured state ρ .

1.6 Fibre-Based Transmission of Polarised Light

A major area of research at this time is quantum communication, i.e. the ability to transmit quantum information over large distances and network quantum computers. Given its speed and how little it interacts with its environment, the photon is the most obvious platform for transferring quantum states from one point to another. As we discussed in Section 1.4.1, photon polarisation states are naturally associated with spin states. In Section 1.7, we will look at how optical selection rules dictate which charge carrier spins are generated when a polarised photon interacts with a semiconductor. It is our goal with the photon-to-spin project to encode quantum bits in photon polarisation states, which will then be transmitted to a semiconductor quantum dot by optical fibre.

The ability of optical fibres to transmit light over long distances with minimal signal loss has made them a common telecommunications technology and a prime candidate for use as a means of photon delivery in quantum communication applications. It is a well-known drawback of standard optical fibres, however, that the birefringence of their cores prevents them from preserving the polarisation of the light they transmit. Some form of compensation is therefore required to accurately deliver the desired photon polarisation states through a fibre [42, 44–46].

In this section, we will examine the effects of optical fibres on polarisation and review past attempts to compensate for those effects. While compensation technique used in the SPIN Lab differs significantly from the methods used by other groups, this general overview of the field should provide a better understanding of polarisation in fibres and may help the reader understand the choices we made when developing our own procedure. In particular, many of our system’s unique features were designed so that we could send linear, circular and elliptical polarisations by fibre; past experiments by other groups have almost always involved linear polarisations alone.

1.6.1 Fibre Birefringence

In an ideal fibre, the cross-section of the core has perfect circular symmetry, resulting in identical transmission velocities for the two orthogonal components of the polarisation. Real fibres, however, have defects which introduce a small optical anisotropy in the core and result in two group velocities; these are the fast and slow modes, oriented along the horizontal and vertical axes of the core [47]. This property, called birefringence, is also the mechanism by which waveplates operate.³⁹ There is some probability of randomly-occurring coupling between the two fibre modes, with a mean coupling length on the order of about 20-30 m [47]. The difference between the arrival times of light travelling in the two modes is called the differential group delay (DGD) [45]. Since the DGD in most fibres (even those which are many km long) is usually on the order of a picosecond or less,⁴⁰

³⁹See Section 1.4.2 for information about how waveplates work.

⁴⁰A DGD of 0.28ps was measured by Xavier et al. in their 16 km, SMF-28 optical fibre spool [45].

the delay between arrival times is typically not measurable by modern detectors, which have resolutions on the order of a few hundred picoseconds [44]. The observable effects of this birefringence-induced delay, however, are the phase shifts and rotations in the photon polarisation states at the fibre exit.

The portion of the fibre birefringence which is left over from the manufacturing process is called a residual birefringence. It is usually considered to be linear, with negligible circular birefringence. In fibres which have only a small residual birefringence, however, the birefringence which is induced locally by bends or twists in the fibre tends to dominate [45, 47, 48]. The core's birefringence may therefore vary significantly over its length, depending on how the fibre is arranged. Further changes in birefringence will occur each time the fibre experiences a local or global change in applied strain. This process, whereby the polarisation is altered by the random, spatially-varying birefringence along the fibre, is called polarisation mode dispersion (PMD) [44, 47].

The time-dependence of the PMD can have significant effects on experiments involving polarisation delivery. External factors such as temperature fluctuations and mechanical vibrations can cause the shape and/or length of the fibre to change over time. When under the influence of such forces, the relative speed and distance travelled by light in each of the two fibre modes will vary, resulting in a temporally-varying birefringence and a changing output polarisation. For the short fibres used in a laboratory setting, environmental factors may have a negligible effect on polarisation stability, so the change in polarisation for a given input state is likely to be small and to occur relatively slowly. Longer fibres, especially those which are at least a few kilometers in length, tend to experience larger, more frequent fluctuations in birefringence, as there are more opportunities for localised events to alter the birefringence [40, 46, 49]. Several past experiments have shown that significant changes in output polarisation may occur in such fibres over a period of milliseconds [40], seconds [40], or minutes [40, 46, 49], depending on the fibre's length and environment.

When it comes to delivering polarised light, standard telecom fibres installed within or between cities present several challenges which are not found in the laboratory environment. For example, fibres buried near highways may experience vibrations due to traffic. In countries prone to earthquakes, such as Japan, fibres are often suspended from utility poles over long distances, making them susceptible to vibrations caused by wind and precipitation [40]. Temperature fluctuations and temperature gradients are more likely to occur in outdoor fibres which stretch across tens of kilometers, especially if portions of them are exposed to direct sunlight. For example, Dynes et al. demonstrated in their experiments that day/night cycles, vibrations from rain, and localized heating and air conditioning were all observable in the polarisation states of photons travelling along a 45 km fibre in Tokyo [49]. Crosstalk from other fibres in a bundle will also produce noise in the channel [49].

The experiments which have been and will be discussed in this section, as well as our own photon-to-spin experiment, use only standard, single mode fibres with no built-in polarisation-maintaining properties. Polarisation-maintaining (PM) fibres do exist, however, operating on a principle of extreme optical anisotropy along the fibre modes, to the

degree that any birefringence induced by bends in the fibre (or by other external factors) is negligible by comparison. The downside to these fibres is that they only preserve perpendicular linear polarisations which are aligned with the core’s fast and slow axes, i.e. the horizontal and vertical states. A linear polarisation which is aligned with one of these modes will therefore be preserved as it propagates along the fibre. All other polarisations, however, will be projected onto these two modes, with each wave packet travelling at a different speed. The result is that one cannot use PM fibres to send linear diagonal, circular or elliptical polarisations [41, 47]. (A special type of fibre, called a spun optical fibre, can preserve circular polarisation, but there are no fibres which will preserve all polarisation states simultaneously.) PM fibres are also not guaranteed to preserve polarisation when they are exposed to cryogenic temperatures, as the stresses induced by cooling can be on the same order as the artificial stresses designed to maintain the polarisation [50].⁴¹ Therefore, we shall deal only with non-polarisation-maintaining fibres in the following sections, and more specifically, we will focus on those which have just a single spatial mode. Multi-mode fibres are far more complicated, as each spatial mode is capable of hosting an independent polarisation state from the others [51].

1.6.2 Compensating for Fibre Birefringence

The polarisation exiting a fibre can be adjusted by using polarisation controllers [42, 44–46, 48, 49] to twist, bend, squeeze or stretch the fibre, thereby artificially altering its birefringence [51]. Electronic versions of these devices generally employ feedback and optimization systems to manipulate the fibre and produce the appropriate polarisation at the fibre output. Manual polarisation controllers, such as fibre ‘paddles’ (in which a fibre spooled around a movable disc mimics the behaviour of a retarder), can also be used for brief experiments involving short fibres.⁴² Several groups have also developed alternative methods of compensating for fibre birefringence and polarisation mode dispersion. These techniques include temporal multiplexing [40, 45, 46], wavelength multiplexing [42, 44, 45] and two-fibre systems, where a second channel neighbouring the primary is used as a reference [45, 49].

One of the main methods to compensate for a fibre’s effects on polarisation is to transmit an ancillary signal with a known polarisation. The second signal provides a measure of the change in fibre birefringence and acts as a reference when attempting to correct the output polarisation. There are several different ways in which this reference signal can be

⁴¹Experimental tests of PM fibres at low temperatures are currently underway (see [50], for example). Regardless, they still do not maintain all polarisation states simultaneously, which is the main requirement for our experiments.

⁴²We attempted to use manual paddle controllers in our early polarisation experiments, but we found them to be insufficient for long-term experiments. The devices depend on the twisting of the fibre as each paddle is rotated in order to mimic the effects of a rotating waveplate. Unfortunately, the fibre appears to relax into its new position in the hours after the paddle has been moved, which changes the strain on the fibre core and slowly alters the output polarisation.

delivered. For example, the primary signal may be periodically interrupted so that reference signals can be sent instead [45, 46], and the fibre birefringence is then adjusted until the output polarisation of the reference signal matches its known input state. Once the control procedure has sufficiently compensated for the drift in fibre birefringence, transmission of the primary signal resumes. While this method is fairly reliable for short fibres in temperature-controlled laboratory conditions, the random, time-dependent variation in the birefringence of long fibres requires active compensation in order to simultaneously maintain various polarisation states [42, 44–46]. This means longer, more frequent interruptions, resulting in lower transmission rates [40, 46]. For example, to achieve proper output polarisations during their experiments, Chen et al. were only able to transmit the primary signal for 4.7, 3.1, and 1.6 minutes at a time for fibre lengths of 50, 75, and 100 km respectively [46]. They then had to interrupt the main signal to perform polarisation correction procedures. The duration of the interruptions themselves need not necessarily be long, however, as Ding et al. were able to demonstrate that they could accomplish proper polarisation alignment in a few seconds for buried fibres within cities ($\lesssim 20$ km in length) [40].

Rather than interrupting the primary signal to perform polarisation adjustments, reference pulses can be sent between signal pulses, allowing for real-time compensation. This form of temporal multiplexing has the advantage of continuous signal transmission, although it has the potential to yield poor-quality states at the fibre output, as there is little time to perform polarisation alignment procedures between signal pulses [45]. This issue can be resolved by sending the reference pulses via a second channel neighbouring the first. As the two fibres follow the same path, they are expected to experience very similar environmental conditions. It is therefore expected that the desired change in the primary signal can be achieved by adjusting both channels in the same way, using the behaviour of the reference pulses alone as a measure of success [45, 49]. While this method has been shown to work in some experiments [44], it does make the assumption that the fibres have identical or very small residual birefringence and that neither has any additional twists/rotations over its length compared to the other. Were either of these assumptions to be false (which is possible), applying the same form of adjustment to both fibres might not necessarily result in the same effects on their respective signals. The end result would be that while the fibres might be adjusted such that the reference signals are transmitted properly, the channel carrying the primary signal has additional, hidden birefringence which is still distorting its signal.

Although fibre birefringence is wavelength dependent [44, 45], an alternative to multiplexing the two signals in time is to send reference pulses at slightly different wavelengths to the primary signal [42, 44, 45]. This allows simultaneous transmission of both the primary and secondary signals, as they can be spectrally distinguished at the fibre output, and they may be transmitted either in the same channel or in neighbouring channels. Two reference signals at different wavelengths are generally deployed, one with a wavelength slightly longer than the primary signal and the other with a slightly shorter wavelength. A wavelength difference of 2-3 nm between the primary and reference signals is common and

seems to ensure that all the signals experience similar birefringences as they pass through the fibre(s) [44, 45]. Accurate, real-time polarisation control is thus possible using this method. In one experiment, Xavier et al. demonstrated that their wavelength-multiplexed system was capable of automatically correcting in a few milliseconds for a significant polarisation rotation which was artificially-induced in the fibre [44].

In many experimental tests, electronic polarisation controllers (EPC) are employed at the receiving end to quickly make the necessary polarisation adjustments [42, 44–46, 49]. Examples of such devices include the General Photonics PolarITE III, the EOSPACE Lithium Niobate Polarisation Controller, or the many varieties of OZ Optics EPCs. They are typically composed of a single fibre enclosed in a box containing piezoelectric actuators. The application of DC voltages to the actuators causes them to twist, bend, squeeze or stretch the internal fibre along different axes [51]. This allows for control of the fibre’s birefringence and hence allows for compensation of the fibre’s effects on the polarisation. These devices employ feedback and optimization systems to automate the process of manipulating the fibre and producing the appropriate polarisation at the fibre output [46, 49]. It should be noted, however, that these procedures do not actually characterise the effects of the fibre on the polarisation, but merely adjust the fibre in a feedback loop until the desired output state is reached.

The frequency with which adjustments to the fibre are required depends on the rate of change in the output polarisation. The speed of polarisation change is typically measured in radians per second, referring to a moving vector on the Poincaré sphere [40, 45]. These units are especially used when the change is induced by a polarisation stabilization system. An EPC allows the user to apply a voltage to individual piezoelectric actuators, each of which applies a force to the fibre along a given axis. One full voltage ramp of an actuator should cause a polarisation rotation of 2π in a given direction on the Poincaré sphere, followed by another 2π rotation in the opposite direction. The frequency at which the voltage is ramped then defines the speed of the polarisation change, so that 4π rad/s indicates a ramping frequency of 1 Hz [45].

To help determine how frequently compensation is required, two parameters have been defined to characterise polarisation fluctuations in optical fibres [40]. The first is the polarisation drift time, which is the average duration of stable communication before compensation is required due to polarisation drift. This parameter may be determined from the average amount of time it takes for the fidelity of a reference signal to fall below a threshold value. When performing real-time compensation for fibre effects, the tracking speed is the speed at which the bit error rate increases above a threshold value. This is therefore the speed at which the polarisation output of the fibre must be adjusted [40].

1.6.3 Long-Distance Transmission of Polarisation States

For experiments using long optical fibres, active compensation for polarisation-altering effects becomes increasingly necessary, as these fibres tend to have larger DGD values.

Most fibres, even those up to 100 km in length, have a DGD of well below 1 ps [45], but suppose a beam of light with a wavelength of 800 nm is travelling along a 100 km fibre which has a 1 ps DGD. The light in the slow mode will then be $300 \mu\text{m}$ (or 375 wavelengths) behind when the fast mode reaches the end of the fibre. There is therefore a 750π phase shift between the fast and slow modes. Given how long the fibre is, there will be quite a lot of opportunity for localised changes in length or strain due to vibrations or temperature fluctuations. Even if these are just minor influences, the totality of these effects could significantly alter the polarisation, which is already phase shifted by 750π due to the built-in properties of the fibre alone. Active compensation for these effects is therefore even more important for long fibres, particularly those which stretch over long distances in an uncontrolled outdoor environment.

Many quantum key distribution (QKD) experiments have been conducted over the past decade, often involving long-distance, fibre-based transmission of photonic qubits encoded in linear polarisation states. In spite of the need for some form of polarisation correction mechanism, many of these experiments yielded good results, with most able to achieve a quantum bit error rate target of under 1% [40, 44, 45]. Xavier et al. were able to demonstrate the stability of their system by installing a polarisation scrambler in order to purposely increase the degree of polarisation drift in their experiment. Even when performing polarisation scrambling procedures at a rate of up to 16π rad/s, they were able to maintain a quantum bit error rate of less than 6% [45]. In the work performed by Dynes et al., a total of 6.33×10^{10} quantum bits were securely and continuously transmitted via a 45 km, above-ground optical fibre over a period of 60 hours [49].

Due to the inherent transmission losses in a fibre, increasing the fibre length results in fewer photons reaching their destination. A fibre length of 100 km is typically the limit for reliable QKD experiments, as longer fibres cause the photon count rate to drop below the error rate of the receiving detector [1]. Loss rates also increase with increasing numbers of fibre splices and connectors [49]. With improvements to detector dark count noise and a reduction in the number of stray photons, however, it is expected that fibre-based, point-to-point quantum communication could be extended beyond the 100 km limit, possibly up to 165 km [1]. In one experiment conducted by Gobby et al., the experimentalists were able to achieve 88.4% visibility of quantum interference fringes for photons travelling over a 122 km optical fibre [1].

1.7 The Photon-to-Spin Qubit Transfer Protocol

With the basic concepts having been introduced, we can now gather together the information provided in the previous sections and apply it all to a single technology: the photon-to-spin hybrid quantum repeater. This system, which combines the photonic and spin qubit platforms, is designed to overcome some of the fundamental limitations of all-optical quantum repeater schemes [2]. Simultaneously, it is able to take advantage of the non-destructive heralding and long coherence times afforded by charge carrier spins trapped in lateral quantum dots.

In this section, we will introduce the concept of the quantum repeater through a review of quantum teleportation and entanglement distribution. We will then examine the NRC's proposed design for a photon-to-spin hybrid quantum repeater network, with specific focus on the interface where qubits encoded in photon polarisation states are transferred to carrier spin states. This interface is the main project upon which I have been working during my degree, and the remaining chapters of this thesis will describe the work which has been done on it so far. Much of my work has been with the optical components of this system, so the optical aspects will be covered here in much greater detail than the spin/electrical components. For an excellent overview of this project, including proposals for alternative paths should some features be too difficult to achieve, see the article *Entanglement distribution schemes employing coherent photon-to-spin conversion in semiconductor quantum dot circuits* by Gaudreau et al. of the NRC's quantum physics group [2].

1.7.1 Quantum Teleportation

We begin our study of the photon-to-spin system with one of the major research areas in the field of quantum communication: the long-distance transmission of quantum information. If a global quantum network is to be created, delicate quantum states will need to be sent to distant locations without (significant) loss of information. Photons are generally considered to be the prime candidate for sending quantum information over long distances, given the speed at which they travel, as well as their minimal interactions with their environment. Photons also have several properties which could potentially be used to encode quantum states, including polarisation, wavelength, and arrival time.

The free-space transmission of photons over great distances is obviously impractical for a variety of reasons, especially in the context of a highly-interconnected global quantum network. We must therefore rely upon optical fibres, which are already widely used and installed for regular telecommunications purposes. As with any transmission medium, however, fibres attenuate signals, and the losses increase with travel distance. Unlike classical communication, in which beams of light which are modulated in amplitude or frequency, quantum communication generally relies upon the transmission of individual photons. The attenuating effects of a fibre can therefore result in complete signal loss for long transmission distances. A standard, 100km-long, single-mode fibre is expected to attenuate a

1550 nm signal by 99% (0.2 dB/km) [1]. The maximum reasonable distance for fibre-based transmission of quantum states is therefore considered to be 100 km.

To overcome these difficulties, we must employ quantum teleportation, whereby quantum information is transmitted from point A to point B without physically moving through the intervening space. This is accomplished using quantum entanglement, which creates a connection between distantly-separated particles. In particular, the particles used for quantum teleportation are in Bell states, which are maximally-entangled two-particle states.

Bell States and Bell State Measurements

Unlike classical communication, where signals can be amplified and are encoded on beams of light, quantum states are generally encoded on individual particles. According to the no-cloning theorem, these single-particle states cannot be copied, only transferred, so it is not possible to amplify them. It is also impossible to determine the complete state of a single-particle qubit, as measurement destroys superposition and reveals only one of the basis states; information about the probability amplitudes is lost after just one measurement.

Quantum communication must therefore rely on the uniquely quantum property of entanglement, where the states of two particles are correlated (or anti-correlated) with each other.⁴³ This two-particle superposition ensures that the measurement of one particle's state forces the other particle to *instantly* collapse to a specific state, regardless of the distance separating them.⁴⁴ For example, let us consider the four maximally-entangled two-qubit states, called Bell states. They are:

$$\begin{aligned}
 |\Psi^+\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) \\
 |\Psi^-\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle) \\
 |\Phi^+\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \\
 |\Phi^-\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)
 \end{aligned}
 \tag{1.53}$$

⁴³Partial entanglement and entanglement between more than two states is also possible. See Appendix A for a closer look at these possibilities, specifically in the context of quantum teleportation.

⁴⁴As we shall see, this does not violate the ultimate speed limit (light speed) for the transmission of information. The particles are generally entangled through local operations performed when they are close together, then they are separated by travelling at or below the speed of light. (There are, at least theoretically, indications that it may be possible for particles with space-like separation to be entangled [52], but this still does not allow superluminal transmission of information.) While a measurement of one particle may collapse the states of both, the information of such a measurement can only be transmitted at or below the speed of light. Even if both particles are measured (which would automatically yield correlated results), the relativity of simultaneity ensures that it is impossible to know which one was measured first, and again the fact of each measurement can only be communicated at light speed or slower.

where we are representing the bipartite system of qubits A and B as $|i\rangle_A \otimes |j\rangle_B = |ij\rangle$ for simplicity. These four states form a basis, called the Bell basis, for two-particle entangled systems. We can see the effects of entanglement just by looking at the Bell states. If the system is in entangled state $|\Psi^+\rangle$, for example, a measurement of qubit A which yields $|0\rangle_A$ automatically forces B to be in state $|1\rangle_B$, and vice versa. The measured states of qubits in the $|\Psi^\pm\rangle$ systems will always be anti-correlated, although note that the relative phase is different for these two Bell states. For $|\Phi^\pm\rangle$, the measured qubit states are always correlated.

If we wish to know the Bell state representing a maximally-entangled two-particle system, we can perform a Bell state measurement (BSM). While we won't go into great detail about this type of measurement here, one example of a quantum circuit which could perform a Bell state measurement is the combination of a controlled-NOT (CNOT) gate, followed by a Hadamard gate [53]. The CNOT gate uses qubit A as the control and flips the state of qubit B only if A is in state $|1\rangle_A$. So for example, the application of a CNOT gate to $|\Psi^+\rangle$ yields:

$$\begin{aligned} U_{\text{CNOT}} |\Psi^+\rangle &= \frac{1}{\sqrt{2}} U_{\text{CNOT}} (|01\rangle + |10\rangle) \\ &= \frac{1}{\sqrt{2}} (|01\rangle + |11\rangle) \end{aligned} \tag{1.54}$$

where $|01\rangle$ remains unchanged because A is in state $|0\rangle_A$, and $|10\rangle$ becomes $|11\rangle$. Note that after performing this operation, the states of the two particles are no longer anti-correlated; qubit B now has a definite value of 1. This result will be true for both $|\Psi^\pm\rangle$ states, whereas if we perform the CNOT operation on one of the $|\Phi^\pm\rangle$ states, qubit B will always be zero. While this operation clearly destroys the entanglement, it does reveal the parity of the Bell state, i.e. whether the qubits were originally correlated or anti-correlated.

The Hadamard gate acts to transform one qubit from a basis state into a superposition state. If we perform this operation on qubit A, it has the following effect on the qubit states:

$$\begin{aligned} U_{\text{Had}} |0\rangle_A &= \frac{1}{\sqrt{2}} (|0\rangle_A + |1\rangle_A) \\ U_{\text{Had}} |1\rangle_A &= \frac{1}{\sqrt{2}} (|0\rangle_A - |1\rangle_A) \end{aligned} \tag{1.55}$$

It can easily be shown mathematically that the successive combination of the CNOT and Hadamard gates on the four Bell states yields the following output states:

$$\begin{aligned}
|\Psi^+\rangle &\xrightarrow{\text{BSM}} |01\rangle \\
|\Psi^-\rangle &\xrightarrow{\text{BSM}} |11\rangle \\
|\Phi^+\rangle &\xrightarrow{\text{BSM}} |00\rangle \\
|\Phi^-\rangle &\xrightarrow{\text{BSM}} |10\rangle
\end{aligned}
\tag{1.56}$$

So by measuring the states of the two qubits after the CNOT and Hadamard gates have been applied to them, we can determine which of the four Bell states formerly represented the system. (As with any measurement, the BSM destroys the entanglement, so the particles are no longer in a Bell state afterwards.) Importantly, the Hadamard gate reveals the relative phase of the original Bell state, which will be encoded on qubit A, while the CNOT gate indicates the parity of the original state, which will be encoded on qubit B. The BSM therefore expresses the initial parity and relative phase of the system, which can be used to identify the original state.

Using Bell States for Quantum Teleportation

To return to the topic of quantum teleportation, we can use the phenomenon of entanglement to transmit quantum information over long distances. Rather than physically transmitting the particle containing a desired quantum state $|\psi\rangle$, we use entanglement to transform the system at the destination into $|\psi\rangle$, while simultaneously destroying $|\psi\rangle$ at its original location.

To understand this concept, let's work through an example. Suppose we wish to send state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ from point A to point B, where the amplitudes α and β are unknown. This qubit is encoded on a particle which we will label C. If the system is delicate, perhaps susceptible to vibrations or interactions with its environment, then moving the particle over a long distance could destroy the state $|\psi\rangle$. Ideally, we would want to simply recreate the state at its destination, but the values of α and β cannot be measured without destroying $|\psi\rangle$. We must therefore find some method of transferring the unknown state without making copies or moving the particle C.

To accomplish this, let's assume we have a pair of entangled particles in one of the four Bell states, say $|\Psi^+\rangle$. Whatever these particles are, their states are robust enough that we may send one to point A and one to point B without destroying the entanglement. Photons could potentially be used, as they are fast and have minimal interaction with their environment. We will name these photons A and B for the locations to which they are sent. Then the multipartite system incorporating the entangled photons A and B, as well as the unknown state $|\psi\rangle$ on particle C, is given by:

$$\begin{aligned}
|\Psi^+\rangle_{AB} \otimes |\psi\rangle_C &= \frac{1}{\sqrt{2}}(|01\rangle_{AB} + |10\rangle_{AB}) \otimes (\alpha|0\rangle_C + \beta|1\rangle_C) \\
&= \frac{1}{\sqrt{2}}(\alpha|010\rangle_{ABC} + \alpha|100\rangle_{ABC} + \beta|011\rangle_{ABC} + \beta|101\rangle_{ABC})
\end{aligned} \tag{1.57}$$

A simple mathematical manipulation of these terms yields the following:

$$\begin{aligned}
|\Psi^+\rangle_{AB} |\psi\rangle_C &= \frac{1}{\sqrt{2}}(\alpha|010\rangle_{ABC} + \alpha|100\rangle_{ABC} + \beta|011\rangle_{ABC} + \beta|101\rangle_{ABC}) \\
&= \frac{1}{\sqrt{8}}(2\alpha|010\rangle_{ABC} + 2\alpha|100\rangle_{ABC} + 2\beta|011\rangle_{ABC} + 2\beta|101\rangle_{ABC} \\
&\quad + \alpha|111\rangle_{ABC} - \alpha|111\rangle_{ABC} + \alpha|001\rangle_{ABC} - \alpha|001\rangle_{ABC} \\
&\quad + \beta|110\rangle_{ABC} - \beta|110\rangle_{ABC} + \beta|000\rangle_{ABC} - \beta|000\rangle_{ABC}) \\
&= \frac{1}{2}(\alpha|\Phi^-\rangle_{AC}|1\rangle_B + \alpha|\Phi^+\rangle_{AC}|1\rangle_B + \alpha|\Psi^+\rangle_{AC}|0\rangle_B \\
&\quad - \alpha|\Psi^-\rangle_{AC}|0\rangle_B + \beta|\Psi^+\rangle_{AC}|1\rangle_B + \beta|\Psi^-\rangle_{AC}|1\rangle_B \\
&\quad + \beta|\Phi^+\rangle_{AC}|0\rangle_B - \beta|\Phi^-\rangle_{AC}|0\rangle_B)
\end{aligned} \tag{1.58}$$

and collecting together the Bell state terms, we get:

$$\begin{aligned}
|\Psi^+\rangle_{AB} |\psi\rangle_C &= \frac{1}{2} [|\Psi^+\rangle_{AC} (\beta|1\rangle_B + \alpha|0\rangle_B) \\
&\quad + |\Psi^-\rangle_{AC} (\beta|1\rangle_B - \alpha|0\rangle_B) \\
&\quad + |\Phi^+\rangle_{AC} (\alpha|1\rangle_B + \beta|0\rangle_B) \\
&\quad + |\Phi^-\rangle_{AC} (\alpha|1\rangle_B - \beta|0\rangle_B)]
\end{aligned} \tag{1.59}$$

We have now written particles A and C in terms of the entangled Bell states, while particle B is written as a superposition state with the unknown amplitudes α and β . Note that we have not actually performed any operations on these particles yet; we have merely rewritten Equation 1.57. This means that A and B are actually still entangled, and C is not entangled with A. This equation shows very clearly, however, that if a BSM is performed on particles A and C, which are both in the same location, the system will collapse into a state where the amplitudes α and β are immediately transferred to the distant particle B, which is forced into a superposition state. (See Appendix A for a more detailed analysis of this procedure.)

In the case where the Bell state measurement on A and C reveals them to be in state $|\Psi^+\rangle_{AC}$, particle B will automatically assume the desired state $|\psi\rangle$, which we wanted to

send. If the BSM reveals one of the other Bell states, however, we will need to perform a unitary transformation on particle B in order to reproduce state $|\psi\rangle$ (in particular, one of the Pauli operators). The results of the BSM must therefore be communicated from point A to point B via classical channels so that we know what type of operation we need to perform on particle B to reproduce $|\psi\rangle$. Even in the case where the BSM reveals $|\Psi^+\rangle_{AC}$, we must still communicate that particle B is already in state $|\psi\rangle$ and does not require transformation. We are therefore always limited to light speed communication or slower, since we must always send the BSM results via a classical channel.

This is the fundamental protocol underlying quantum teleportation. To transfer a delicate state from one location to another, we require a pair of entangled particles, a BSM, and communication over a classical channel (which necessarily limits the transmission speed). Particles A and B may initially be in any of the four Bell states, provided we know which one it is. Equation 1.59 looks similar (but not identical) if A and B are entangled in one of the other three Bell states. Note also that once this procedure is complete, particle C has been projected onto either $|0\rangle_C$ or $|1\rangle_C$; the state $|\psi\rangle$ exists only on B, in accordance with the no-cloning theorem.

If particles A and B are photons transmitted through a fibre, then the attenuating properties of the fibre limit them to travel distances of roughly 100 km each [1]. We would therefore be limited to quantum communication over a total distance of ≤ 200 km if each particle is sent through a separate 100 km fibre. In order to transmit quantum states over greater distances, we will need a different protocol. This extended form of quantum teleportation requires the use of quantum repeaters.

1.7.2 Quantum Repeaters and Entanglement Distribution

If we are to create a global quantum network, quantum information must of course be transmitted over much greater distances than the 100 km limit set by optical fibres. A quantum repeater network allows for long-distance quantum teleportation by employing a series of BSM nodes and sources of entangled photons (SEPs), all connected by relatively short fibres. The goal of the repeaters is to extend the distance separating a pair of entangled photons through a process called entanglement swapping or entanglement distribution.

Entanglement Distribution

We can investigate entanglement distribution by considering two pairs of entangled particles, pair A and pair B. Each pair is in one of the Bell states of Equation 1.53. Let's say they are both in state $|\Psi^+\rangle$. The combined state of these two pairs is therefore:

$$\begin{aligned}
|\Psi^+\rangle_A |\Psi^+\rangle_B &= \frac{1}{2}(|01\rangle_A + |10\rangle_A)(|01\rangle_B + |10\rangle_B) \\
&= \frac{1}{2}(|01\rangle_A |01\rangle_B + |10\rangle_A |01\rangle_B + |01\rangle_A |10\rangle_B + |10\rangle_A |10\rangle_B)
\end{aligned} \tag{1.60}$$

We now write pair A as being composed of particles 1 and 2, while pair B is composed of particles 3 and 4. Let's rearrange Equation 1.60, collecting the states of particles 1 and 4 together and the states of 2 and 3 together:

$$\begin{aligned}
|\Psi^+\Psi^+\rangle_{AB} &= \frac{1}{2}(|01\rangle_{14} |10\rangle_{23} + |11\rangle_{14} |00\rangle_{23} + |00\rangle_{14} |11\rangle_{23} + |10\rangle_{14} |01\rangle_{23}) \\
&= \frac{1}{4}(2|01\rangle_{14} |10\rangle_{23} + 2|11\rangle_{14} |00\rangle_{23} + 2|00\rangle_{14} |11\rangle_{23} + 2|10\rangle_{14} |01\rangle_{23} \\
&\quad + |10\rangle_{14} |10\rangle_{23} - |10\rangle_{14} |10\rangle_{23} + |00\rangle_{14} |00\rangle_{23} - |00\rangle_{14} |00\rangle_{23} \\
&\quad + |11\rangle_{14} |11\rangle_{23} - |11\rangle_{14} |11\rangle_{23} + |01\rangle_{14} |01\rangle_{23} - |01\rangle_{14} |01\rangle_{23}) \\
&= \frac{1}{2}(|\Psi^+\rangle_{14} |\Psi^+\rangle_{23} - |\Psi^-\rangle_{14} |\Psi^-\rangle_{23} + |\Phi^+\rangle_{14} |\Phi^+\rangle_{23} - |\Phi^-\rangle_{14} |\Phi^-\rangle_{23})
\end{aligned} \tag{1.61}$$

From this result, it is very evident that by performing a Bell state measurement on particles 2 and 3, particles 1 and 4 will be projected into an entangled state. Importantly, the new entangled state of 1 and 4 will be identifiable from the results of the BSM, since we know the original Bell states of the four particles. Therefore, performing a BSM on two particles, each from a different entangled pair, will entangle the remaining two particles (also from different pairs). This quantum mechanical phenomenon is known as entanglement distribution, or entanglement swapping, and we can now look at an important practical implementation of it.

The Quantum Repeater

Figure 1.16 shows the basic concept of entanglement distribution via a quantum repeater network. Two sources of entangled photons each send their photons through separate optical fibres. One photon of each pair is sent to a quantum memory (QM), which is necessary in case one of the photons arrives before the other. Once both photons have arrived, a Bell state measurement is performed on them. In accordance with Equation 1.61, the BSM entangles the two remaining photons at points A and B. This setup, with the two SEPs, the two QMs, and the BSM, is a quantum repeater. If we suppose each fibre to be 50 km long, we now have a pair of entangled photons separated by 200 km, and we can perform quantum teleportation on them using the protocol discussed in Section 1.7.1.

Evidently, quantum repeaters allow quantum teleportation over much greater distances than the 100 km limit of a single fibre. The repeater of Figure 1.16 can be expanded by adding more QM/BSM nodes and SEPs, allowing for quantum communication over even larger distances.




Figure 1.16: A generic quantum repeater system for fibre-based entanglement distribution, featuring two SEPs, two QMs, and one BSM. By performing the BSM on one photon from each SEP, we can entangle the photons at points A and B.

1.7.3 The Photon-to-Spin Interface

One common version of the quantum repeater is the all-optical system, which features a 50/50 beam splitter for performing Bell state measurements on photons. The first experimental demonstration of quantum teleportation occurred in 1997, using a polarising beam splitter and qubits encoded in photon polarisation states [54]. In this setup, a single-photon detector is placed at each output of the beam splitter. Two photons enter the beam splitter simultaneously at the two input ports. This projects them into one of the four Bell states, as their polarisation determines which way they exit the splitter. While we won't go into the full details here, the physics of the beam splitter ensures that the $|\Phi^\pm\rangle$ and $|\Psi^+\rangle$ Bell states cause both photons to exit on the same side. Only the $|\Psi^-\rangle$ state results in the photons exiting on opposite sides of the splitter, due to constructive interference. This is therefore the only state which results in both detectors detecting a photon. This type of repeater is therefore limited to an overall efficiency of 25%, as it can only detect one of the four Bell states [2].⁴⁵ It can be shown, however, that similar techniques using three beam splitters (called the Innsbruck method) or qubits encoded in photon number states (called Fock states) are capable of signaling the creation of both the $|\Psi^\pm\rangle$ states, but the overall efficiency is still limited to 50% [55].⁴⁶

To overcome these limitations, we instead propose to merge the photonic qubit platform, which provides us with the speeds necessary for long-distance communication, and the spin qubit platform, which has the potential to allow BSMs with 100% efficiency [2]. Figure 1.17 shows the setup for this hybrid quantum repeater. Again, two SEPs each entangled photon pairs through separate fibres, with the qubits encoded in the photon polarisation states. One photon of each pair reaches a photon-to-spin conversion interface, where the photonic

⁴⁵We need both detectors to click because it indicates the successful arrival of both photons. When both photons exit the same side of the beam splitter, only one detector clicks, and we don't know if both photons were measured simultaneously by the detector or if one was lost before it arrived at the beam splitter. The $|\Psi^-\rangle$ is therefore the only one which definitively tells us that the photons were properly entangled.

⁴⁶Even using auxiliary photons, it is impossible to ever achieve 100% BSM efficiency for an all-optical system [2].




Figure 1.17: The photon-to-spin hybrid quantum repeater, with two conversion interfaces to transfer qubits from photonic states to spin states. The spin states, which will act as the quantum memory, should allow for a full BSM to be performed, improving upon the 50% efficiency of the all-optical system.

qubit is transferred to the spin state of a carrier trapped in a lateral quantum dot. The spin states have long coherence times, so they act as the quantum memories. A Bell state measurement is then performed on the two spin states, with a theoretical efficiency of 100% [56].⁴⁷ The result is that the remaining two photons at points A and B are entangled and available for quantum teleportation procedures.

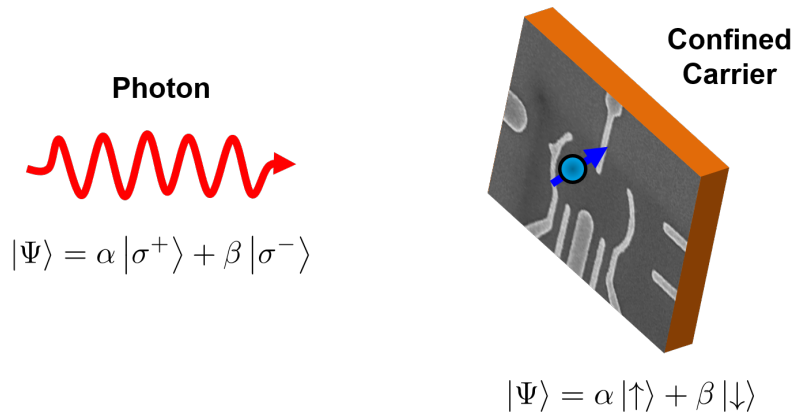


Figure 1.18: The transfer of a photonic qubit, $|\Psi\rangle$, to the spin state of a charge carrier trapped in a quantum dot.

The concept of photon-to-spin qubit transfer is summarised in Figure 1.18, where we see a photon incident on a lateral DQD. The qubit state $|\Psi\rangle$ is encoded in the photon polarisation, written as a superposition of right- (σ^+) and left- (σ^-) circular polarisation.⁴⁸

⁴⁷Experimental BSM efficiencies may be lower, due to the fact that the protocol requires the system be projected onto all four Bell states, one at a time. The time required for this procedure, as well as the changes in the system required to perform these projections, suggests that the spin states could decohere before they are successfully entangled.

⁴⁸Recall that any polarisation can be written in the basis of right- and left-circular states, so the amplitudes α and β may be encoding any arbitrary qubit.

This is the state which is to be transferred to the spin of a carrier in one of the quantum dots.

An optical mask is generally placed over the lateral DQD sample, with a small hole above one of the dots to ensure the photon can only be absorbed by that dot. As we saw in Section 1.4.1, the σ^+ state corresponds to a photon with spin +1, while σ^- indicates spin -1 . From Section 1.2.3, we know that the absorption of the photon by a direct band gap semiconductor material generates an electron-hole pair with total angular momentum 1. Optical selection rules will dictate the angular momentum of each carrier. Depending on the construction of our DQD sample, we can use this photoexcitation to transfer the photonic qubit to either an electron spin or the angular momentum of a hole. We will look at both of these options in a little more detail, but in either case, we use a magnetic field in Voigt configuration, i.e. with the field direction perpendicular to the sample growth direction [2].




Figure 1.19: Transfer of a qubit from a photon polarisation state to the spin state of a) an electron, or b) a heavy hole. The optical selection rules dictate which spin state is generated for a given polarisation state.

In the case of a sample designed to trap electrons, the effective, in-plane electron g-factor must be engineered to be zero. This ensures that the energies of the s-shell electron spin states are degenerate and can be labelled by their projection along the z-axis (i.e. the growth direction): $|\uparrow\rangle$ and $|\downarrow\rangle$. The heavy holes are also degenerate, since their in-plane g-factor is nearly zero, but Zeeman splitting ensures the light holes are separated in energy. Labelling the light hole angular momenta as \uparrow_L for $+\hbar/2$ and \downarrow_L for $-\hbar/2$, their in-plane states are $|LH^\pm\rangle = (|\downarrow_L\rangle \pm |\uparrow_L\rangle)/\sqrt{2}$. Figure 1.19a) shows the energy diagram for this scenario. By ensuring the arriving photon has the energy of the gap between the electrons

and the $|LH^+\rangle$ state, the absorption of a σ^+ photon will generate a spin-up electron, while σ^- produces a spin-down electron.⁴⁹ A $|LH^+\rangle$ hole is generated in both cases, which means the light hole state is not entangled with the electron state. We can see this from the mathematical description of the system:

$$\alpha |\sigma^+\rangle + \beta |\sigma^-\rangle \rightarrow (\alpha |\uparrow\rangle + \beta |\downarrow\rangle) |LH^+\rangle \quad (1.62)$$

The hole state is evidently uncoupled from the electron state. As the dot is only designed to trap electrons, the hole escapes to elsewhere in the sample, and the photonic qubit has therefore been coherently transferred to the electron spin state.

Unfortunately, there is a slight issue with this scenario. Since the photon energy is so large, a heavy hole could be generated instead, with the corresponding electron having an energy well above the conduction band ground state. This would ruin our transfer protocol, as either heavy hole state could be generated, and the process of electron relaxation to the ground state could disturb its spin state. Therefore, an alternative option is to use a sample designed to trap holes in the quantum dot, again in Voigt configuration [2]. With the heavy hole g-factor being nearly zero, the heavy hole states are degenerate and can be represented by $|\uparrow_H\rangle$ and $|\downarrow_H\rangle$ for total angular momenta of $+3\hbar/2$ and $-3\hbar/2$ respectively. In this scenario, the electron states are non-degenerate, and they are represented by $|e^\pm\rangle = (|\downarrow\rangle \pm |\uparrow\rangle)/\sqrt{2}$ because the electron g-factor is non-zero. With $|e^+\rangle$ being lower in energy, we will focus on the optical transition between $|e^+\rangle$ and the two heavy hole states. From the energy diagram in Figure 1.19b), we see that the absorption of a σ^+ photon now generates a spin-up hole, while σ^- produces a spin-down hole. The absorption event can be represented by:

$$\alpha |\sigma^+\rangle + \beta |\sigma^-\rangle \rightarrow (\alpha |\uparrow_H\rangle + \beta |\downarrow_H\rangle) |e^+\rangle \quad (1.63)$$

Again, as $|e^+\rangle$ is always produced, the electron and hole are not entangled. This type of quantum dot traps only the heavy hole, whose angular momentum is now encoding the desired qubit. This version of photon-to-spin does not suffer from the same issue as the electron-trapping version, as both the holes and electrons are generated in their respective ground states.

The basics of photon-to-spin qubit transfer have already been demonstrated experimentally. Free-space, horizontally/vertically-polarised single photons have been shown to transfer their momenta to electron-hole pairs in an electron-trapping quantum dot (albeit using heavy holes rather than light holes in a variant of the protocol given here) [57]. Our goal is to move to the next steps, where we shall attempt to demonstrate that this procedure can be achieved using optical fibres for photon delivery and that a full BSM can be performed on the photo-generated spins without causing them to decohere.

⁴⁹We won't go into the mathematics behind these optical transitions, as they are not overly relevant to the rest of the thesis, but the optical selection rules dictating them are well known.

1.7.4 Advantages of a Photon-to-Spin Quantum Repeater

The photon-to-spin repeater has many advantages over the all-optical system, including its potential to allow for full Bell state measurements [56]. We can also take advantage of the inherently long coherence times of the spin qubits, which will act as our quantum memory. Optical versions of a quantum repeater might typically rely on atomic ensembles to store the photons [2]. This presents several problems, as there is no way of detecting whether the photons have been successfully absorbed by the ensemble (any measurements would cause the stored qubit to be lost), and there is also the additional step of re-emission prior to the BSM (which introduces yet more possibility for losses in the system). The photon-to-spin system is able to solve both of these problems by using the carrier spins themselves as the memory, with no re-emission necessary, and a nearby quantum point contact can be used to detect the presence of the photo-generated electrons in the dots. As we saw in Section 1.3.1, the QPC detects the charge of the carrier only and does not affect the spin state, since spin and charge are not coupled. We can therefore herald the arrival and successful absorption of a polarised photon through QPC charge detection.

In Section 1.6, we learned that optical fibres alter the polarisation of the light sent through them. This makes it difficult to test the successful transfer of a polarisation state to a spin state, since the polarisation is altered in a random way. Fortunately, we have recently developed a system which is capable of compensating for the fibre's effects on the polarisation. This will be the topic of Chapter 5.

Incidentally, should our attempts to use polarisation encoding fail, this system is also compatible with time-bin encoding [2]. A time-bin qubit is encoded based on its arrival time, which may be in one of two time slots. The qubit is prepared using the setup of a Mach-Zehnder interferometer, which is comprised of two beam splitters and a delay line.

Overall, this system has quite a lot of potential to surpass the all-optical setup as a far more efficient quantum repeater. Indeed, theoretical modelling has already shown it to have a much higher qubit transmission rate over long distances (>500 km) than the all-optical system [2]. In the following chapters, we will explore some of the stages involved in creating the photon-to-spin system, with particular focus on its optical components.

Chapter 2

The SPIN Lab

The SPIN Lab, or Spin-Photon Interactions in Nanostructures Laboratory, is a recently-established quantum research facility at the National Research Council (NRC) in Ottawa. Featuring milli-Kelvin temperatures, high-resolution optics, magnetic fields up to 8T, and high-frequency electronics, the facility is designed for complex experiments involving interactions between single photons and individual electrons/holes. Its versatility and wide-ranging capabilities make this system ideal for testing a great variety of different nanostructure samples and for conducting experiments with many potential applications in the fields of quantum computing and quantum communication.



Figure 2.1: The SPIN Laboratory is a specially-designed facility for performing optical and electrical experiments on different types of nanostructures.

Combining quantum optics and quantum electronics technologies, the design and construction of the lab was undertaken with the goal of testing the concept of a fibre-based photon-to-spin quantum repeater, which would merge the photonic and spin qubit platforms. This project involves the delivery of single, polarised photons via fibre (Chapter 5) to lateral quantum dots in a dilution refrigerator. Although this project remains on-going, we have also demonstrated the versatility of this system by testing other types of samples, such as two dimensional materials and high-mobility Hall bars. In particular, we have de-

voted a lot of time to the study of nanowire quantum dots, where we sought to create a fibre-coupled source of single photons (Chapter 4).

The SPIN Lab is owned by the Quantum Physics group at the NRC, under the supervision of group leader Dr. Andrew Sachrajda. The majority of the lab construction and equipment installation was conducted by Dr. Louis Gaudreau, Piotr Zawadzki, and Dr. Sergei Studenikin, with assistance from Joseph McKee, Dr. Alex Bogan, Dr. Guy Austing, and myself. We would also like to extend our thanks to Professor Jan Kycia of the University of Waterloo for his many helpful discussions and suggestions as we constructed and designed the lab components, particularly in the case of the high-frequency electrical systems.

In this chapter, I will describe the various features of the SPIN Lab from a technical standpoint, focusing on its electrical (Section 2.3) and optical (Section 2.4) systems. We will also examine in detail the cryogenic SPIN setup, which is a specially-designed unit attached to the mixing chamber of the lab’s dilution refrigerator. The cryogenic SPIN setup hosts most of the components necessary for us to perform photon-to-spin experiments, including a cryogenic optical table, a RF-reflectometry circuit, a sample holder, and a fibre-based photon delivery system. This chapter therefore introduces the laboratory systems and equipment which were used to conduct the experiments described in the remainder of this thesis.

2.1 The Equipment and Systems of the SPIN Lab

The SPIN Lab’s major components have been designed and/or installed with the express purpose of testing the proposed photon-to-spin system. The experiments take place within a dilution refrigerator, which hosts the quantum dot samples, a reflectometry circuit for fast charge detection, and optical fibres for delivery of polarised light. There are also several room-temperature systems which complement the cryogenic components, such as optical elements for controlling polarisation in an optical fibre and electronics for manipulating the quantum dot samples in the fridge. In this section, we will take a broad look at the main components of the SPIN Lab, with further details being provided in the following sections and chapters as necessary.

2.1.1 The Dilution Refrigerator

The majority of the SPIN Lab’s experimental apparatus is centered around a dilution refrigerator (Bluefors, BF-LD250), which has a base temperature of under 10 mK at the mixing chamber (MC).¹ Images of the fridge are shown in Figure 2.2, with the different levels labelled according to the temperatures they reach when the fridge is running. At

¹The mixing chamber is the coldest part of the fridge. In an attempt to stay on topic with the main subject of this thesis, i.e. the development of the photon-to-spin system, I shall assume the reader has a

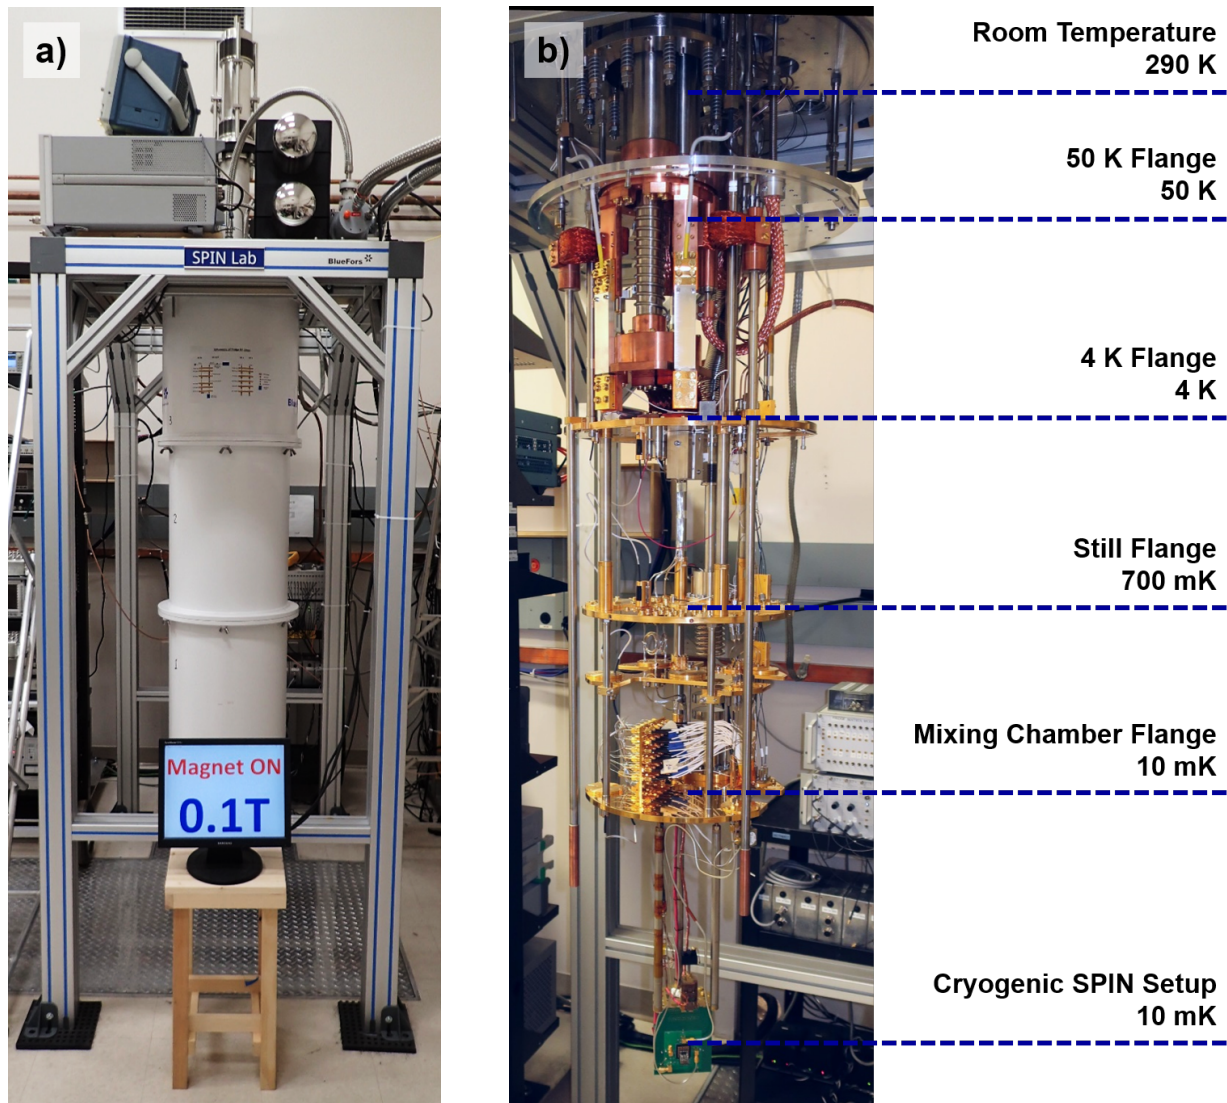


Figure 2.2: a) Closed dilution refrigerator with magnetic field strength indicator. b) Open refrigerator with the various levels labelled according to function and temperature. The sample is loaded onto the cryogenic SPIN setup, which is attached to the mixing chamber. Photo of open fridge provided by Piotr Zawadzki.

the bottom of the fridge, connected to the MC flange, is a custom-built, cryogenic SPIN setup, where our nanostructure samples are loaded and the majority of our experiments are performed. Its components will be described in more detail in the next section. The SPIN setup is positioned within an 8T, 5" bore, superconducting magnet (American Magnetics,

basic understanding of the operation of a dilution refrigerator. Detailed knowledge of such systems will not be needed in order to comprehend the contents of the following sections, however, and basic introductions to the operation of a dilution refrigerator may be easily found online.

custom model). RF and DC lines for electrical control of the sample, as well as optical fibres for delivery of light to the sample, extend from room temperature at the top of the fridge to the experimental stage at the bottom.

Installed in 2015, the SPIN Lab’s cryogen-free dilution refrigeration unit circulates a 20% ^3He - 80% ^4He mixture, allowing it to reach a base temperature of under 10 mK. The MC thermometer stops operating below 7 mK, a limit which we often surpass during normal operation. The system typically requires 2 days to cool from room temperature to 4 K by pulse tube cryocooler (or a little over 1 day without the magnet). Once the dilution unit is started, it takes another 4 hours to cool to base temperature, which we can usually reach without applying heat to the still. The fridge is very stable at base temperature, with fluctuations of just 0.1 mK recorded by the thermometer. Heating the fridge from base temperature to room temperature can take up to three days. In total, it therefore requires about 5 days to switch samples in the fridge.

The success of our experiments (and in some cases, our ability to even perform them) depends upon the smooth operation of the dilution refrigerator. A lot of attention is therefore given to the behaviour of the system to ensure that any issues are noticed and resolved before they can interfere with our experiments. Regular recording of all system parameters and periodic comparisons to previous cool-downs are key to ensuring that the system is operating properly.

Care is also taken to ensure that unnecessarily large amounts of heat are not dumped into the system while it is operating at base temperature. Certain experiments dissipate heat in the fridge and may raise the temperature significantly (sometimes rising above 200 mK). Heat sources at the mixing chamber include electrical currents, moving parts, and light from optical experiments. We have therefore worked hard to ensure that all components in the fridge, especially those on the cryogenic SPIN setup, have the proper degree of thermal anchoring required in order to sustain milli-Kelvin temperatures during experiments. In general, we attempt to keep the system below 500 mK when the dilution unit is operating, as the properties of the ^3He - ^4He mixture begin to change dramatically beyond that point. If we expect a great deal of heat to be generated in the fridge during an experiment, sometimes we simply shut down the dilution unit and conduct the experiment at a temperature of 4K instead.

2.1.2 The Cryogenic SPIN Setup

The majority of the experimental apparatus in the fridge consists of a single unit containing all necessary mechanical, electrical, and optical components, as shown in Figure 2.3. We call this the cryogenic SPIN setup, and it was designed and constructed at the NRC by Dr. Louis Gaudreau, Piotr Zawadzki, and Dr. Sergei Studenikin. The main components of the SPIN setup are the following:

- 1) **Sample Holder & Reflectometry Circuit:** Nanostructure samples are mounted on a specially-designed printed circuit board (PCB). The PCB hosts the DC chan-

nels for electrical control of the nanostructures, as well as the RF reflectometry circuit used for fast charge detection of quantum dot populations. Section 2.3 describes our preliminary tests of the reflectometry system.

- 2) **Fibre Positioning System:** Located at the front of the system, three piezoelectric positioners are used to focus a lensed optical fibre on the sample for accurate photon delivery. Chapter 3 provides details about our techniques for focusing the fibre on nanostructures.
- 3) **Cryogenic Optical Setup:** Consisting of a rotating quarter-wave plate, a linear polariser, and a photodiode, this system is used to identify the polarisation of light reaching the fridge through the lensed fibre. See Section 2.4 and Chapter 5 for more information.




Figure 2.3: a) The cryogenic SPIN setup, featuring (left to right) the optical components, printed circuit board, and piezoelectric positioners. (We will look at these components in more detail in Figure 2.4.) b) Front view of the SPIN setup, showing the reflectometry circuit, the sample, and the polarisation delivery system. The photos in both a) and b) were provided by Piotr Zawadzki. c) Full experimental structure in fridge, with the SPIN setup attached to the mixing chamber by stainless steel tubes and a copper ring.

All of these components are mounted on a custom-made brass plate (1/16" thick, 4.64" maximum wide), which has a ‘Swiss cheese’ pattern of holes and slots. These openings help

minimize Joule heating due to eddy currents in the presence of a changing magnetic field. The brass plate is connected to a 1/4"-thick copper ring by three thin-walled, stainless steel tubes (1/2" diameter, 12.4" long), which are positioned at 120° intervals around the ring for mechanical stability (Fig. 2.3c). The copper ring is bolted to the bottom of the MC flange in the refrigerator, and the stainless steel tubes center the experimental apparatus within the fridge magnet. As the thermal conductivity of the stainless steel tubes is too low at milli-Kelvin temperatures, the brass plate and all components are thermally linked to the copper ring (and therefore to the MC) by an oxygen-free copper cable (4mm diameter), screwed down at both ends. All wiring from the components and sample is thermally connected to the MC via a copper thermal sink bolted to the copper ring.




Figure 2.4: The cryogenic SPIN setup has a variety of components to allow many different types samples to be tested. a) Sample mounted on circuit board: 1. mounted sample on wire-bonded chip carrier, 2. RF reflectometry circuit, 3. one of the RF lines. Red circle indicates hole in chip carrier/PCB for polarisation identification. b) Fibre positioning system: 4. piezoelectric XYZ nanopositioners, 5. fibre clamp and lensed fibre. c) Polarisation identification system: 6. piezoelectric rotator with mounted quarter-wave plate, 7. linear polariser, 8. photodiode, 9. DC lines for electrical control of sample. The photos in both b) and c) were provided by Piotr Zawadzki.

The nanostructure samples are loaded onto a printed circuit board (3.5" × 4.5") attached to the cryogenic SPIN setup. (We'll look at sample mounting in more detail in Section 2.2.) Figure 2.4a) shows a sample on the PCB, where it is connected to the RF reflectometry circuit for charge sensing and to the DC lines for electrical control of the nanostructure. Five of the pins on the chip carrier are connected to high-frequency lines, two of which are

part of the RF-QPC, and the remainder of which are used for the application of microwave pulses to the sample. Some details of the system’s electrical components, particularly the reflectometry circuit, will be described in Section 2.3. As the focus of this thesis is on the optical components of the SPIN Lab, however, I will leave detailed descriptions of the electronics for future publications.

The optical components of the cryogenic SPIN setup consist primarily of a fibre-based light delivery system and a polarisation identification system. The former delivers photons to the sample via a lensed fibre, which is focused on the nanostructure by three XYZ nanopositioning stages from Attocube Systems (two ANPx101/RES, and an ANPz101/RES). These are the copper-coloured devices which can be seen in Figure 2.4b). By scanning the fibre in a grid pattern and measuring the sample’s response to light, we can use this system to perform high-resolution imaging of various nanostructure devices, as we shall explore in Chapter 3. Once we have imaged a nanostructure, we can use the positioners to focus the fibre at the appropriate point on the sample for light delivery.

The positioning stages can also move the fibre so that it directs light through a small hole in the PCB, which is marked by the red circle in Fig. 2.4a). This sends the light to the back of the circuit board, where we have placed the optical components required for identifying the polarisation state of the light. Figure 2.4c) shows the rear of the cryogenic SPIN setup, which holds a quarter-wave plate in a piezoelectric rotator, a linear polariser, and a photodiode. The combination of these three elements allows one to uniquely identify the polarisation of the light exiting the lensed fibre, a necessary feature of our system given that optical fibres alter the polarisation of the light they transmit. Section 2.4 has more information about our cryogenic and room-temperature optical setups, while in Chapter 5, I will describe our method of compensating for the polarisation-altering effects of optical fibres. This technique allows us to use a fibre to deliver photons with any polarisation state to a sample in the dilution refrigerator.

2.1.3 Equipment Control and Experiment Automation

All equipment in the SPIN Lab is controlled using custom-made Python software called Platypus, which was created by Dr. Alex Bogan during his time as a graduate student at NRC. Platypus is a framework conceptually similar to LabVIEW, except the user creates their own custom drivers for remote control of their equipment. They can then write experiment automation files, which use the device drivers to call upon the instruments to perform their functions in a specific order. Platypus’ functions are accessed through a Python shell, from which all equipment and experiments can be controlled. A second window displays the results of the most recent measurements made by the equipment. A button in this second window allows the user to pause an experiment so that the measurement parameters can be adjusted before the experiment resumes. Specially-written scripts allow for multiple experiments to be run automatically overnight or over a weekend without the need for user input.

The adaptable, customizable nature of Platypus, as well as the wide variety of options it provides for equipment and experiment control, has made this software invaluable to our experimental efforts. All of the experiments described in this document were conducted using Platypus (except where manual control of equipment is noted), and it is doubtful that we would have been able to progress as rapidly through so many different types of experiments without it. Most of the equipment drivers and experiment automation files used in our lab were written by myself and Dr. Louis Gaudreau of NRC on an as-needed basis, with periodic updates to expand the capabilities of our system. For further details about the programming of device drivers and experiment automation files using Platypus, see Appendix B.

In addition to Platypus, another program often used in the SPIN Lab is an interactive data plotting and analysis program, which I have been slowly constructing over the course of my work at NRC. Upon startup, the program automatically plots the most recently-acquired data set, be it a line plot or a heat map, and provides the user with customizable options for performing operations on the data (e.g. normalization, background subtraction, differentiation, etc.). The user may easily scan backwards through previous data sets, all of which are automatically plotted, and many options are provided for altering the appearance of the graph for easier viewing. This program, which is described in more detail in Appendix C, has saved us a great deal of time over the past few years, as all data plotting is now automatic and preliminary data analysis can be achieved with the press of a few buttons.

2.2 Quantum Dot Samples

We have so far tested several different types of samples in the SPIN Lab fridge, but we have particularly focused on quantum dots. All samples are loaded onto the circuit board in the fridge, where systems are provided for electrical and optical control of the nanostructures. Our setup enables quick sample swapping; each sample is attached to its own chip carrier, which can be easily installed and removed from the circuit board in a few minutes.²

In this section, we will start by looking at how samples are mounted in the SPIN Lab fridge. We will then discuss the photon-to-spin lateral quantum dot samples. Unfortunately, fabrication issues caused these samples to malfunction, so this will be just a brief look at them to round out the photon-to-spin concept. In general, this thesis will focus on the optical side of the photon-to-spin project, while future publications will deal with the quantum dot samples.

²The entire sample swapping process does take a little more time than that, however, as the time for the fridge to warm up and cool down again is about 5 days. The new sample must also be wire-bonded to its chip carrier, although that can be done as the fridge is warming.

2.2.1 Mounting a Sample in the Refrigerator

Using GE varnish, each nanostructure sample is mounted on its own 40-pin chip carrier ($7/8'' \times 7/8''$) designed at NRC, and the sample is wire-bonded to the pins if electrical control is required. Figure 2.5a) shows one of these chip carriers, as well as an example of the smaller, 19-pin chip carriers ($1/2'' \times 1/2''$) which our system can also accommodate. The chip carrier may have a hole drilled into it to allow light to pass through and reach the polarisation identification setup at the back of the cryogenic SPIN setup. (The circuit board has a corresponding hole to let the light through.) A socket designed to receive the chip carrier pins is attached to the PCB on the SPIN setup. An unmounted socket is shown in Fig. 2.5a).

Mounting the sample onto the circuit board is a delicate procedure. If too much force is applied to the board, it will bend and break. Unfortunately, the optical components mounted at the rear of the PCB prevent one from pushing from behind while mounting the sample. A small, home-made device is therefore used to carefully load the chip carrier, as shown in Figure 2.5c). Four screws are used to gently press the chip carrier into the socket, while two small lips on the bottom of the device slide under the edges of the socket and press against its underside. We can therefore press the chip carrier and socket together without pushing against the PCB. Figure 2.5c) shows a sample mounted on a chip carrier, wire-bonded to the pins, and loaded onto the PCB. Note the grounded gold plate below the sample, which acts as a transmission line ground if microwaves are applied to the nanostructure. All lateral quantum dot samples used in the photon-to-spin project, as well as the nanowire samples discussed in Chapter 4, are mounted in the fridge in this manner.³

The vertical orientation of the socket means that the sample must be loaded vertically, i.e. with the growth direction horizontal (see Fig 2.5c). While the orientation of the sample might normally be irrelevant, it becomes important during magnetic field tests, as we currently have a single-axis magnet with a vertically-oriented magnetic field. The system is therefore in Voigt configuration, where the magnetic field is perpendicular to the growth direction of the sample (and the confinement direction of the quantum dots). To perform experiments which require Faraday configuration (magnetic field parallel with sample growth direction), we have created the adapter in Figure 2.5d). This adapter allows the sample to be loaded horizontally, but it is only designed to accommodate the smaller chip carriers. Additionally, with the B-field adapter in place, we cannot use the optical fibre for light delivery due to limited space in the fridge. As we can see in the figure, however, the use of the small chip carriers does allow additional components, such as an LED, to use the free socket pins.

³Unlike the lateral dots, however, the nanowire samples required no wire bonding as we were only performing optical experiments on them.




Figure 2.5: a) Sample chip carriers. Left: Top view of chip carrier, with hole for polarisation identification. Right: Socket and large/small chip carriers. b) Sample loading device. The four screws press the chip carrier into the socket on the PCB. c) Circuit board with loaded sample. d) B-field adapter for Faraday configuration. Right-hand image shows the adapter in use, with an LED connected to two of the available socket pins.

2.2.2 Lateral Quantum Dots for Photon-to-Spin Qubit Transfer

As we saw in Section 1.7, the photon-to-spin protocol uses a lateral quantum dot to trap a photo-generated carrier with a specific spin state. In working towards the creation of this system, we have tested several double quantum dot samples in the fridge. These are GaAs-AlGaAs devices grown by the Advanced Technology Fabrication Facility at NRC. An example of one of these samples is shown in Figure 2.6. Each sample has an optical

mask over the dots to prevent unwanted light from reaching the sample. A small hole (400 nm diameter) over one of the dots is designed to allow the polarised photons to reach that dot and transfer their qubit state to an electron spin.




Figure 2.6: Images of a good photon-to-spin DQD sample. a) Full sample imaged under a microscope prior to wire bonding. b) Zooming in on the centre of the sample, we can now see the optical mask (yellow square). The black spot in the centre of the optical mask is the hole over the left dot. c) A schematic of the sample gates, showing additional gates on either side of the DQD for charge detection. They may be used to form either a SET or a QPC. The circle over the left dot represents the size of the hole in the optical mask.




Figure 2.7: Images of a photon-to-spin DQD sample with broken, merged and missing gates. a) Full sample imaged under a microscope, with gates missing. Compare to Fig. 2.6a). b) An SEM image of the sample gates, which are broken and merged. Compare to Fig. 2.6c).

In preliminary tests of these samples, we were able to show that the gates could deplete the 2DEG and form quantum point contacts (QPCs). Beyond that, however, we began to experience several issues. In all of the samples we received, the gates were broken, merged, or missing, which prevented us from properly forming quantum dots. Figure 2.7 shows examples of samples with broken and merged gates. In some samples, the gates also tended to exhibit current leakage at voltages below -0.35 V, which prevented proper QPC measurements from occurring. These fabrication issues also extended to the ohmics, which had poor conductance, and to the optical mask, which did not have a properly-formed hole over the quantum dot.

While waiting for these fabrication difficulties to be resolved, we have tested other samples, such as nanowire quantum dots. We have also recently begun testing new photon-to-spin samples at the SPIN Lab in conjunction with the Oiwa Lab at Osaka University, who are working on similar research. While the results are still preliminary, these samples look promising so far, and have demonstrated the ability to form dots. As we shall see in Chapter 3, we have also managed to locate the optical mask holes in these samples using the lensed fibre. Our photon-to-spin research will therefore progress using these samples for the near future.

2.3 Preliminary Testing of the RF-QPC Charge Sensor

Experiments on quantum dot samples, especially lateral quantum dots, require some method of determining the electron (or hole) population within the dot. Mechanisms such as the single-electron transistor (SET) and QPC have been used for quantum dot charge detection in the past. Recently, however, these have been superseded by the RF-SET [30], the RF-QPC [33], and the dispersive gate sensor [58], which act as highly sensitive electrometers and use high frequency signals to resolve single-electron transitions at frequencies of greater than 1 MHz [29]. The lab’s cryogenic SPIN setup has therefore been equipped with a reflectometry circuit which can be used for fast RF-QPC charge detection.

While many of the details of an RF-QPC’s operation are covered in Section 1.3.2, in this section I will describe the full layout of the SPIN Lab’s charge detection system and some of our tests of and improvements to it. As discussed in the previous section, the lateral quantum dot samples we attempted to use during the course of my degree did not function properly, with most being unable to form dots at all. We were able to use these devices to form QPCs, however, so although we have yet to test the RF-QPC’s response on an actual quantum dot, we have been able to conduct several preliminary tests of the system’s response to a changing QPC resistance.

2.3.1 Layout of the RF Reflectometry Circuit

The SPIN Lab is equipped with an RF reflectometry circuit for fast charge detection of quantum dot samples in the dilution refrigerator. Although constructed at NRC, the SPIN Lab’s reflectometry circuit was designed with input from Professor Jan Kycia (University of Waterloo) and his former PhD student Jeffrey Mason (University of Waterloo). Figure 2.8a) is a diagram of the main components of the circuit, which are on the PCB board in the fridge. Outlined in red is the tank circuit, the purpose of which is to produce a reflected signal with an amplitude proportional to the charge state of the quantum dot sample.⁴ The tank circuit contains the QPC (which we treat as a resistor), a capacitor

⁴The operation of the tank circuit is discussed in more detail in Section 1.3.2.

(C), and an inductor (L). For DC measurements, a bias is placed across the QPC using the source and drain bias tees outlined by the green and pink boxes. The RF signal, V_{IN} , is fed into the circuit from the top, where it passes through the bias tee outlined in blue. The purpose of this tee is to allow a bias to be placed across the tank circuit capacitor, which, in our initial design, was a variable capacitor. By tuning the applied bias, one can control the capacitance of the variable capacitor. (See Section 2.3.3 for more information on our tests of the variable capacitor.)

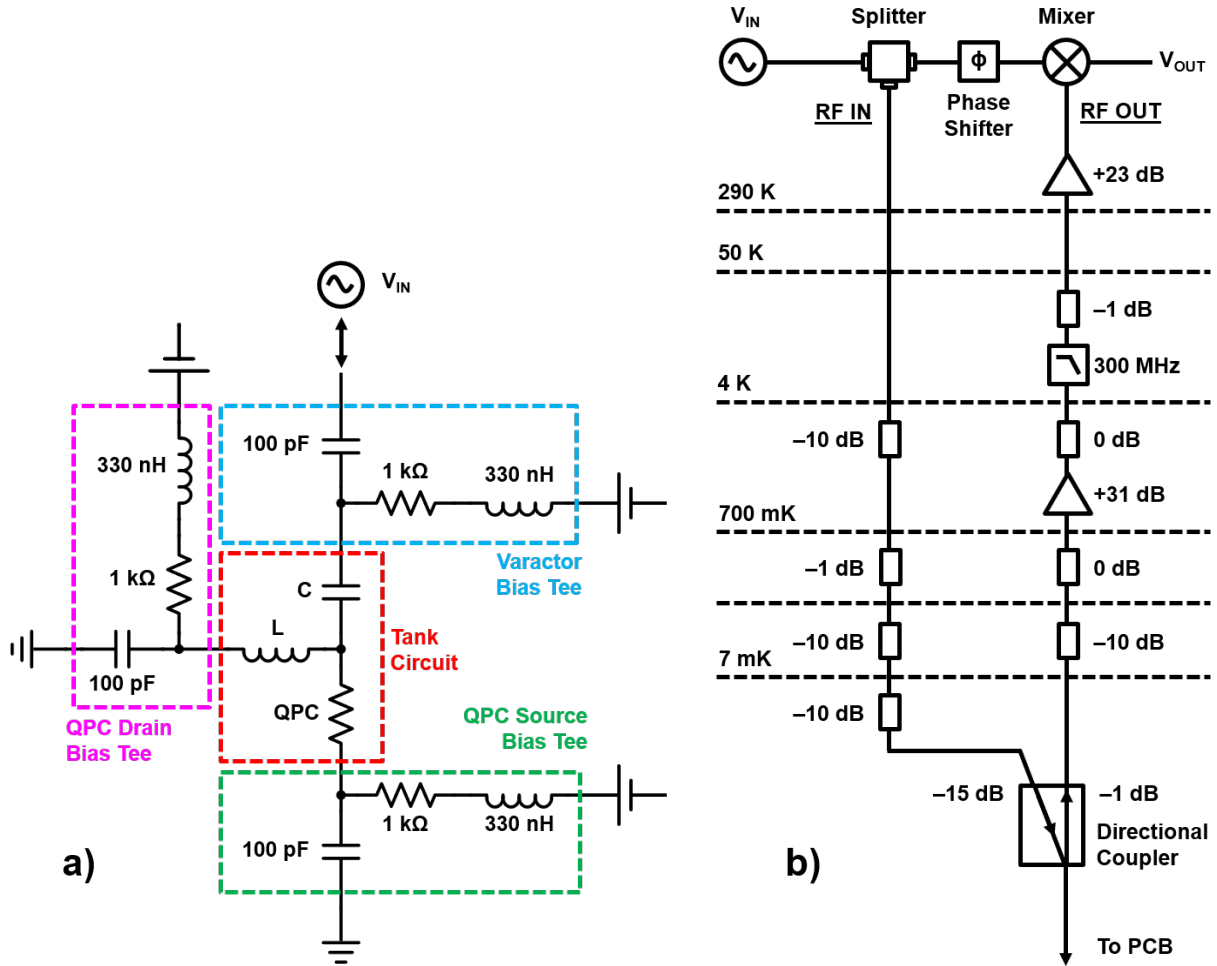


Figure 2.8: a) Schematic of the RF reflectometry circuit on the PCB, showing the tank circuit and bias tees. After simulating and testing the system, we chose a tank circuit capacitance of $C = 1.5$ pF and a tank circuit inductance of $L = 180$ nH. b) Schematic of the high frequency lines used for RF-QPC measurements in the fridge.

A major advantage of this setup, with the RF reflectometry circuit on a PCB, is the ease with which we can switch samples. Designs used by other groups include constructing the circuit directly on the sample chip [29] or integrating the circuit and qubit into a single

chip [59]. These designs likely have reduced signal loss, as there is less distance between the circuit and sample. It becomes much more difficult to change samples, however, as the reflectometry circuit must be reconstructed each time a new sample is tested, and new circuit components are required if the old ones were glued to the sample. The decision of whether to build the circuit separately or on-chip is principally determined by the frequency with which the sample is expected to be changed and the degree of signal loss that can be tolerated. With our system, we chose to have the reflectometry circuit separate to allow for quick swapping of the many samples we expected to be testing. This setup also allows us to consistently maintain the exact same circuit components, which reduces the number of issues that can appear each time a new sample is tested.

Two high-frequency, BeCu coaxial wires run from room temperature at the top of the fridge to the circuit board at the mixing chamber, with attenuators and amplifiers at various levels.⁵ These high-frequency lines, which are shown in Figure 2.8b), deliver the RF signal to the circuit board and receive the reflected portion of the signal. We can see at the top that an RF signal (with the tank circuit resonance frequency) is generated at room temperature, and a power splitter sends half of the signal into the fridge along the RF IN line. Attenuators reduce the signal amplitude so that we avoid damaging the sample and heating the fridge. As we see from the figure, there is currently 31 dB of attenuation on the RF IN line, although we have used up to 40 dB in the past. A directional coupler sends the incoming signal to the circuit board (Fig. 2.8a).⁶ As we saw in Section 1.3.2, the circuit may reflect some portion of the signal if its impedance does not match that of the RF lines, while the remainder of the signal is transmitted to ground. The factor which determines the circuit impedance is the QPC, the resistance of which is dependent upon the quantum dot population. The amplitude of the reflected signal therefore indicates the charge state of the quantum dot sample.

The reflected signal is directed up a second high-frequency line, which we label RF OUT. A 10 dB attenuator helps reduce the relative amplitudes of oscillations produced by standing waves. Several 0 dB ‘attenuators’ are present because they provide thermal coupling between the inner and outer components of the coaxial wires. Amplifiers are used to return the amplitude of the signal (when fully reflected) to its original power. The reason we split the incident and reflected signals with the direction coupler is therefore to allow us to reduce the amplitude of the signal before it reaches the sample and then amplify the reflected signal.

An electronic mixer combines the RF OUT signal with a phase-shifted version of the original to produce an output voltage, V_{OUT} . The amplitude of V_{OUT} is dependent upon the relative amplitude of the reflected signal; a smaller reflected amplitude produces a smaller V_{OUT} voltage, and vice versa. V_{OUT} therefore provides a measure of the tank circuit impedance and hence the QPC resistance. As the resonance frequency of our tank circuit is

⁵Two other high-frequency wires are also present, although they are used to apply microwaves to the sample.

⁶The directional coupler attenuates a 200 MHz signal travelling down RF IN by roughly 15 dB, but it only attenuates the returning signal by 1 dB as it begins travelling up the RF OUT line.

typically about 200 MHz, this system allows for fast detection of changes in the populations of dots which are adjacent to the QPC.

When performing tests of this system, we generally use a network analyser (Agilent, E5071C) instead of the four components at the top of the fridge (RF generator, splitter, phase shifter, and mixer). The network analyser performs much the same function as these components, but it continuously scans through a range of signal frequencies, which allows us to observe the behaviour of the tank circuit resonance under various conditions. An AC signal is output from Port 1 of the network analyser, it travels along RF IN to the PCB, and any reflected signal from RF OUT is measured at Port 2 of the analyser. Signals received at Port 2 are compared to those output at Port 1, and data is then plotted in terms of the S_{21} element of the scattering matrix. The S_{21} element is defined as follows:

$$S_{21} = \frac{V_2}{V_1} \quad (2.1)$$

where V_1 is the amplitude of the signal emitted at Port 1 and V_2 is the amplitude of the signal measured at Port 2. We are using a directional coupler to separate the reflected signal $V_2 = V_-$ from the emitted signal $V_1 = V_+$, but if we sent them along the same line, the Port 1 line for example, we would have:

$$S_{11} = \frac{V_1^{\text{reflected}}}{V_1^{\text{incident}}} = \frac{V_-}{V_+} = \Gamma \quad (2.2)$$

The scattering matrix element S_{11} is therefore the reflection coefficient Γ , which we defined in Section 1.3.3. In reality, of course, the directional coupler splits the incident and reflected signal, so we equate the S_{21} element with the reflection coefficient in our system. We typically plot the reflection coefficient in decibels ($S_{21}[\text{dB}] = 20 \log(S_{21})$) as a function of the RF signal's frequency.⁷ The power of the signal output at Port 1, P_1 , is usually between 0 dBm and -20 dBm, with an additional 46 dB of attenuation by the time it reaches the tank circuit. The signal power is related to the reflection coefficient by the equation:

$$\frac{P_2}{P_1} = \left[\frac{V_2}{V_1} \right]^2 = |\Gamma|^2 \quad (2.3)$$

We can therefore calculate the power of the reflected signal, P_2 , using the known input power and the measured reflection coefficient. Overall, the network analyser is very useful for studying the behaviour of the system in preliminary tests, as we can watch the frequency

⁷Keep in mind, however, that the data in these plots is generated after the signal has passed through multiple attenuators on the RF IN line and multiple amplifiers on the RF OUT line. The plots, therefore, are not showing the exact magnitude of the reflection coefficient at the tank circuit, but rather the ratio of the voltages measured/output by the network analyser at the top of the fridge (after the attenuators and amplifiers). The tank circuit impedance will determine the overall shape of S_{21} as a function of frequency, however, and we are currently only concerned with generating a sharp, deep resonance.

and depth of the tank circuit resonance simultaneously. Once we begin performing quantum dot experiments, however, we will only require a signal at the resonance frequency to measure the changes in V_{out} .

2.3.2 Computer Modelling of the RF Reflectometry Circuit

Our goal in designing the reflectometry circuit was to use components which, in conjunction with typical QPC resistances, could achieve a 200 MHz resonance frequency and a 50Ω circuit impedance. This impedance was of course selected to ensure impedance-matching with the transmission lines.⁸ The resonance frequency of 200 MHz was chosen for several reasons. Larger frequencies (e.g. 1GHz) require smaller electrical components, which can be difficult to manufacture properly and lead to complications. Parasitic capacitance also tends to play a larger role in the behaviour of the circuit at high frequencies. Lower resonance frequencies (<100 MHz) can limit the bandwidth, resulting in a slower RF-QPC readout.

Prior to purchasing the components for the reflectometry circuit, we decided to model it to determine which elements would produce the desired results. While early-stage modelling can be done theoretically, the actual system is more complicated than the simple tank circuit we looked at in Section 1.3.3. In order to account for the effects of the bias tees, the system was modelled using the circuit simulation software PSpice Lite by OrCAD Systems. Using this software, I recreated the circuit shown previously in Figure 2.8a). By applying RF signals of different frequencies and measuring transmission through the circuit, we could observe the effects of the various components as their strengths were altered.

The main circuit elements we were looking at were the tank circuit capacitor (C), the tank circuit inductor (L), and the QPC, which we modelled as a resistor (R).⁹ The data presented in Figure 2.9a) was obtained by observing the transmitted signal amplitude versus signal frequency. The fact that we measured the transmitted portion of the signal explains why we get peaks in these plots, rather than the valleys which we usually observe for the reflected component. This test was done for a range of different tank circuit capacitance values (from 0.3 pF to 1.5 pF), and three values of the tank circuit inductor (800 nH, 1100 nH and 1350 nH). The QPC resistance was held constant at 25 k Ω . We can clearly see that as the tank circuit capacitance decreases, the resonance frequency increases. The same can be said for the inductance, which causes a decrease in resonance frequency as it increases. This is exactly the type of behaviour predicted by our calculations in Section 1.3.3. The capacitance value which is closest to 200 MHz has been outlined in orange for all plots in Figure 2.9.

In Figure 2.9b), we are again changing L and C, but now we are looking at the effect of the QPC resistance as well. We can see that as we increase R from 25 k Ω to 100 k Ω , we get much larger peaks. This improvement in transmission suggests that for these values of

⁸See Section 1.3.3 for more information on impedance matching.

⁹We also examined the bias tee components, to ensure that they behaved as expected.

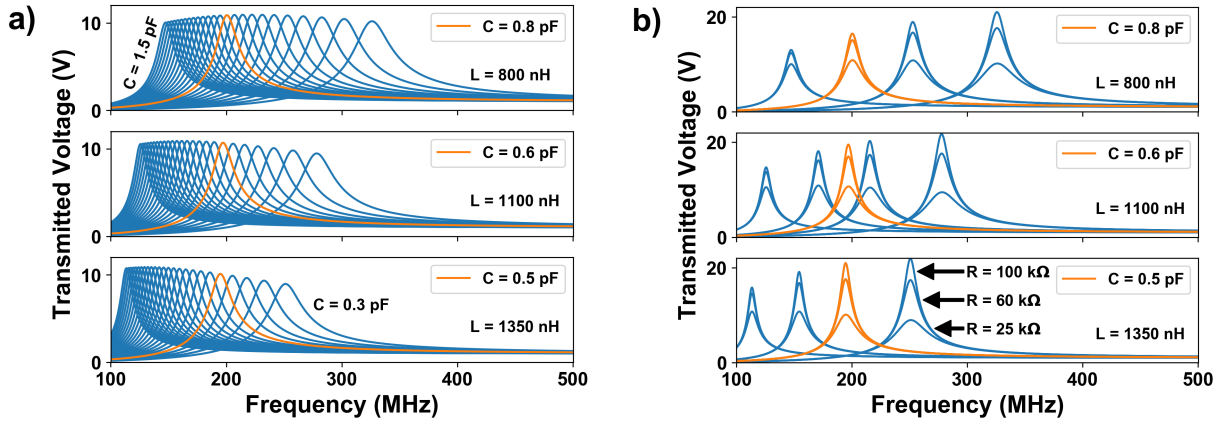


Figure 2.9: Frequency response of the system for different values of the tank circuit inductor (L), capacitor (C), and QPC (R). We have highlighted in orange the result which yields a resonance frequency of 200 MHz. In a), we are changing just C and L, while keeping $R = 25 \text{ k}\Omega$. The value of C ranges from 0.3 pF (right-most peak in each plot) to 1.5 pF (left-most peak in each plot). In b), we are now changing the QPC resistance as well, setting it to values of 25 kΩ, 60 kΩ, and 100 kΩ. The same capacitance range is used, but with fewer points.

C and L, better impedance matching is achieved for higher QPC resistances. As we found in Section 1.3.3, changing R affects only the amplitude of the resonance; it has little or no effect on the resonance frequency.

With these simulations, we were able to show that we can alter the resonance frequency using the tank circuit capacitor and inductor, while the QPC resistance controls the resonance amplitude. Based on these results, we decided to choose a tank circuit inductor with an inductance of 180 nH. A variable capacitor covering a capacitance range of about 0.4 pF to 6pF was then installed in the circuit so that we could tune the resonance frequency, while a QPC already allows manual adjustment of the resonance amplitude by altering the gate voltages. As we shall see in the following sections, this combination enabled us to achieve relatively deep resonances at frequencies near 200 MHz.¹⁰

Finally, one additional use for these simulations was that they allowed us to calculate the parasitic capacitance to ground, C_p , by adding to the simulated circuit a capacitor to ground in parallel with the QPC and inductor. Since the parasitic capacitance has a large effect on the resonance frequency (see Section 1.3.3), we compared the resonance frequency of the real circuit to that of the simulated circuit for a range of simulated C_p values. The point at which the real and simulated resonance frequencies match gives us the actual C_p value. The parasitic capacitance is sample-dependent, but for one of our photon-to-spin samples, we were able to use this technique to show that the parasitic capacitance was

¹⁰In the following section, I will describe the behaviour of the variable capacitor at low temperatures and our subsequent decision to switch to a fixed, 1.5 pF capacitor.

roughly 2.4 pF. In general, the parasitic capacitance of our system seems to be about 2–3 pF.

2.3.3 In-Situ Tuning of the Resonance Frequency Using a Variable Capacitor

Initial tests of the RF-QPC were conducted with a varactor diode (M/A-COM Technology Solutions, MA46473-186) as the tank circuit capacitor. A varactor diode (or just varactor, for short) consists of a reverse-biased p-n junction which acts as a variable capacitor. With the capacitance determined by the applied DC bias, it has been shown that these devices can be used to tune the resonance frequency of an RF reflectometry circuit [60]. Such an ability would enable us to compensate for the effects of parasitic capacitance and achieve our desired resonance frequency of 200 MHz.

Our initial RF-QPC tests involving the varactor diode were promising. Tests conducted at room temperature and low temperatures demonstrated our ability to tune the resonance frequency of the circuit over a total range of about 50-60 MHz. This range of resonance frequencies, extending from 140 MHz to 200 MHz, was achieved by varying the reverse bias on the varactor from 0 to 12 V, which corresponds to a (non-linear) change in the capacitance from 6 pF to 0.4 pF [61]. These tests were mainly conducted using our photon-to-spin lateral dot samples (see Section 2.2).¹¹ Figure 2.10a) shows the frequency response of the system as a function of the varactor’s reverse bias, with decreasing varactor capacitance corresponding to an increase in resonance frequency. This is in agreement with what we found in our simulations of the circuit.¹² The change in capacitance also caused a slight change in the depth of the resonance, allowing us to choose a capacitance which produced a deep resonance near 200 MHz.

As we began testing the RF-QPC at low temperatures, we immediately noticed that the resonance would disappear as the temperature dropped below 10 K. Figure 2.10b) shows an example of the resonance vanishing as a function of temperature, an event which occurred consistently across multiple cool-down attempts. After further testing, research, and consultation with researchers at other facilities, it was determined that the varactor was experiencing a ‘carrier freeze-out’ effect. A varactor operates through the reverse-biasing of a p-n junction. As the temperature approaches absolute zero, however, conduction band electrons drop to their lowest energy state and bind to the semiconductor dopants, preventing proper operation of the device [62].

In order to resolve this freeze-out issue, a total of four different varactors were tested at low temperatures. In addition to the MA46473-186, we also looked at the MGV100-26-P55 and MGV100-27-E28 from MACOM, and an Aeroflex A17041. Each varactor was installed in turn on a replica tank circuit mounted on a dipstick, and the system’s frequency

¹¹Although we could not properly form quantum dots with them, we could still form QPCs.

¹²See Sections 1.3.3 and 2.3.2.

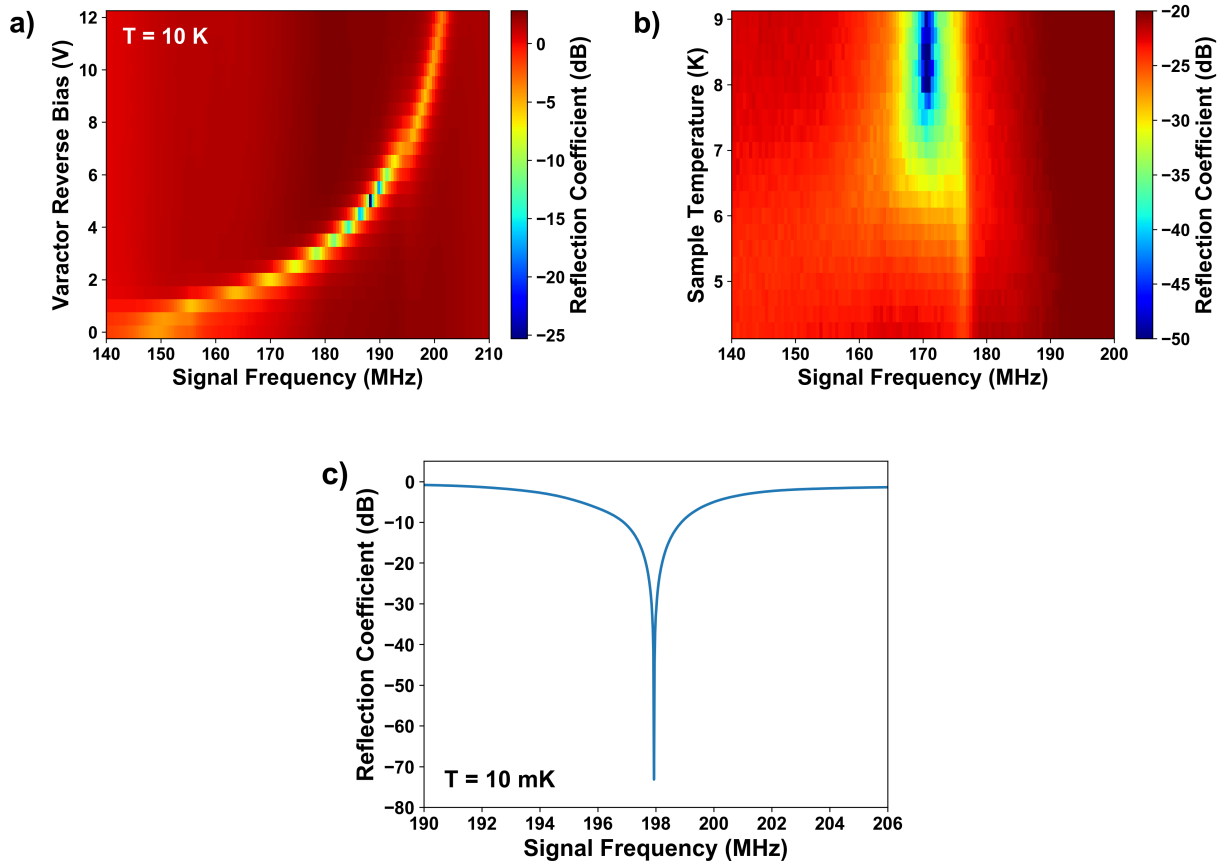


Figure 2.10: a) Tuning the tank circuit resonance frequency using the variable capacitor. Capacitance decreases from 6 pF to ~ 0.4 pF as reverse bias is ramped from 0 V to 12 V. The temperature is 10 K. b) With a varactor (MA46473-186) as the tank circuit capacitor, the RF-QPC resonance reliably disappears at low temperatures. c) With a fixed, 1.5 pF capacitor in the tank circuit, we can generate deep and narrow resonances near 200 MHz while operating at base temperature.

response was observed both at room temperature and with the circuit dipped in a liquid helium dewar (4 K). An infinite resistor (i.e. open circuit) was used to replicate the effects of a completely pinched-off QPC, which should produce a decent resonance when the circuit is operating properly. Our goal in these tests was therefore to determine if any of these devices would produce a deep resonance for a low signal power at a temperature of 4 K. The MG100-26-P55 was the only varactor which produced a reasonably deep resonance (20 dB) at both room temperature and 4 K, but even it was unable to maintain this performance when tested in the fridge at milli-Kelvin temperatures.

After failing to successfully find a suitable varactor, we resorted to installing a low-temperature 1.5 pF capacitor in the tank circuit instead. We have continued to use this capacitor ever since, as it is able to achieve a much deeper resonance than any produced

by the varactors, both at room temperature and milli-Kelvin temperatures. It also yields a resonance frequency which is very close to 200 MHz for almost every sample, which is fortunate because we can no longer tune the resonance frequency. Figure 2.10e) shows an example of the type of frequency response we observe at base temperature (10 mK) with this fixed capacitor.

While we do appear to have found an ideal capacitor, one final varactor may be tested in the future. Obtained after the installation of the 1.5 pF capacitor, this new varactor is maximally-doped, meaning that it has overlapping valence and conduction bands [62]. This results in free conduction band electrons always being present, even at very low temperatures. This device might therefore allow us to tune the resonance frequency with the fridge at base temperature.

2.3.4 Testing the RF-QPC System on QPC Samples

Due to a persistent lack of functioning quantum dot samples, we have yet to test the RF-QPC on an actual dot and demonstrate its functionality as a fast charge detector. We can form QPCs with our samples, however, so we have been able to show that the RF-QPC responds appropriately to a (manual) change in QPC resistance. As we await the arrival of new samples, this has proven to be a good preliminary test of the system.

With a QPC established,¹³ we are able to verify that the resonance depth was dependent upon the QPC resistance, as expected from the model. Figure 2.11a) shows the RF-QPC frequency response as a function of the voltage V_g applied to one of the QPC gates. As the gate voltage decreases, the QPC narrows and its resistance increases. Figure 2.11b) shows the corresponding QPC resistance over this range of gate voltages, which was obtained by applying a 2 mV bias across the QPC and monitoring the DC current passing through it.

Figure 2.11c) shows the three traces marked by the white lines in a). We can see from these plots that as the resistance is manually increased, the RF-QPC resonance becomes narrower and deeper, until it passes through a critical point at a QPC resistance of $R = 145$ k Ω , whereupon it becomes slightly shallower. This critical point occurs when the impedance of the reflectometry circuit most closely matches the 50 Ω transmission line impedance. At that point, most of the RF signal is transmitted to ground. The deepest point of the resonance is -63 dB (at this resolution, anyway). As the system's noise level is roughly -80 dB at the level of the network analyser (possibly due to the amplifiers in the RF OUT line), we have not achieved perfect impedance matching, and we should attempt to produce a resonance which extends down to that level. Note that the temperature also affects the resonance depth (and width), with lower temperatures producing deeper, narrower resonances. The data in Figure 2.11 was taken at a temperature of 4 K. (Since we were unable to form a dot with this sample, we did not perform any tests at base temperature.)

¹³In this case, we were using a triple lateral quantum dot sample instead of a photon-to-spin DQD sample. Due to damaged gates, we were also unable to form dots with this new sample.




Figure 2.11: a) Frequency response of the RF-QPC for a manually-changing QPC resistance. The white lines and associated QPC resistances (R) indicate the traces shown in c). b) QPC resistance as a function of gate voltage, with the resistance values for the traces in c) marked by orange dots. c) Three traces from a) showing the evolution of the resonance as the QPC resistance increases.

As we were led to expect by the simulations in Section 2.3.2, the resonance frequency does not appear to depend significantly on the QPC resistance. When the resonance is deepest, with a depth of 53 dB, the resonance frequency was 197.4 MHz, just as we intended. (Recall that our resonance frequency goal was 200 MHz.)

In addition to the effect of the QPC resistance on the system's frequency response, we can also look at the effect of the signal power. Looking again at the frequency response as a function of QPC gate voltage, Figure 2.12 shows the behaviour of the system for three different powers of the applied signal at the tank circuit: -55 dBm, -65 dBm, and -75 dBm. This data was taken with the same triple-dot sample as in Figure 2.11. Although the resonance depth and frequency remain largely unchanged, we can see that as the power decreases, there is a narrowing in the range of QPC resistances which produce a significant

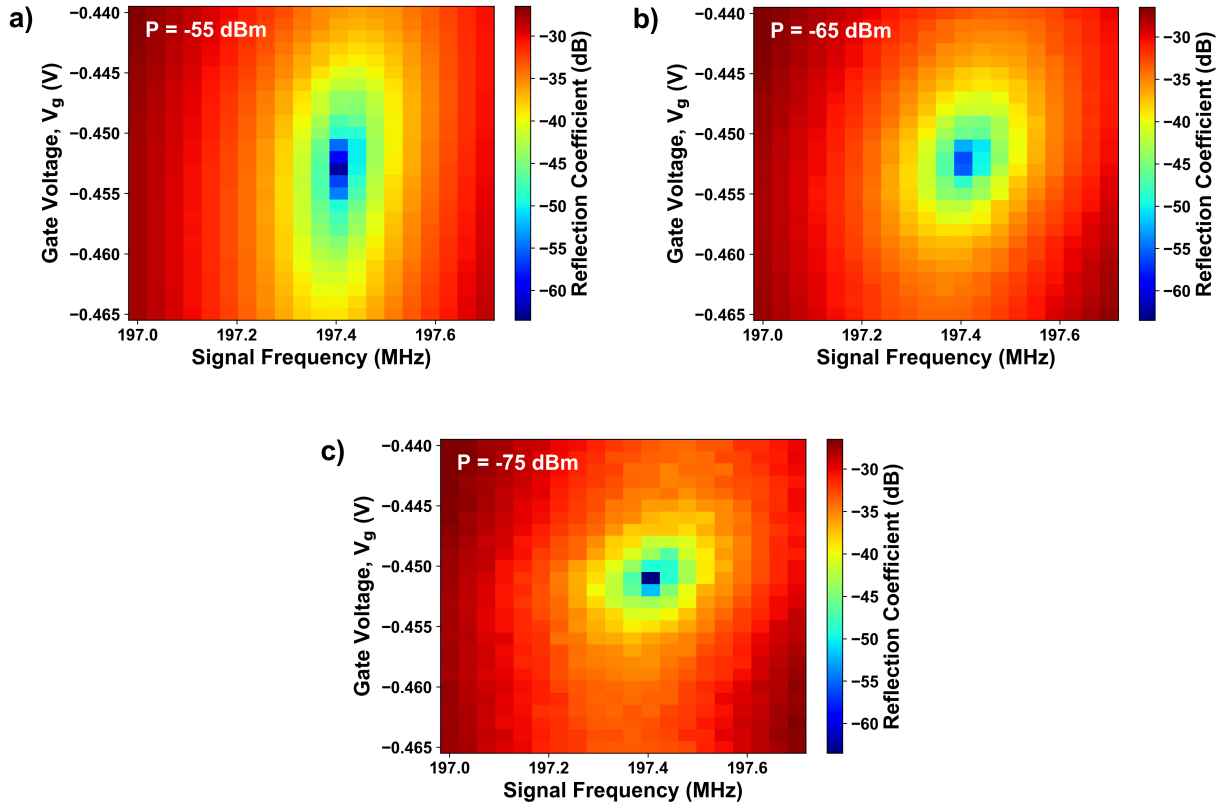


Figure 2.12: The frequency response of the RF-QPC as a function of QPC gate voltage for three different RF signal powers: a) -55 dBm, b) -65 dBm, and c) -75 dBm.

resonance. This is important because it means that at lower powers, a smaller change in QPC resistance will yield a large change in the amplitude of the reflected signal. We can therefore use a lower signal power to achieve a more sensitive charge sensor or a higher power to look over a larger range of QPC states. Also, note that these plots highlight the *very* slight dependence of the resonance frequency on resistance, which is something we expected to see based on the calculations in Section 1.3.3.¹⁴

Putting all these results together, we have recently tested the RF-QPC of a photon-to-spin DQD sample provided by the Oiwa Lab of Osaka University, and the result was a very nice resonance signal. The data in Figure 2.13 was taken at a mixing chamber temperature of 10 mK. Again, changing the QPC resistance manually, we found the point at which the resonance was deepest. This yielded a resonance depth of 72 dB, which begins to approach the noise limit of our system (-80 dB). The incident RF signal had a power of -75 dBm at the tank circuit to improve sensitivity.

Recall from Section 1.3.2 that a preliminary method of determining RF-QPC sensitivity

¹⁴Although we made an assumption in Equation 1.15 which eliminated the QPC resistance from the final resonance frequency equation, it should in reality contribute ever so slightly.

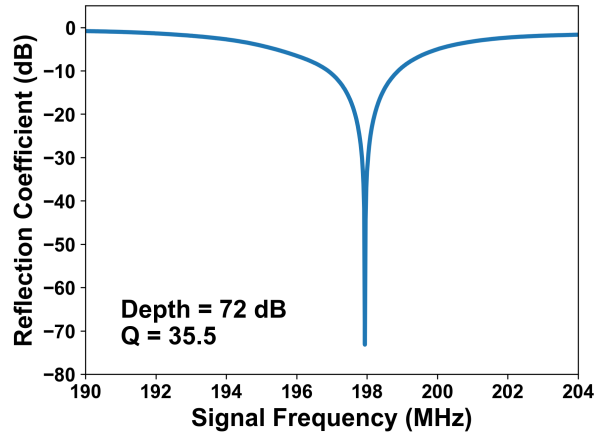


Figure 2.13: Frequency response for a photon-to-spin double-dot sample at a temperature of 10 mK and incident signal power of -75 dBm. The resonance depth is 72 dB and the quality factor is 35.3.

is the quality factor:

$$Q = \frac{f_0}{\text{FWHM}} \quad (2.4)$$

where f_0 is the resonance frequency and FWHM is the full width at half maximum. From this, we calculate that our RF-QPC system currently has a quality factor of 35.5. While still not quite as high as we would like it, this is the highest value we have yet been able to achieve, indicating that we are slowly improving the RF-QPC sensitivity over time.

As we move forward, the RF-QPC will become a major component of our system, enabling us to achieve fast and highly sensitive detection of carrier photo-generation in photon-to-spin quantum dots. The design of this system also makes it very versatile. Proper wire bonding of the sample would allow it to function as a dispersive gate sensor for charge detection [58], or as an RF-SET, something which we briefly attempted several months ago.¹⁵ Once we are able to test the circuit on functioning quantum dot samples, we should be able to determine its charge sensitivity (Equation 1.10), but in the meantime, we have at the very least managed to demonstrate that our RF-QPC is operational and ready for new samples.

¹⁵The RF-SET tests were promising, but we were unable to form quantum dots with those samples.

2.4 Optical Setup for Fibre-Based Light Delivery and Polarisation Identification

One of the main features of the SPIN lab is the fibre-based delivery of light to the mixing chamber for optical measurements of nanostructures. Of particular importance to the photon-to-spin protocol is ensuring that polarised light delivered through a standard optical fibre has the desired polarisation state upon arriving in the fridge. To that end, we have invented a system comprised of standard room-temperature and cryogenic optical components to compensate for the fibre's effects on polarisation.




Figure 2.14: A schematic of the major components in the setup used to deliver polarised light to the dilution refrigerator by fibre, including the three polarisation identification components in the fridge. The curved arrows on the waveplates represent the fact that they are mounted in rotators, while the horizontal arrows on the linear polarisers indicate their polarisation direction (i.e. horizontal).

Figure 2.14 is a schematic of the optical setup used for polarisation control in a fibre. All components within the green box are part of the cryogenic SPIN setup in the fridge, while the components in front of the 30 m fibre are located in the optics lab across the hall from the SPIN lab. The 30 m fibre in the figure is actually a series of fibres, including a long fibre stretching between the two labs, a second fibre which extends from room temperature at the top of the fridge to 10 mK at the mixing chamber, and the lensed fibre which focuses the light on the sample.

In this section, we will discuss the optical components of the SPIN lab, which were used to accomplish the experiments described in the remaining chapters. We will begin by looking at the room-temperature setup, which is used to control the polarisation of the light reaching the fridge. In Section 2.4.2, we will take a look at the cryogenic optical components mounted on the cryogenic SPIN setup. Their purpose is to identify the polarisation of light reaching the sample in the fridge.

2.4.1 Room-Temperature Optical Setup

Preliminary preparation of the light is performed at room temperature in our optics lab. The room-temperature setup is shown schematically on the left side of Figure 2.14, and a picture of it is presented in Figure 2.15. The laser source used in our experiments is a tunable Ti:sapphire (Spectra-Physics, Tsunami model), which is usually tuned to about 780–800 nm. The beam is directed through a linear polariser (Thorlabs, LPNIR050) to produce pure linear polarisation. We label this polariser P1 to distinguish it from the polariser in the fridge. A rotating half-wave plate (labelled HWP*) in front of the polariser controls the amount of light which passes through it.¹⁶ The polariser is followed by another half-wave plate (Thorlabs, AHWP05M-980) and a quarter-wave plate (Thorlabs, AQWP05M-980), both mounted in computer-controlled mechanical rotators. The combination of these two elements, which we label HWP and QWP1 respectively, allows for the generation of any polarisation state at the input of the fibre. A collimator (Thorlabs, TC06APC-780) couples the free space beam into the single-mode, 780HP fibre. A 50:50 fibre beam splitter (Oz Optics, Part #: FUSED-12-850-5/125-50/50-3A3A3A-3-1) and photodiode (Thorlabs, S130C) are used to monitor the intensity of light entering the fibre so that fluctuations may be filtered out of the data during post-processing. We label this photodiode PD1. The light is then transmitted from our optics lab to the refrigerator in our cryogenics lab by a 30-metre-long, 780HP, single mode fibre.

As we will see in Chapter 5, the beam splitter and PD1 are critical to our ability to eliminate noise from polarisation data, as most of the noise occurs in the room temperature setup prior to the fibre. Some of this noise may simply be fluctuations in the laser power, an issue which we have experienced with several of our laser sources. Additionally, the rotating waveplates may cause the laser beam to wobble, which changes the amount of light coupling into the collimator. Careful alignment of the waveplates (ensuring that no wobble occurs as they rotate) can be time-consuming, but it is critical to the success of our experiments. Fortunately, the effects of any residual beam wobble can be eliminated by using PD1 to monitor the amount of light entering the fibre. Polarisation noise arising after the collimator may occur due to vibrations in the fibres, temperature fluctuations in the fridge, or dirty fibre tips. Dirty fibres can also cause extreme intensity fluctuations over time, so the fibre tips are always cleaned before use.

Variations on the room-temperature setup are occasionally used, depending on the type of experiment we wish to perform. When testing nanowire quantum dots, for example, we were not concerned with polarisation, and we therefore did not use any waveplates. The optical setup used for the nanowire experiments will be described in more detail in Chapter 4.

¹⁶HWP* is not included in Fig. 2.14 as it is not one of the main components; its purpose is merely to control the amount of laser light used in the experiment.




Figure 2.15: The room-temperature optical setup used for delivery of polarised light. All components from Figure 2.14 are labelled, and the free-space path of the laser is depicted by the red lines. The half-wave plate labelled HWP* is used to align the initial polarisation with the orientation of P1, to maximize throughput.

2.4.2 Cryogenic Optical Setup

Polarized light is transmitted from the top of the refrigerator at room temperature to the MC flange by a single mode, 750 nm fibre. A lensed fibre (Adamant, SLF R4 SMF FC/APC) delivers the light from the MC flange to the nanostructure sample mounted on the PCB. The total length of the fridge fibres is ~ 2.5 m. The lensed fibre has an expected focal length of $8 \mu\text{m}$ and a calculated spot size which is roughly $0.75 \mu\text{m}$ in diameter.

The three piezoelectric positioners on the cryogenic SPIN setup provide sub-micron positioning of the lensed fibre in the x , y and z directions. The positioners are used both to focus light on our samples and to direct the beam through a hole in the PCB. Mounted behind the circuit board is a compact optical setup for polarisation identification. It consists of a zero-order quarter-wave plate (Thorlabs, WPQ05M-808), a linear polariser (Thorlabs, LPNIR100), and a photodiode (Thorlabs, FDS1010). The quarter-wave plate, which we label QWP2, is mounted in a piezoelectric rotator (Attocube, ANR240/RES). As the photodiode, labelled PD2, responds to incident light even without a reverse bias applied, the bias is set to 0 V throughout our experiments, which results in less Joule heating. It maintains a linear relationship between the power of the incident light and the resulting photocurrent for power levels ranging from 0 to $\sim 30 \mu\text{W}$. As we shall see in Chapter 5, this system allows us to monitor the intensity of the light passing through the polariser,

P2, as a function of QWP2 angle, which produces a unique intensity signature for each polarisation state.

From the schematic in Figure 2.14, we can see that the beam emitted by the lensed fibre is not collimated, but diverges before it reaches the polarisation identification components. (This is because its focal length is so small.) We initially used a gradient-index (GRIN) lens in the PCB hole to improve the collimation of the beam, but as we will discuss in Section 5.2.2, it also had the effect of partially depolarising the beam. We therefore removed the GRIN lens and use just the lensed fibre now, in spite of the beam divergence it produces. Using the fibre alone, we are still able to collect almost all of the light, and we do not observe any noticeable scattering during our polarisation measurements.

2.5 Conclusions

With the unique and varied capabilities of the cryogenic setup's combined optical and electrical components, the SPIN Lab's systems are extremely versatile. As we will demonstrate in the following chapters, this setup is capable of performing a great variety of measurements on different types of nanostructures at low temperatures and high magnetic fields. Our fibre-based system of polarisation delivery will enable us to test the potential of the photon-to-spin hybrid quantum repeater. Using this system, experiments conducted on nanowire quantum dots have already demonstrated that they are promising candidates for fibre-coupled sources of single photons. As we continue to use the SPIN Lab's experimental apparatus, we will work on improving its current design and adding additional capabilities.

Chapter 3

Distributed Imaging of Nanostructures

The optical experiments performed at the SPIN lab are made possible through the use of a single-mode, lensed fibre, which delivers light to nanostructure samples in the dilution refrigerator. Three piezoelectric positioning stages from Attocube Systems (two ANPx101/RES and an ANPz101/RES) provide sub-micron positioning of the fibre in the x, y and z directions, so light can be focused at specific locations on the sample with high precision. Figure 3.1 shows a simple diagram of the lensed fibre focusing light on a lateral DQD sample.




Figure 3.1: Schematic of lensed fibre delivering light to a quantum dot sample. (Fibre not to scale with sample.)

With the system enclosed in a dilution refrigerator, focusing the fibre at the appropriate point on the sample becomes difficult. The fridge has no installed camera to provide direct

observation of the sample, and there are also no windows, as they could cause unwanted heating and background illumination. The experimentalist must therefore attempt to align the fibre while virtually blind; the only available information about the fibre position is the electrical response of the sample and any light re-captured by the fibre lens in the fridge.

To overcome this challenge, a high-resolution imaging technique has been developed. The positioners scan the fibre over the sample in a grid pattern, while low-intensity light is directed at the sample. Topographical information about the sample is collected as a function of fibre position using one of three methods: photocurrent generated in the sample, light reflected off the surface, or photoluminescence¹ (PL) of the semiconductor material. With this process, we are able to map out different types of samples and properly focus the light on the nanostructures.

This chapter provides an overview of our imaging system and the methods by which it was developed. The details of operating the linear positioners are discussed, and the three methods of gathering topographical information are summarised. This imaging system is a vital component of our experimental capabilities, and the nanowire experiments described in Chapter 4 would not have been possible had we been unable to develop these techniques. The success of the photon-to-spin project will also rely upon our ability to map out lateral quantum dot samples and accurately deliver single photons to the hole in the optical mask (see Section 2.2.2 for descriptions of photon-to-spin samples).

The cryogenic positioning setup described in this chapter was designed and constructed by Dr. Louis Gaudreau and Piotr Zawadzki, with some minor assistance from myself. The imaging and positioner control software/procedure was largely developed by me, with suggestions and assistance from Dr. Gaudreau.

3.1 Operating Principles of Piezoelectric Nanopositioners

The imaging system we have developed for our lab depends entirely upon three piezoelectric nanopositioners, with which we are able to move a lensed fibre over the sample. As many of the challenges we encountered in developing our imaging techniques arose due to the very nature of the nanopositioners themselves, we will begin this chapter with a review of the basic operating principle of these devices.

Piezoelectricity is a property of certain solids where the application of mechanical stress, such as compression or stretching, generates an accumulation of electric charge. The result is a measurable electric potential across the solid, with the strength of the potential in direct proportion to the degree of applied stress. The inverse process is also possible. Piezoelectric materials, examples of which include various types of crystals and ceramics, respond to

¹Photoluminescence is the emission of light by an object after optical stimulation, such as exposure to a laser of sufficiently high energy.

the external application of an electric potential by undergoing a physical deformation, a property known as reverse piezoelectricity. The inherent electrical-mechanical relationship within these materials has resulted in the use of piezoelectric solids for a wide variety of applications, from pressure sensors in nuclear warheads to diaphragms in loudspeakers.

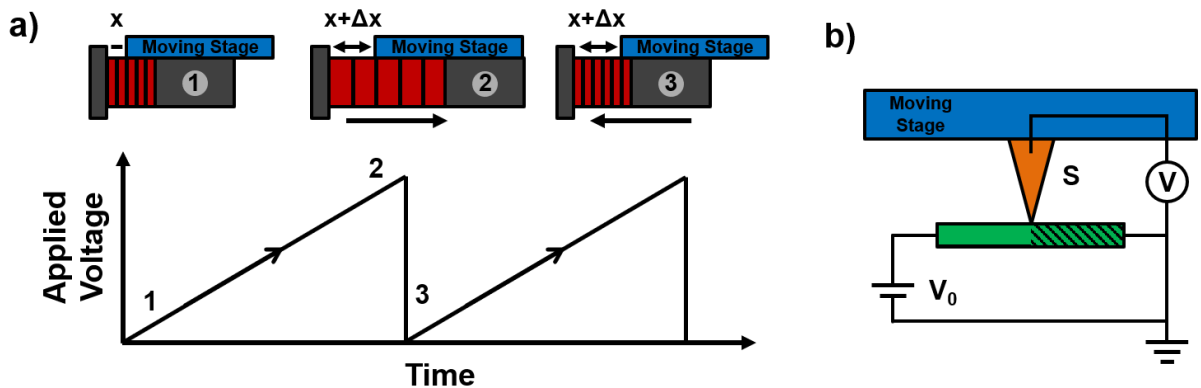


Figure 3.2: a) Slip-stick process of a nanopositioning stage. In steps 1 and 2, the piezostack (red) expands by an amount Δx as the applied voltage is ramped. When the voltage suddenly drops in step 3, the crystals rapidly contract to their original shape. Step 3 is too fast for friction to pull the stage (blue) back, so it remains in its new position, $x + \Delta x$. The stage has therefore moved one step forwards. b) Resistive position encoding system. A voltage V_0 is applied across the resistive strip (green). The probe (orange) attached to the moving stage (blue) measures a voltage S , represented by the hashed lines on the resistive strip. The value of the measured voltage S varies linearly as the stage and probe move.

Piezoelectric crystals are commonly used as mechanical actuators in devices requiring small, highly precise movements, such as nanopositioning systems. A single crystal will expand only slightly, even when over 100 V is applied to it, so stacks of crystals (called piezostacks) and mechanical amplifiers are often used in nanopositioners to increase the range of motion. Many piezoelectric positioning devices, such as those produced by Attocube Systems, operate on a ‘slip-stick’ principle [63]. We can understand this type of operation through a simplified description, as shown in Figure 3.2a). In this (highly-simplified) example, the moving stage sits atop a piezostack, and a sawtooth voltage is applied to the stack. At the start of the sawtooth waveform, the slow ramp in voltage expands the crystals, and friction moves the stage with them. This is the ‘stick’ phase. The sudden drop in voltage at the end of the waveform contracts the crystals, pulling the piezostack out from under the stage. (The ‘slip’ phase.) The stack moves fast enough in this step that the stage’s inertia prevents it from being pulled back to its original position by friction. The application of a single cycle of the sawtooth waveform therefore corresponds to a single slip-stick step in the forwards direction. Reverse motion is achieved simply by reversing the direction of the sawtooth wave, so the stack first expands quickly under the stage, then drags it backwards during the slow contraction phase.

The amplitude of the applied voltage determines the degree of piezostack expansion and therefore the size of the step taken by the stage. The sawtooth function’s frequency determines the speed of travel. For our experiments, typical values of these parameters range from 300 Hz to 1000 Hz in stepping frequency and 24 V to 28 V for the step amplitude, with a maximum allowed amplitude of 60 V. The full range of motion for our positioners is 5 mm, with the smallest step size being about 10 nm at 4 Kelvin [64, 65].

These positioners also have a ‘DC mode’ or ‘fine positioning mode’, where a DC voltage is applied and causes the stack to expand to a certain size. There is no actual ‘step’ of the positioner in this mode because the voltage is only ramped slowly, unlike the sawtooth waveform used during slip-stick movement (also called ‘AC’ or ‘coarse positioning’ mode). DC mode allows for sub-nanometer displacement of the stage [64, 65], and it is therefore used for our high-resolution imaging procedures. Maximal expansion of the piezostack, which occurs at an applied DC voltage of 150 V, will move the stage by 800 nm when at a temperature of 4 K. The fine and coarse positioning modes can be used simultaneously, provided the crystals are not already fully expanded by the DC voltage when attempting to step the positioner.

Our nanopositioners are equipped with resistive position encoders to allow the stage position to be determined. Figure 3.2b) shows a schematic diagram of the resistive encoder system. Each positioner contains a single resistive strip which stretches along its length. A voltage V_0 is applied across the resistor, and a probe attached to the bottom of the stage touches the resistor. Since it is in contact with the resistive strip, the probe is at a voltage S , where $0 \leq S \leq V_0$. As the stage moves, the probe is dragged along the resistor, and the changing value of S provides a measure of the stage location.²

Finally, we should note the temperature-dependent behaviour of the positioners. Piezoelectricity relates electric field strength and charge density to the strain and stress experienced by a material. As the temperature decreases, so too do the piezoelectric coefficients which relate these properties [63]. The result is a diminished voltage response to the application of stress, and a diminished deformation response to an applied voltage. Appendix D will cover additional details about our positioning system, including the reduction in positioner step size that we observe at lower temperatures.

²To understand this device, consider the operation of a voltage divider. The output voltage is determined by the input voltage and the resistances of the two resistors. With the resistive encoder, the total resistance on either side of the probe S changes as the probe moves along the resistor. This is analogous to the resistance of the two resistors in a voltage divider changing, with the resistance of one increasing while the other resistance decreases. In that case, the output voltage of the voltage divider would also change according to the change in the resistors. In the same way, we are able to determine the position of the moving stage based on the measured voltage S .

3.2 The SPIN Lab's Piezoelectric Nanopositioners

In the SPIN Lab, three linear, piezoelectric nanopositioners are used to move a lensed fibre over a nanostructure sample, as shown in Figure 3.3. A fibre clamp (Thorlabs, T711-250) with custom alterations is mounted on the positioners and holds the fibre, the tip of which protrudes from the clamp by 1/4" (Fig. 3.3b). Each positioner moves along a different axis (x, y, and z), and in this chapter, we will refer to the positioners by the directions in which they move the fibre: horizontal, vertical, and focusing.




Figure 3.3: a) Cryogenic SPIN setup showing: 1. a mounted sample, 2. the fibre clamp, and 3. the three XYZ piezoelectric positioners, labelled according to their directions of motion. b) Front of fibre clamp showing protruding fibre (red box). c) Fibre delivering laser to a sample.

Our choice of nanopositioner model was predominantly determined by the temperatures at which we operate, the need for precision focusing of the fibre, and the total travel distance required by our experiments. While Attocube does provide a range of ‘scanners’ for high-resolution motion, they are limited to ranges of less than 50 μm . In order to reach both

the sample and the optics setup (described in Chapters 2 and 5), we require a travel range of several millimeters.

Each positioner is controlled by its own Attocube ANM300 module, which allows for independent manipulation of its associated positioner. These modules generate the appropriate sawtooth function for step motion and allow one to manually set the step amplitude and frequency. Manual stepping of the positioner (either continuous or single-step) is enabled by toggle switches on the front of each module, with the direction of motion determined by the direction of the toggle switch’s activation.

Readout of the positioner location is accomplished using the resistive encoder. While Attocube does provide advanced control modules which come with built-in position readout systems, we did not possess such equipment when performing the experiments described in this thesis. We therefore devised a method of determining the location of each positioner using a separate voltage source (to supply the voltage V_0 across the resistor) and a multimeter (to measure the probe voltage S). Two multi-way electrical switchers (Electro Standards Laboratories, Path Way 7334) allow for connection of the voltage supply and multimeter to one positioner at a time. A calibration procedure was developed to determine the relationship between the actual stage position and the measured voltage S .

3.2.1 Calibration of the Resistive Position Encoding System

To determine the location of a positioning stage using the resistive encoder, the relationship between the measured voltage S and the position of the stage must be characterised. We therefore produced a calibration curve for each device, where the voltage S was recorded as a function of position.

For this calibration procedure, 1 V was applied to the resistive element within the positioner (this is the typical value of V_0 used in our experiments). Starting at one end of its range, the positioner moved in 100-step increments, making a measurement of S after each movement. Figure 3.4 shows the result, which is a roughly linear relationship between position and measured voltage. Fig. 3.4a) is the graph produced for one of the ANPx101/RES positioners, while part b) has the curve for the ANPz101/RES positioner. As our cryogenic optical system includes a piezoelectric rotator (Attocube, ANR240/RES), Fig. 3.4c) shows the rotator’s calibration curve for the sake of comparison.³ With these

³Since the rotator has no end point and can circulate continuously, its resistive element cannot cover its entire range. If it were to do so, the two ends of the resistor would be touching, and the system would not function properly. The resistive strip in the rotator therefore stretches only $\sim 340^\circ$ around the ring. (Note that the specifications sheet states the encoded travel range to be 315° [66], but repeated calibration measurements very clearly show ours to be 340° .) As the rotator’s circumference is far greater than 5 mm, the rotator requires many more steps to perform one full revolution than the linear positioners require to travel their full movement range. Consequently, the rotator also has very good resolution (micro-degrees, according to the spec. sheet).

curves, we are able to determine the exact position of the stage for a given measured voltage S .⁴




Figure 3.4: Calibration curves for the a) ANPx101/RES horizontal positioner, b) ANPz101/RES vertical positioner, and c) ANR240/RES rotator.

Further details about our calibration procedures and the reliability of these calibration plots are provided in Appendix D.

3.2.2 Computer Control of the Positioning Stages

In our system, computer control of the piezoelectric positioners is achieved through communication with an ANC300 control box. The ANC300 houses the ANM300 modules, which induce positioner movement. It also has a touchscreen to provide full control of all positioner functions in both coarse and fine positioning modes, with each positioner being individually addressable.

Commands are sent from a computer to the ANC300, which relays them to the appropriate positioners. The control software we created for these devices allows for direct control of all positioner functions from a command-line interface.⁵ The basic abilities of the positioners, such as continuous stepping and position readout, have been combined in this program to enable the positioners to perform more complex operations. The most important of these new functions is the ability for each positioner to move itself to any specified position within its 5 mm movement range. The positioner accomplishes this task by intermittently stepping and reading its position. The distance travelled between position measurements becomes progressively shorter as it nears the target location, to avoid overshooting.

When setting the positioner to a specific location, the program moves the stage as close as possible to the desired point. There is a certain window of acceptability about the target location, however, and when the positioner is within that window, movement will

⁴Details about the effects of noise on our ability to resolve the position will be discussed in Section 3.3.

⁵See Appendix B for full details of the homemade software we use to remotely control our equipment and automate our experiments.

cease and the program considers the device to have arrived at its destination. For example, if the tolerance window is $5\ \mu\text{m}$ and we ask the positioner to move $2.35\ \text{mm}$ from the ‘zero’ position, it will move to position $2.35 \pm 0.005\ \text{mm}$. The width of the tolerance range is primarily dictated by the noise in the resistive positioning system, which determines the uncertainty in our position measurements. Greater noise levels necessitate a larger acceptance window. It is therefore very important to reduce the noise as much as possible in order to achieve highly accurate positioning.

3.3 Imaging with Piezoelectric Nanopositioners

Distributed imaging of nanostructures is necessary for light delivery because we have no other method of viewing the sample when it is enclosed in the fridge. Focusing the fibre at the appropriate point therefore requires one to image the relevant region, creating a map of the sample’s recognizable features, and stop the fibre when it reaches the correct location.

To image a nanostructure sample in the dilution refrigerator, the piezoelectric positioners are used to scan the lensed fibre over the device. Light emitted by the fibre provides information on the sample surface, either through photocurrent detection or analysis of the PL spectrum. (We’ll look at the details of these imaging options in Section 3.4.) Throughout our development of this imaging technique, it has been necessary to find solutions to the many challenges presented by the properties and limitations of the piezoelectric positioners themselves.

In this section we will look at the development of our imaging procedure, the methods used to focus the fibre on the sample, and issues such as sample damage and heating.

3.3.1 Development of a High-Resolution Imaging Procedure

The basic procedure for imaging a nanostructure involves the use of custom-made, automated imaging software, which employs the resistive position encoder to move the fibre in a grid pattern. At each point in the grid, either photocurrent or PL is used to determine the structure of the sample directly below the fibre, allowing us to form a map of the sample as the fibre moves. In early tests of the imaging system, we performed large-scale scans, approximately $1\ \text{mm} \times 1\ \text{mm}$ in dimension, and the sample was imaged using a zigzag pattern, with the horizontal positioner switching directions with each successive row in the image grid.

Several related issues became immediately evident during the initial room-temperature tests of the system. The quality of many images was severely reduced due to an ‘interlacing’ effect, where alternating rows of the image appeared offset from each other. Figure 3.5a) shows an example of a nanostructure image distorted by the interlacing of consecutive rows. The source of this distortion was determined to be the control program for the

nanopositioners, which specified a large tolerance window about the target location during movement. With such a large acceptance range, the positioner was stopping at the leading edge of the window for each point in the grid. While travelling in the other direction for the next row, the positioner stopped at the opposite edge of the tolerance window for each grid point. The acceptance windows were so large that they overlapped with other points on the grid, creating an interlacing effect in the final image.

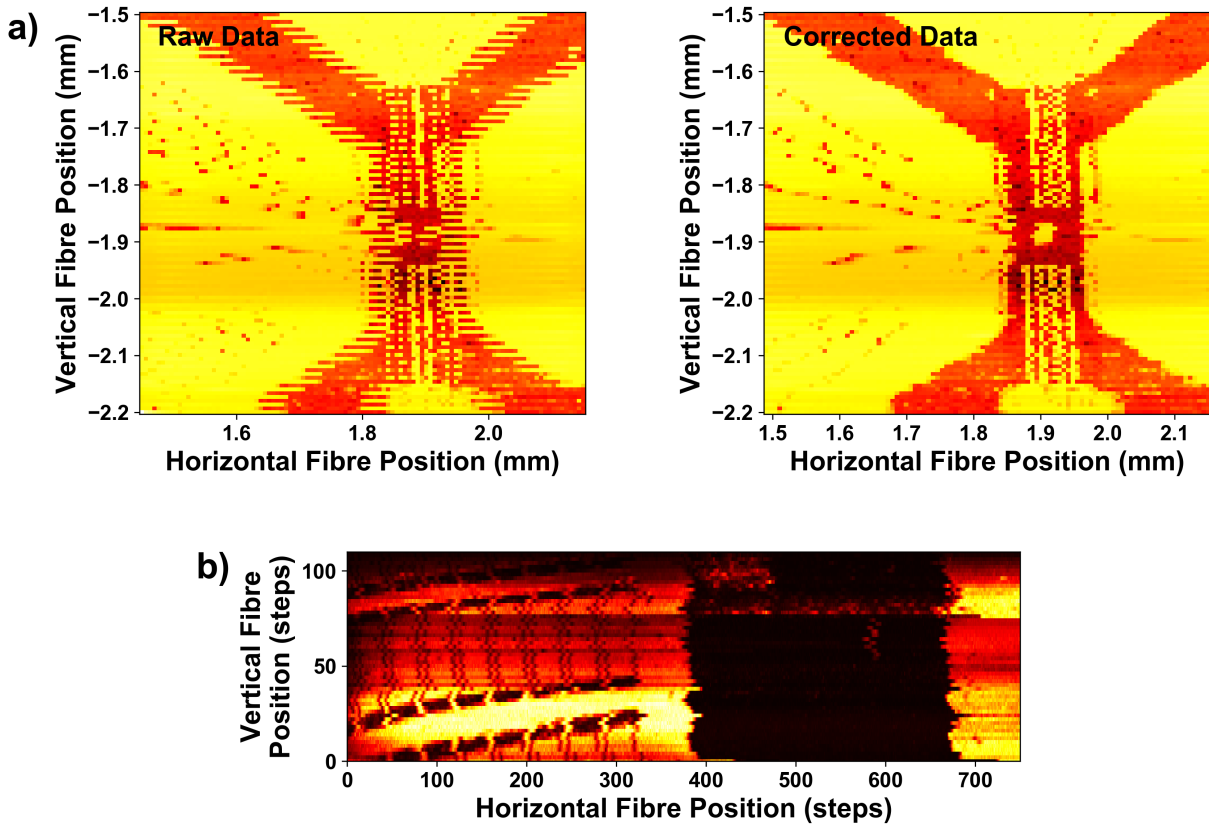


Figure 3.5: a) Interlacing distortion due to large position tolerance window. Left: Distorted image. Right: Artificially corrected image, with every other row shifted by ~ 0.05 mm. b) Distortion due to 1-2 μm repeatability limit of positioning system.

The reason for such wide tolerance windows was due to the variation in the distance travelled for a given number of steps. Successive movements of 100 steps each, for example, yielded an average travel distance of $13.5 \pm 5.0 \mu\text{m}$, where the uncertainty is nearly half as large as the distance travelled. This uncertainty was partially attributed to noise in the measured voltage S , but was mainly considered to be caused by the random variation in positioner step size, an issue from which these devices are known to suffer. The uncertainty in the travel distance becomes progressively worse for smaller numbers of steps; for a movement of 10 steps, the measured uncertainty is nearly as large as the average travel distance (roughly $2\mu\text{m}$). Since the positioner moves in successively smaller step increments

as it approaches a target location, the position tolerance window needs to be at least as large as the uncertainty in the smallest step increment. (Any smaller and we risked the positioner endlessly jumping back and forth over the desired position due to the smallest step increment being overly large.) We tended to err on the side of caution during preliminary tests, hence the $50\ \mu\text{m}$ interlacing in Fig. 3.5a), but we were eventually able to narrow the position acceptance window to a few microns.

Unfortunately, with the tolerance window at a few microns, we effectively reached the limit of our home-made positioning system. Although the noise level in the measured voltage S was eventually reduced to approximately $50\ \mu\text{V}$ (corresponding to a position uncertainty of roughly $400\ \text{nm}$), we could not reach this level with the positioning system. Increments of smaller than 10 steps (roughly $\sim 2\ \mu\text{m}$) tend to vary wildly in travel distance and sometimes appear to result in no observable motion according to the positioning system. Although we required the ability to resolve sample features on the order of $400\ \text{nm}$ or less, we were apparently limited to an image grid resolution of $2\ \mu\text{m}$.

This problem was solved by measuring the distance between grid points by number of steps, rather than measured positions. Whereas we originally instructed the system to acquire data at points separated by some absolute distance (e.g. $5\ \mu\text{m}$), the positioner would now be commanded to move a certain number of steps between measurements instead. Figure 3.5b) shows the result of such a scan. To avoid the issue of step size hysteresis (we estimate the positioners have a forward-backward step size ratio of roughly 5:6 at low temperatures), the zigzag pattern of imaging was abandoned in favour of rastering, with the horizontal positioner instructed to return to a specific position at the start of each row in the scan.

The image in Fig. 3.5b) shows no apparent distortion due to varying travel distances between points. This would seem to indicate that the large uncertainty in travel distance arose primarily from our measurement system, and the step size of the positioners is fairly constant in fact. The devices have a fundamental position repeatability limit of $1\text{-}2\ \mu\text{m}$ [64, 65], however, so this process did lead to some distortions of the image. With the horizontal positioner starting each row at a slightly different position, the straight lines of the sample appear wandering and jagged in Fig. 3.5b).

Importantly, the repeatability issue also means that we must re-scan the sample if we wish to return the fibre to a position which it has previously inhabited. The positioner coordinates measured when the fibre was previously at the desired location will give us a general idea of the region which needs to be scanned, but the repeatability limit ensures that we cannot use the measured coordinates to return to the same position.

To avoid image distortion due to repeatability issues, we have now switched entirely to step motion for the imaging process. At the end of each row in the image grid, the horizontal positioner quickly moves a certain number of steps back to the start of the row. (This is in contrast to the previous procedure, in which the positioner was instructed to return to a specific position at the end of each row.) Due to step size hysteresis, the

total number of reverse steps must differ from the number of forward steps.⁶ The number of backwards steps required to ensure each row starts in the same location is unknown beforehand, however, and must be estimated. Our imaging program allows for the number of reverse steps to be adjusted during the scan, allowing the user to observe whether the initial value is correct and make changes accordingly. The first few rows of every scan are therefore throwaway lines, as the experimentalist alters the number of reverse steps in order to properly align the features in the image.

A property of the positioners which has major implications for our imaging system is the stability of the step size. For a constant step voltage, the step size has a tendency to change significantly over time, with such changes typically occurring over a period of about 30 minutes to an hour. (The stability is highly variable though, and the step size has in some cases been observed to remain constant for several hours.) This is the main reason why one cannot predict in advance how many reverse steps to take at the end of each row, even if one has recently done a scan of exactly the same size. This phenomenon has a greater effect when taking many steps, as opposed to a few steps at a time. Our imaging scans usually take at least a couple of hours to complete, and as a result, the features in each row of an image appear to maintain a constant separation, but the starting point of each row begins to shift over time. The process must therefore be monitored so that the number of reverse steps can be adjusted if the image begins to drift. We currently have no general solution to this issue. Image recognition software could potentially be employed to adjust for deviations from the expected pattern, but the programming of such software would be complicated by the variety of nanostructures we test, the variation within individual samples (including unexpected damage), and the deformation of the image due to the different step sizes of the vertical and horizontal positioners. In spite of the temporal instability in the step size, however, the images we produce are generally of sufficiently high quality to allow us to locate the appropriate point on the sample where the fibre is to be focused and the laser light delivered.

3.3.2 Focusing the Fibre on a Sample

Once an imaging scan has revealed the desired feature, the scan is halted with the fibre directly above the appropriate location. Accurate delivery of light to the nanostructure then requires the lensed fibre to be precisely focused, which is a delicate procedure due to the 8 μm focal length of the fibre. With such a small distance between the fibre tip and the surface, one can easily damage the sample while attempting to focus. There are, however, several precautionary measures, such as manual operation of the positioners, which can be employed to help avoid sample damage. Additionally, focusing is only attempted after a sample has been cooled to base temperature and the fridge components have ceased to shift under the influence of the changing temperature. Two procedures have been developed to

⁶See Appendix D for details on the vertical positioner's rather extreme step size hysteresis.

focus the fibre on different types of samples, with one method based on image sharpness and the other involving the intensity of reflected light.

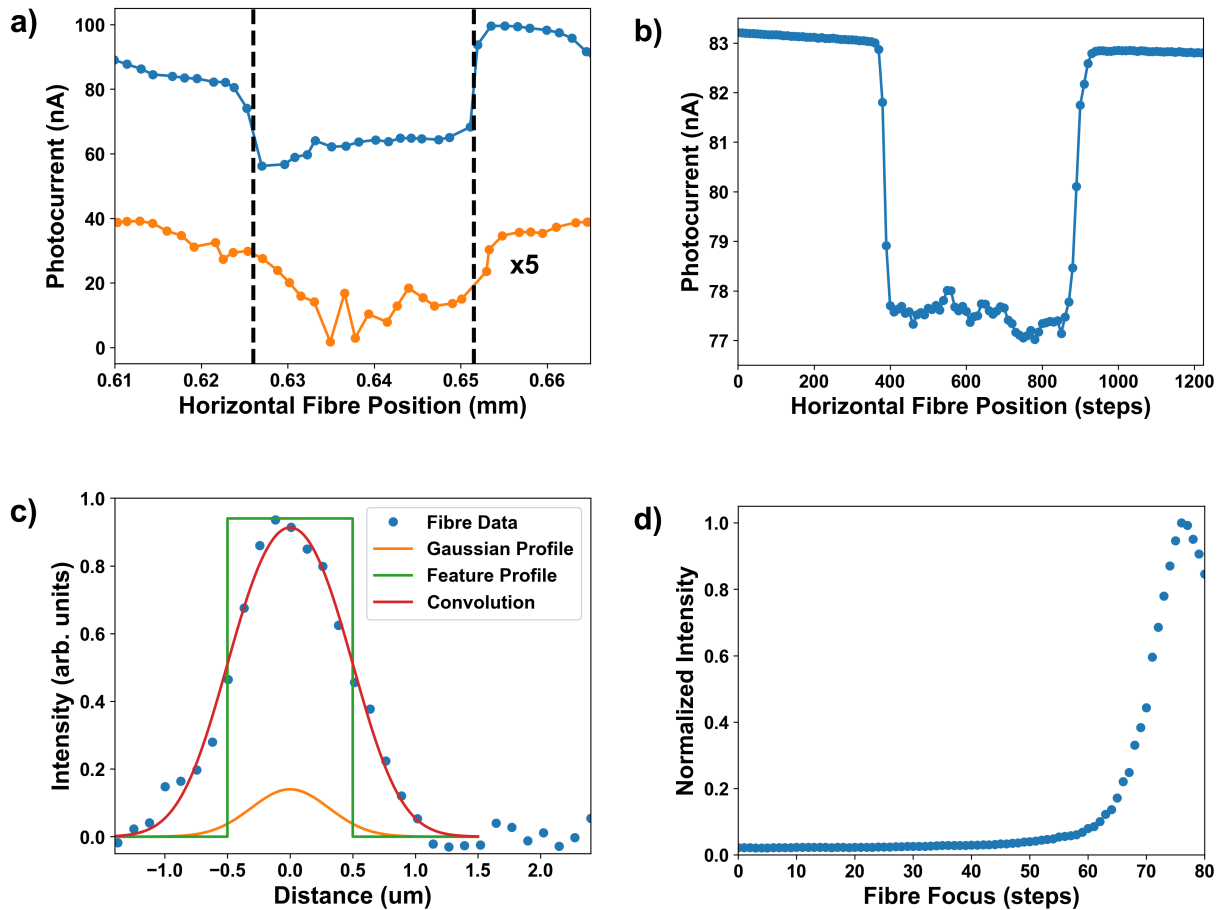


Figure 3.6: a) Image of a nanostructure feature becoming progressively clearer as it comes into focus. When out of focus (orange), the feature is indistinct and barely visible. This unfocused data has been multiplied by a factor a five to show that the dip in photocurrent is still visible, but only just. When in focus (blue), the image is much clearer, and the transitions are sharper. b) With the fibre highly focused on the sample, the transition between the two levels is very sharp. c) Data (blue) from fibre scanning over a $1\mu\text{m}$ -wide feature of a nanowire sample. The data is a convolution of the structure itself (green) and the Gaussian emission profile of the fibre (orange). The red curve is this convolution, from which we can extract the laser spot size. d) Intensity of light reflected by the sample substrate and collected by the fibre as it approaches the surface. The intensity reaches a maximum when the fibre is in focus, then it begins to drop again.

In the first type of focusing procedure, the fibre is repeatedly scanned over some feature of the sample, moving closer to the surface with each pass. The image of the feature becomes sharper as the fibre begins to focus on the nanostructure and blurs again once it passes

the focal point. Figure 3.6a) shows a scan over a lateral quantum dot, where the light from the fibre generated a photocurrent in the 2DEG. As the fibre moves into focus, it becomes much easier to distinguish between the two sections of the sample because the laser spot size on the surface narrows.⁷ The goal of this procedure is to find the point where the image is clearest, as this indicates the spot size is at its smallest and the fibre is in focus. Fig. 3.6b) shows a very sharp transition between the two sample layers, with the fibre well-focused. When examining a nanostructure feature which protrudes from the sample surface, one must decide upon which layer the fibre will be focused.

The sharpness of the features in one of these images is limited by the finite width of the laser beam waist. If the beam waist were infinitely small, then the image resolution would be limited only by the step size of the positioners. In reality, however, the finite beam waist means that the images produced are always a convolution of the sample features and the Gaussian emission profile of the fibre. Figure 3.6c) shows an example of data (blue points) for two rows of nanowires roughly 1 μm apart. Rather than observing a sharp difference between the sample surface and the presence of nanowires, the data shows a smooth curve instead. The green line shows the ideal scenario, with perfectly sharp features. The smooth curve of the data is the convolution (red line) of the true structure and the Gaussian profile of the fibre mode (orange curve). From a fit of this data, we find that the fibre has a spot size of approximately 0.75 μm ,⁸ which is very close to the beam waist value (0.7 μm) given by the product information sheet for the fibre. This spot size limits our image resolution, although we can still vaguely distinguish the presence of smaller features by a change in measured light intensity or photocurrent.

In the second focusing technique, the fibre delivers laser light to the sample and collects light which is reflected off the surface or emitted by nanostructures. The intensity of the collected light is then observed as a function of fibre position. When using reflected laser light, for example, Figure 3.6d) shows how the intensity increases until the fibre is in focus. At this point, most of the light should be perfectly reflected by the surface and collected by the fibre. The intensity begins to drop again as the fibre passes the focal point. Using reflected laser light can often be the fastest method of focusing on a sample, as it is simple to set up and does not require the fibre to be directly over the nanostructure.⁹

As with reflection, the measured intensity of light emitted by a nanostructure will increase as the fibre moves into focus. Figure 3.7 shows several low-resolution images of the PL emissions from a nanowire quantum dot as the fibre moves towards it. The ideal position seems to be somewhere around a focus position of 600-800 nm, where the collected intensity is greatest.

⁷The apparent shift in the position of the feature is a fairly common phenomenon and occurs because the fibre is not perfectly normal to the surface.

⁸This is an upper limit, as the fibre was likely not perfectly focused during this scan to avoid accidentally knocking over the taller wires. The fact that it is two rows and not one, as in the ‘ideal’ scenario indicated by the green line, may also affect our results.

⁹Provided the nanostructure does not protrude too far from the surface and the sample does not tilt towards/away from the fibre, this is probably the best method of focusing the fibre.

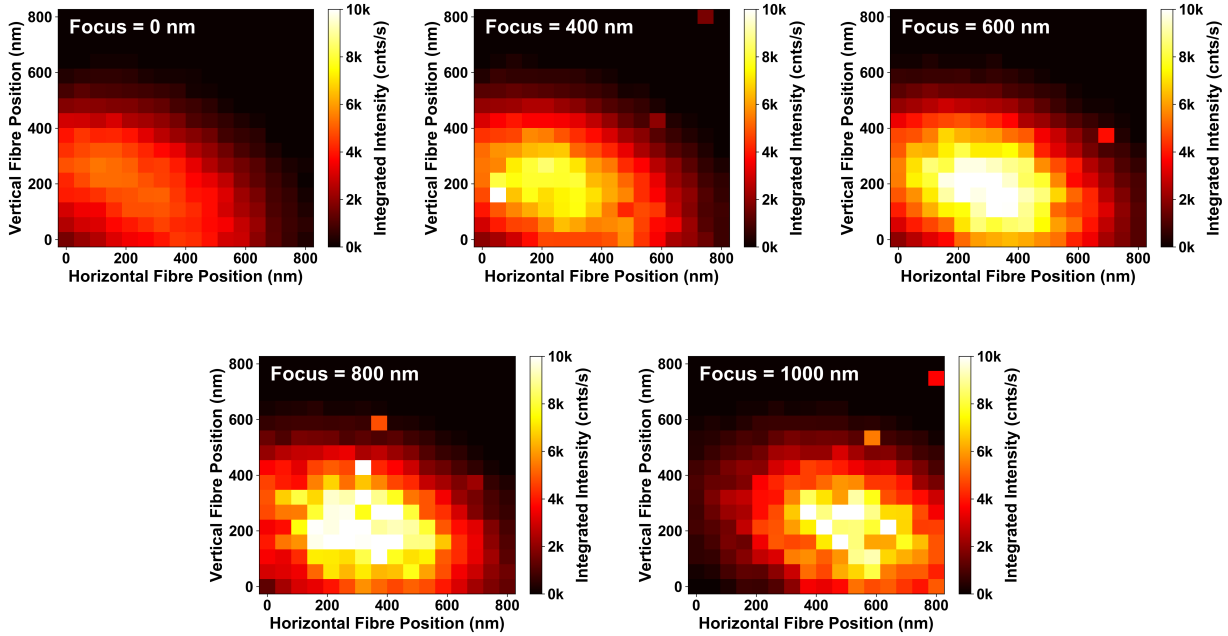


Figure 3.7: Measured intensity of nanowire emissions as a function of fibre position and focus. The fibre moves closer to the sample as the focus value increases from 0 nm to 1000 nm.

When using PL for focusing, however, there are several factors which make it difficult to determine the true focal point. In order to collect the most light, the fibre should be focused on the tip of the nanowire, where the nanowire emissions begin to diverge. Since the quantum dot is embedded near the base of the wire in our samples, the dot needs to be pumped at saturation, where we know that an increase in pump power (due to the approach of the fibre) will not result in an increase in dot emissions. There is also the matter of the difference between the wavelengths of the laser and dot emissions, which will affect how the light is collected/emitted by the fibre.¹⁰ To focus the fibre using PL, the fibre should be moved to the position where the greatest intensity is observed, with the laser power adjusted to ensure the device is always emitting at saturation.

3.3.3 Heat Output of the Piezoelectric Positioners

Two major sources of heat in the dilution refrigerator are: 1. the mechanical motion of the positioners, and 2. the Joule heating of their resistive elements. While these heat sources have little effect at temperatures above ~ 2 K, one must monitor the heat output of the positioners closely when performing experiments at milli-Kelvin temperatures. Extreme heat output can alter the state of the sample and disturb the ^3He - ^4He mixture used to cool

¹⁰This issue would be resolved if resonant pumping were used.

the fridge, which could require the system to be shut down. To reduce the probability of the positioners heating the fridge above safe levels, a temperature monitoring feature has been included in all of our imaging software. When the fridge temperature rises above a certain threshold (usually 500 mK), the scan is automatically halted until the temperature falls low enough that the imaging process can be resumed safely.

A potential of 1 V is typically applied across the resistive position encoders. This voltage is large enough to give us decent position information with minimal noise, but it is also small enough that we can read the position at low temperatures without overwhelming the fridge with heat dissipated by the resistor. With 1 V applied to the resistor of a linear positioner, the fridge temperature will rise from base temperature to about 130 mK within 2–3 minutes. The position readout system is therefore only turned on when the fibre is in motion, and electric switchers are used so that the voltage is applied to only one positioner at a time.

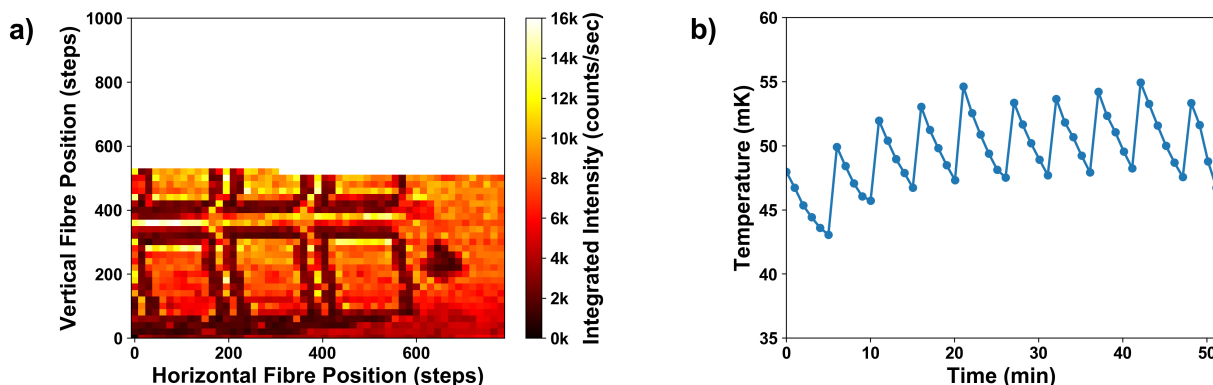


Figure 3.8: a) A scan of a nanowire sample acquired with a raster pattern. b) Fluctuations in the mixing chamber temperature caused by the movement of the positioners as several rows of the image were acquired.

High-resolution images, which require the positioners to stop more frequently, reduce heat dissipation in the fridge. Scanning small regions and including periodic wait times during the scan can also help maintain lower temperatures. If a scan is using step number, rather than distance, to move the fibre to the appropriate points in the image grid, the resistive encoders may be turned off. Figure 3.8a) shows an image of a nanowire sample, which was obtained by moving the fibre in a raster pattern without the use of the resistive encoders.¹¹ Fig. 3.8b) displays the temperature of the fridge during a portion of this scan. Note the slow decay of the temperature during the acquisition of each row, where the fibre moves slowly enough and stops frequently enough that the cooling power of the fridge exceeds the heating due to the positioner. When a row of the scan is complete, the horizontal positioner quickly returns to the start and begins the next row. These fast

¹¹The image represents a scan of dimension $80 \mu\text{m} \times 30 \mu\text{m}$.

backwards movements, which are done without stopping to acquire any data, produce a lot of heat and correspond to the abrupt temperature increases in the plot. The combination of slow forward motion followed by quick backwards motion results in a roughly stable MC temperature of about 50 mK.

To help reduce unwanted heating of the sample, all heat sources in the fridge are thermally anchored to the MC to ensure they are properly heat sunk. For more details on the thermal anchoring of the various experimental components, see Section 2.1.2.

3.3.4 Sample Damage Caused by Fibre During Imaging Scan

In generating an image with the lensed fibre, great care must be taken to avoid damaging the sample. From melting the nanostructure by overheating it with a focused laser [67] to scratching or gouging the surface with the fibre, unintentional damage comes in many forms. Control of the nanopositioners must be precise enough to focus the fibre during imaging scans without moving it so close as to touch the sample. Nanowire samples present particularly unique challenges, as the small, upright crystals are difficult to locate and easy to accidentally destroy with just a few steps of the positioners.

Figure 3.9 presents several micrographs of a nanowire sample which was damaged during our imaging scans. An undamaged sample is seen in Fig. 3.9a), while Fig. 3.9b) shows our sample after 5 months of experimentation and imaging. All features on this sample, be they numbers, letters, or boxes, are formed from nanowires. In our sample, the lower set of boxes (outlined in red) show very obvious damage where the fibre was scratched along the sample. At the centre of each of these boxes is a single nanowire, which is visible in the picture as a small, bright dot. These are the nanowires which we study, and many are absent in this picture, having been accidentally destroyed by the fibre during our scans and experiments.

The principle reason for this accidental destruction is the focal length of the lensed fibre, which is just 8 μm . This is nearly the same length as the nanowires themselves, which stand approximately 5 μm tall. Consequently, there are several factors which can, individually or in combination, result in our unintentionally and unknowingly placing the fibre too close to the sample during a scan. Such factors include the laser power being more attenuated than expected, the light collected by the fibre having an unexpectedly low intensity, or the sample being slightly tilted towards the fibre. Additionally, if we are trying to focus on a wire, but the fibre is not properly aligned, the collected wire emissions may not get significantly brighter as the fibre moves towards the sample. This could be interpreted as the fibre being far away from the wire, leading us to move it too close to the sample, destroying the wire.

Figure 3.9c) is a SEM image of a dense nanowire array, which appears as a black rectangle in the bottom right corner of Figures 3.9a) and b). This image clearly shows that during several of the imaging scans, the fibre passed too close over the array, knocking




Figure 3.9: a) Microscope image of an undamaged nanowire sample (which was not used in our experiments). b) Microscope image of our sample after 5 months of experimentation and imaging. c) SEM image of the dense nanowire array, with lines of wires knocked over by the fibre. d) SEM image of damage to the dense nanowire array. Inset: micrograph of the array showing gradual appearance of the lines, indicating the sample was tilted. e) Microscope image and f) SEM image of an area where the fibre penetrated the surface. g) Holes in the substrate caused by the fibre making repeated contact with the surface.

over all the nanowires in its path. Figure 3.9d) shows another section of the array where the fibre made three horizontal lines. The small inset,¹² highlighted by a white box, is a micrograph showing how these three lines appear gradually in the middle of the array. This gentle appearance suggests that the sample was tilted with respect to the imaging plane of the positioners, so as the fibre was scanned along a row, it moved slowly closer to the surface until it was knocking over the nanowires. The angle between the sample and the imaging plane is due to the alignment of the positioners themselves. The positioners must be removed in order to mount a sample and reconnected after the sample is in place. Placement of the positioners is done manually, and they are aligned with the sample by eye. Given the awkwardness and delicacy of mounting these expensive devices with tiny screws while kneeling on the floor, while simultaneously attempting to prevent the fibre tip from touching the sample, accurate alignment of the positioners can be quite difficult. It is therefore common to have a small offset angle between the sample and imaging plane, which increases the probability of accidental contact between the fibre and nanostructure.¹³

Due to fluctuating laser intensities and sample tilt, it can often be difficult to determine whether the fibre is too close to the sample, especially if it has already moved past the focal point. Consequently, there have been several instances where the fibre has been pushed too far forward, resulting in direct contact with the surface. Figures 3.9e) and f) show an optical microscope image and a SEM image respectively of an instance where the fibre was accidentally pushed into the sample. (This hole can also be seen in Fig. 3.9b), where it spreads across the upper set of boxes.) Based on the diameter of the hole and the shape of the fibre lens, the tip of the fibre likely penetrated approximately 25-50 μm into the sample. Smaller indentations due to the fibre also appear elsewhere on the sample, such as those seen in Figure 3.9g). While the fibre is fairly flexible at room temperature, and although we would expect it to remain so at cryogenic temperatures [50], it in fact appears to be quite solid at low temperatures. Rather than bending upon contact with the sample, it retains its shape as it is pushed through the surface layers.

The damage is also not entirely limited to the sample. When scratched along the surface or pushed into it, the fibre tip becomes chipped. Figure 3.10a) shows three SEM images of the fibre tip, with each image from left to right progressively zooming in. The image on the right is of the fibre lens itself, which shows obvious damage. This is the fibre which was used to test the damaged nanowire sample shown in Fig. 3.9, and these images were taken at the same time. Although the overall damage to the fibre tip appears relatively minor (in spite of the destruction it inflicted upon the sample), the lens itself (which is 4 μm in radius) is quite damaged and the consequences for our experiments is unknown. While we

¹²The inset provides a wider view of these three lines, as the SEM image did not capture their full extent. It was taken from the microscope image in Fig. 3.9b).

¹³Monitoring an imaging scan to ensure the fibre maintains a safe distance from the sample surface can be a tedious task. On at least one occasion, it is known to have kept a tired grad student up all night, checking on the scan status every few hours and adjusting the focus as necessary. This event prompted the (perhaps obvious) realisation that the more intelligent course of action would be to only perform scans where the fibre slowly moves away from the sample. Identifying the direction of sample tilt is easily done by observing the intensity of reflected light as the fibre is stepped along the horizontal and vertical axes.




Figure 3.10: a) Zooming in (from left to right) on the damage to the tip of the lensed fibre, which was likely caused during accidental contact with the sample. The right-most image is of the fibre lens itself, which is clearly damaged. b) Broken nanowires attached to the fibre tip.

observed no significant impact on our experiments, it is possible that the broken lens may distort the spatial mode of emitted light or reduce collection efficiency.

In addition to the damage accrued by repeated contact with the sample, the fibre also collected many broken nanowires upon its tip. Figure 3.10b) shows several broken nanowires stuck to the fibre. While there was no noticeable evidence that these wires were emitting light which was then collected by the fibre, it is possible that coupling to the fibre mode could occur if the broken nanowires were properly oriented on the lens. Such coupling would introduce noise and unwanted peaks into the measurement spectrum when studying other nanowires, so care is taken to ensure that there are no persistent peaks in the spectrum as the fibre moves over the sample.

There are a couple options available to help minimize the risk of damaging the sample.

Prior to closing the fridge, one can locate by eye a region on the sample near the nanostructure, but where the surface is empty. After the fridge has been cooled to base temperature, the fibre can be focused on that empty point. This prevents damage to the device if the fibre accidentally makes contact with the surface. If the nanostructure protrudes from the surface, then the fibre can be pulled back slightly from the point where it is focused on the surface, to reduce the risk of hitting the device if the sample is tilted. In the case of the nanowires, which were nearly as tall as the fibre focal length, it was generally sufficient to pull the fibre back until the intensity of the reflected light dropped to 1/10th of the intensity when in focus.

A more accurate, though somewhat risky, method of ensuring the fibre is far enough away from the sample is to focus the fibre directly on the nanostructure itself and then pull back. Ideally, if multiple identical nanostructures are present on the sample, this focus-defocus method should be performed on one which will not be studied but is near the experimental candidate (as a tilt in the sample might otherwise result in the fibre hitting the desired device). This is only necessary when scanning a larger region, however, and one may remain focused on the device if sample tilt is not expected to be significant over the scanning range. With procedures such as this in place, it is possible to produce high-resolution images of a sample without damaging it.

3.4 Options for Measuring Sample Topography

As different samples have different properties, it is necessary to tailor the imaging technique to the type of device. We have therefore developed three methods of imaging samples, each of which has its own advantages in certain scenarios. The first of these options relies upon the photocurrent generated in a sample when exposed to light of the appropriate wavelength, the second involves the collection of laser light which is reflected off the sample surface, and in the third, the fibre collects the photoluminescent emissions of the nanostructure. Of these three, PL imaging has been our most-used method, partially due to happenstance,¹⁴ but also because it provides the most information about the sample. The following sections will review these three methods and provide examples of the images they produce.

3.4.1 Photocurrent Imaging

The devices we imaged using the photocurrent method were lateral DQD samples which were intended for use as part of the Photon-to-Spin project. (See Section 1.1 for information on the properties and structure of these samples.) Figure 3.11a) shows a microscope picture of one of these samples. When illuminated by light of sufficiently high energy, electron-hole pairs are generated in the sample 2DEG, and if the 2DEG is connected to ground through

¹⁴We happened to be performing a lot of experiments with samples compatible with this type of imaging.

an ammeter, the addition of these photo-generated carriers results in a measurable current on the order of tens of picoamps. As the lensed fibre is scanned over the sample using the nanopositioners, the photocurrent is only detected when the fibre is over the sample mesa, under which is located the 2DEG. The gates and optical mask block the laser light, preventing it from generating current in the 2DEG. Figure 3.11b) shows a photocurrent map of a Photon-to-Spin sample, where yellow/orange indicates high photocurrent and red indicates no current. Some scattering of electrons along the gates causes a slight photocurrent to be detectable at the gate edges (these are the lines radiating from the top of the image).




Figure 3.11: a) Microscope image of a DQD sample. b) Partial photocurrent map of the sample, showing a large current (in yellow/orange) when the light is focused on the sample mesa.

Near the top of the photocurrent image, the small dark square at the centre of the mesa represents the optical mask. It is designed to shield the device from the laser light, except for a small hole (400 nm diameter) in the centre so polarised photons can be delivered directly to one of the quantum dots. The goal of imaging these samples was to find that hole, which should have appeared as a bright spot in the centre of the mask, as it is an opening to the 2DEG below. At the time we acquired the image in Fig. 3.11b), however, the resolution of our imaging system was too poor to observe the hole in the optical mask. (Also, as discussed in Section 2.2.2, the optical mask hole was in fact absent from our device due to a fabrication issue.)

Recently, Professor Akira Oiwa's group from Osaka University has provided us with new photon-to-spin samples. These devices have a slightly different design, and are actually quadruple quantum dots, with two sets of quadruple dots per sample. A hole in the optical mask is placed above a single dot for each set of dots, so we should be able to observe two holes in the mask. Myself and Tomohiro Nakagawa, a PhD student from

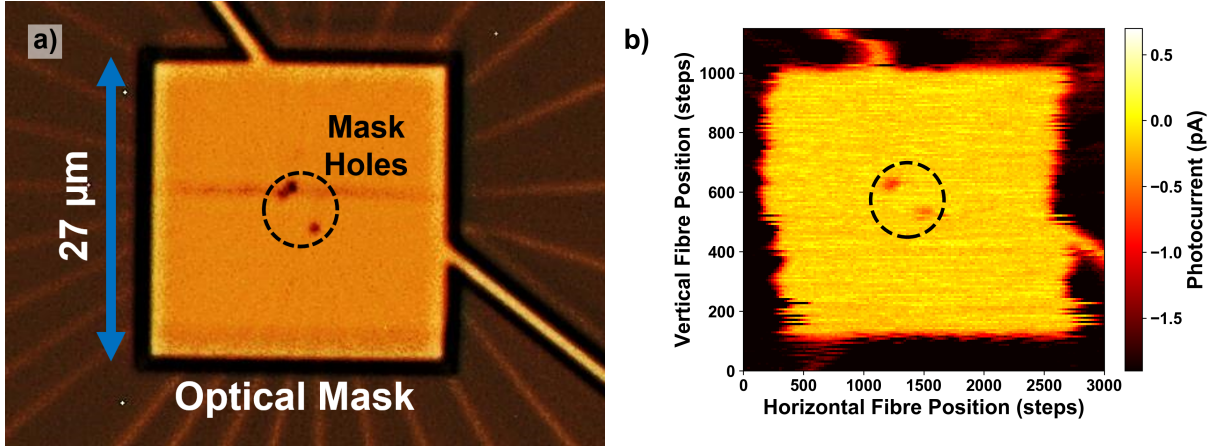


Figure 3.12: a) Microscope image of a photon-to-spin quantum dot sample, showing the optical mask ($27 \mu\text{m} \times 27 \mu\text{m}$) and the gates converging beneath it. The two mask holes in the centre are clearly visible, with a speck of dust visible beside the upper hole. b) Photocurrent scan of the sample from a), showing low current (orange/yellow) when the fibre is over the mask. The two dark spots highlighted by the dashed circle are the 400nm-wide holes in the optical mask. (They appear wider in this image due to the fibre focus and laser spot size.)

Osaka University, have tested the SPIN Lab’s imaging system on two of these samples. Figure 3.12a) is a microscope image of one of the devices. The optical mask ($27 \mu\text{m} \times 27 \mu\text{m}$) blocks the light and prevents a photocurrent from being generated. The two mask holes outlined by the dashed circle, however, should allow the light to pass through and generate a photocurrent.¹⁵ Below and around the mask lies the mesa, which houses the 2DEG and experiences a large photocurrent. Radiating out from beneath the mask are the gates, which reflect light.

Figure 3.12b) shows the photocurrent scan of the optical mask, with the photocurrent scale bar adjusted to highlight features on the mask. At the centre of the mask, highlighted by the dashed circle, we can clearly observe the two holes at the exact locations where we expect them to be. At these points, light passes through the mask and generates electron-hole pairs in the 2DEG, which is necessary for photon-to-spin conversion. This is proof that using our imaging system, we can locate the mask holes through photocurrent detection, and by stopping the imaging scan when the fibre is above one of the holes, we can focus the fibre directly on it. Once electrical tests of this device are complete, we can begin to test our system’s ability to deliver photons to a quantum dot and detect individual photo-generated electrons.

¹⁵A speck of dust next to the upper dot on this particular sample does not affect our photocurrent images, but as we shall see in the following section, it is quite evident in our reflectance scans.

3.4.2 Reflectance Imaging

The reflectance technique produces an image by using a photodiode to measure the intensity of light which is reflected off the sample surface and collected by the lensed fibre.¹⁶ Since the lensed fibre both delivers and collects light, a fibre circulator (or beam splitter, if necessary) is required to separate the incident laser beam and the reflected light. For surfaces covered by different materials which absorb or reflect light differently, we can map out samples according to their composition. Alternatively, this process can also be used to create high-resolution images of samples which have surface features of varying heights, as the change in focus as the fibre passes over these different layers results in different levels of light collection.




Figure 3.13: Reflectance scan of the photon-to-spin sample from Figure 3.12a). The two dark spots highlighted by the dashed circle are the 400nm-wide holes in the optical mask. They appear wider in this image due to the fibre focus and laser spot size, and the upper dot is more prominent due to the presence of a dust particle next to the hole.

Figure 3.13 shows an example reflectance scan of a photon-to-spin lateral dot sample from Osaka university (see previous section). The gold optical mask reflects a great deal of light, while the mesa absorbs light. The two optical mask holes are observable in the centre, highlighted by the dashed line. The upper hole is more noticeable, likely due to light scattering off of the dust particle next to it (see Figure 3.12a). The indistinctness of the lower dot may be closer to what we should normally expect to observe when imaging these samples using reflectance. Comparing this result to the scan in Figure 3.12b), it appears

¹⁶Reflectance is a measure of a surface’s ability to reflect electromagnetic radiation, which we define as the fraction of the power of the incident light reflected by the material surface. In our imaging experiments, we measure the intensity of light reflected by the sample surface for a constant pump power, which allows us to detect changes in the surface topography, such as the presence of a nanostructure.

that photocurrent imaging is the better option for locating the mask holes and optimizing the fibre position for light delivery to the quantum dots below.¹⁷ Reflectance measurements are still useful for large-scale imaging of the sample, however, and can be used to locate the optical mask with the fibre.




Figure 3.14: a) Microscope image of a nanowire sample, where all dark structures (numbers, letters, boxes, etc.) are composed of nanowires. b) Map of nanowire sample compiled using reflected laser light. The individual, 250nm-wide nanowire at the centre of each box is resolvable by our system.

The reflectance imaging technique has also been tested on nanowire samples, such as the one shown in Figure 3.14a), as these devices are not designed to allow for photocurrent measurements. While the sample substrate is highly reflective, the nanowires themselves tend to scatter and absorb the laser, resulting in a decreased reflection intensity as the fibre passes over them. Figure 3.14b) shows an example of this reflectance mapping method, where the dark regions, including the boxes, numbers and letters, indicate the presence of nanowires. Note also the changing intensity of the background along the vertical axis. This is an indication of a tilted sample, which causes the amount of reflected light collected by the fibre to change as it moves vertically. For large scans of tilted samples, the fibre must be periodically refocused in order to maintain the light intensity.

The image of the nanowire sample (Fig. 3.14a) shows 10 boxes, and at the centre of each is a single nanowire (5 μm tall, 250 nm wide). These nanowires are well-formed and have high emission rates, making them perfect candidates for our experiments. Consequently, it is very important that the resolution of the imaging system is sufficient to allow us to locate them and focus the fibre on them. Indeed, Figure 3.14b) shows that the nanowires

¹⁷It also makes more sense to use photocurrent measurements for final fibre alignment: the photocurrent specifically indicates that light is passing through the mask to the 2DEG below, which is what we require for photon-to-spin conversion.

appear as a very clear dark spot at the centre of each box in the reflectance scan. The ease with which we are able to locate the nanowires through reflectance mapping has made this technique the primary method for large-scale imaging and preliminary detection of the nanowires we wish to study. Final alignment and focusing of the fibre is usually accomplished manually using PL measurements.

3.4.3 Photoluminescence Imaging

In addition to photocurrent and reflectance mapping, we are also able to image our samples using PL. As described in Section 1.2.3, optical excitation of a semiconductor generates electron-hole pairs, which emit a photon when they recombine. A PL image of the sample is generated using a similar setup as reflectance, where a circulator is used to direct the nanostructure emissions to a spectrometer. The spectral data allows one to map the sample according to the different wavelengths and intensities of the light collected by the lensed fibre as it moves over the sample. The majority of this light will again be reflected laser light,¹⁸ but a filter placed in front of the spectrometer reduces its intensity to a more manageable level, i.e. a level which will neither damage the spectrometer, nor overpower the PL emissions.




Figure 3.15: Example PL spectrum of a nanowire, showing emissions from the nanowire crystal, emissions from its embedded quantum dot, and reflected laser light. (Laser has been heavily filtered.)

PL imaging was found to be the ideal method for detailed mapping of our nanowire quantum dot samples, as it provides spectral information which can be used for preliminary analysis of nanowire quality. As shown in Figure 3.15, there are three main peaks which

¹⁸We use above band-gap, non-resonant pumping of the wires, so the reflected laser light has a shorter wavelength than the PL photons and can be easily distinguished from them.

we typically observe in a nanowire emission spectrum: InP nanowire emissions, InAs(P) quantum dot emissions, and reflected laser light. If recombination of an electron-hole pair occurs within the InP crystal structure of the nanowire itself, the resulting photon has a wavelength of approximately 830 nm. This peak in the spectrum is broad and very bright. If the recombination occurs within the nanowire’s embedded quantum dot, we observe narrow peaks in the spectrum with wavelengths ranging between 900 and 1000 nm. The exact emission energy varies due to the slightly differing composition and geometry of individual dots, but we generally wish to study those dots which have extremely narrow emission lines at about 950 nm. We also observe the filtered laser peak, of course, which we typically tune to around 790 nm. A more detailed analysis of nanowire emission spectra and our PL collection techniques can be found in Chapter 4.

During a PL imaging scan, the measured spectrum is divided into three separate wavelength ranges based on the main peaks: the laser (780-800 nm), the InP nanowire emissions (825-895 nm), and the InAs(P) quantum dot emissions (905-1000 nm).¹⁹ For each point on the image grid, the measured intensity is integrated over the three spectral ranges individually, providing a single value to represent the amount of light retrieved in each range.²⁰ For a single scan of the fibre, three plots are therefore produced, one for each emission source. The reason for doing this is that each wavelength gives us slightly different information about the sample. Reflected laser light is great for basic, large-scale imaging, while PL emissions from the nanowires are bright and allow for rough alignment of the fibre. The dim quantum dot emissions give information about the quality of the dot itself and can be used for fine-tuning the fibre position to maximize coupling between the wire and fibre.

Figure 3.16 shows three versions of a single image, with each version formed from the integrated counts of a different wavelength range (laser, nanowire, and quantum dot). Due to scattering, the measured intensity of the reflected laser light is reduced when the fibre passes over nanowires, while the images created from PL emissions are brighter in the presence of nanowires. To compare with the large-scale image in Figure 3.14, these images are of the first two boxes in column 3. The wires in these boxes had already been destroyed by the time this image was taken, which is why the scan stops at a wire in column 2 instead (left-hand side of images). Note the attempt to compensate for the hysteresis of the positioner step size, which introduced an offset in the first few columns (right side of each image) as the number of backwards steps was manually adjusted. Each plot is normalized to its largest value.

The waveguide shell and tapered tip of the nanowires produce a Gaussian emission profile [6], ideal for coupling to an optical fibre, and ensure a relatively narrow divergence

¹⁹The exact extent of each range might vary from scan to scan depending on how precisely we wished to study a particular wavelength range and whether we were seeking to exclude unwanted features, such as emission lines produced by defects in the nanowire.

²⁰The maximum value in each range could also be plotted at each point in the image grid, but integration over the spectral range seems to reduce noise and the random, unwanted peaks occasionally produced by cosmic rays. To eliminate background noise, only values above a certain intensity threshold are included in the integrated values.




Figure 3.16: PL image of a nanowire sample created using the integrated counts of a) the reflected laser light, b) the nanowire emissions, and c) the quantum dot emissions.

of the emitted light. While this means that wires can be easily missed during low-resolution scans, the benefit of the narrow light cone is the ability to isolate the emissions of a single nanowire without background light generated by other wires. In searching for an ideal nanowire²¹ to study in a large array, the fibre is scanned along a row of wires, allowing one to examine the unique spectrum of each wire. Figure 3.17a) shows an example of such an array, while Fig. 3.17b) shows the integrated counts and several emission spectra acquired from a scan along a row of nanowires with a $10\ \mu\text{m}$ pitch (i.e. a separation of $10\ \mu\text{m}$ between nanowires). Note how the peaks are sharp in the plot of integrated counts, and each one appears only when the fibre is over a wire. This indicates that we are isolating the emissions of each wire individually, without overlap. Since the emission spectrum is measured at each point in the scan, this system permits simultaneous imaging of the sample and preliminary analysis of each individual wire's emissions. The imaging program can also be set to stop the scan when nanowire emissions are detected, so the experimentalist can refocus the fibre and determine whether the wire merits further study. The scan can then be resumed from that point if the wire is not acceptable.

Recalling that these imaging scans are the only method of finding a nanostructure with the lensed fibre, when the desired feature is found, the scan is immediately stopped and the fibre is focused/aligned by manually stepping the positioners.²² When studying nanowires, the PL intensity of the quantum dot may be relatively weak depending on the laser power and fibre focus/alignment, so the bright nanowire emissions are used for rough positioning of the fibre, while focusing and final alignment adjustments are done

²¹In this case, the 'ideal wire' might be defined by its brightness; the wavelength, number, and spacing of peaks in its quantum dot emissions; and the number of unwanted emission lines caused by defects in the crystal structure.

²²A larger scan of the region might be performed first to determine which nanowire warrants closer examination, but a second scan will then need to be done in order to find that specific wire again with the fibre.

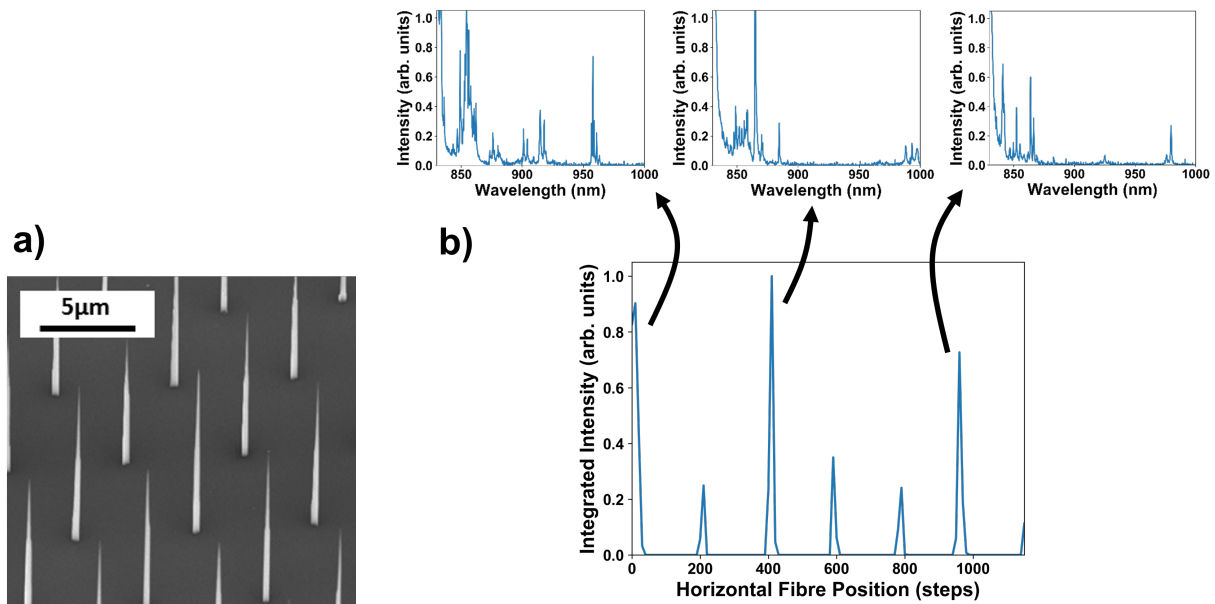


Figure 3.17: a) SEM image of an array of nanowires with 5 μm pitch. Image provided by Dan Dalacu. b) Integrated counts of the nanowire emissions as the fibre is scanned along a row of nanowires with 10 μm pitch. We are able to observe the spectrum of each wire as the fibre passes over it. The spectra of the three brightest nanowires are shown.

(manually) by observing the intensity of the quantum dot emissions.

To assist with alignment of the fibre, the fine positioning mode of the positioners is occasionally used instead of manual control alone. Fine positioning mode permits sub-nanometer movement and can be used to generate high-resolution PL maps, like the image of the nanowire quantum dot emissions shown in Figure 3.18. The sample was imaged in a grid pattern with a resolution of approximately 5 nm. The apparent emission profile in this image, which is far wider than the nanowire itself,²³ is the convolution of the fibre's Gaussian mode and the divergence of the quantum dot emissions (the nanowire has a numerical aperture of 0.5). This image shows that our system can resolve the PL emissions of a single nanowire quantum dot with nanometer-scale resolution. The ability to isolate the PL of a single nanowire is extremely important for our single-photon experiments (see Section 4.3), wherein it is necessary to capture the individual photons emitted by a specific dot. For best results during experiments, the fibre is focused on the centre of the nanowire emission profile, where it collects the most light emitted by the quantum dot.

²³The nanowire's waveguide shell has a diameter of 250 nm at its widest, while the quantum dot is about 20 nm wide.

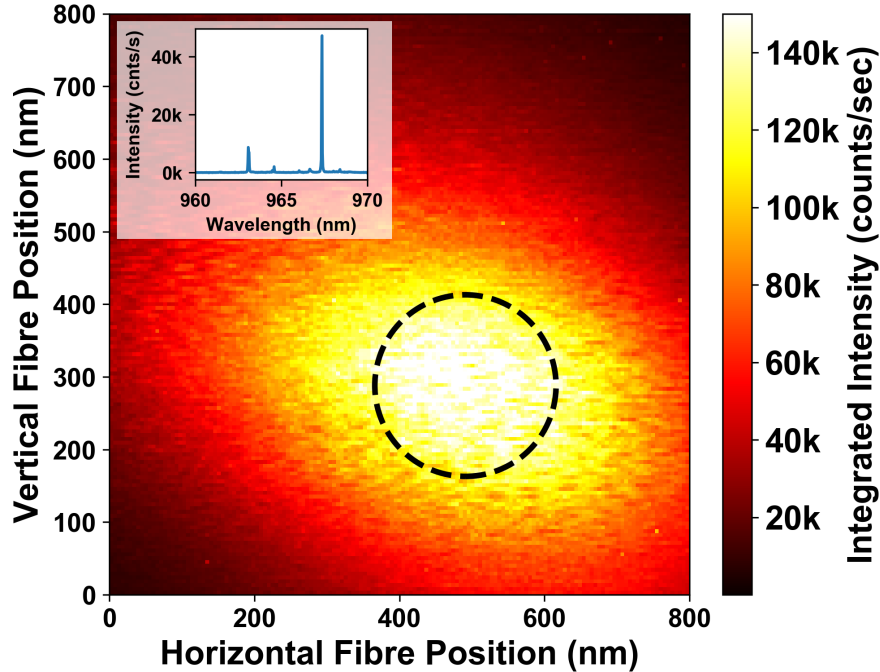


Figure 3.18: A high-resolution spatial map of a single nanowire’s PL emissions, which was acquired using the DC positioning mode of the positioners. The dashed circle with a 250 nm diameter indicates the width of the nanowire base. Inset: The quantum dot’s emission spectrum.

3.5 Conclusions and Future Upgrades to the Imaging System

The diverse optical properties of different semiconductor nanostructures indicate that there may be a wide variety of practical applications for these devices, from providing entangled-photon pairs for quantum communication to acting as highly efficient solar cells. Development of these technologies, however, often requires that the devices be cooled to milli-Kelvin temperatures in a dilution refrigerator, making optical experimentation difficult. Light delivery and collection is commonly achieved through free-space beams, requiring windows in the fridge. To avoid the heating and background illumination resulting from such openings, our system instead uses a lensed optical fibre to deliver laser light to the sample and collect PL emissions.

The scanning techniques we have developed allow for high-resolution imaging of samples. Given that far more precise and advanced imaging technologies exist, however, the purpose of our system is rather more functional: the detection of specific features of a nanostructure sample and the accurate focusing of the fibre on them. To that end, our imaging system performs its intended function very well, and through successive improvements to

our methods, the resolution and accuracy with which these scans can be generated has been vastly improved. While the laser beam waist limits our ability to detect details which are smaller than about 100 nm, we can now achieve sub-nanometer positioning and PL from individual nanostructures can be collected. This system, therefore, has been and will continue to be vital to our ability to successfully perform optical experiments on nanoscale devices.

Nonetheless, further improvement is always possible. Since performing the experiments summarised in this thesis, we have acquired an ANC350 control box for the piezoelectric positioners. This device has a position readout system incorporated into its design, and while still limited by the 1-2 μm repeatability of positions, it performs many of our homemade functions smoothly and automatically, such as setting a positioner to a specific position. It employs calibration files created by Attocube²⁴ which are far more accurate and have higher resolutions than our home-made versions. The ANC350 eliminates many of the issues we experienced when positioning the fibre, such as the interlacing effect described in Section 3.3.1, but it cannot overcome the position uncertainty of ± 400 nm. Until this noise level can be reduced, the high-resolution imaging techniques we have developed will still be needed for accurate positioning of the fibre.

Another potential improvement to our setup is the replacement of the lensed fibre with a fibre connected to a gradient-index (GRIN) lens. This lens would be designed to couple the numerical aperture of the fibre to that of the nanostructure. Conceived as a better method of collecting nanowire PL emissions (it would theoretically have a collection efficiency of 100%), this fibre and GRIN lens combination is currently being tested on nanowire samples in a 4 K cryostat, and the preliminary results appear very promising. The GRIN lens also has a focal length of approximately 200 μm , which is 20 times that of the lensed fibre. Such a large distance between the lens and the sample surface would significantly reduce the chances of accidental damage to the sample. Future upgrades to our imaging setup may therefore include the addition of a GRIN lens to allow for more accurate and efficient collection of nanostructure emissions.

²⁴These are called Look-Up Table files (extension .LUT). They are unique to the individual devices they describe, and we were able to obtain them after temporarily returning the positioners to the manufacturer.

Chapter 4

Fibre-Coupled Nanowire Quantum Dots as Sources of Single Photons

Various quantum technologies, from linear-optical quantum computing to quantum key distribution, require single, indistinguishable photons for proper operation [10–13]. As an efficient and reliable source of such particles has yet to become widely available, many past experiments have instead relied upon the use of highly-attenuated lasers, usually with a statistical 0.1–0.2 photons per laser pulse [68]. Progress has been made in recent years, however; photonic nanowire quantum dots [24, 69], nitrogen-vacancy centres in diamond [70], micropillar cavities [71], and self-assembled quantum dots [72] have all been investigated as true sources of highly-indistinguishable single photons. Nanowires in particular have demonstrated several useful ancillary properties, including controlled positioning of the wires [18], a single, well-formed emitter per wire [5, 15], high quantum and collection efficiencies [15, 73], and a Gaussian emission profile, which is ideal for coupling into optical fibres [6].

Most optical research in the field of nanowires tends to involve the use of a free-space objective to couple nanowire emissions to an external optical system [10]. In our experiments with these devices, we have created a fibre-only system for both exciting them with a laser and collecting their photoluminescence (PL) emissions. While others have employed optical cavities [74] and on-chip ridge waveguides [26] for coupling quantum dot emissions into fibres, our unique setup uses a lensed fibre for direct photon collection, as shown in Figure 4.1. Our intent has been to investigate this setup as a possible first step in creating a fibre-coupled, on-demand source of single, indistinguishable photons. Such a device would have many potential applications, including the generation of entangled photons for a fibre-based photon-to-spin quantum repeater network.

In this chapter, I will provide a description of our experimental setup and discuss the features observed in the PL spectra of single- and double-dot nanowires. Section 4.2 contains a review of our attempts to produce highly indistinguishable photons by reducing phonon interactions within the quantum dot. In Section 4.3, I present the results of our




Figure 4.1: Schematic of a fibre-coupled nanowire. (Fibre not to scale with wire.)

attempts to create a fibre-coupled source of single photons.

I would particularly like to thank Dr. Dan Dalacu of the NRC for providing the nanowire samples described in this chapter, for giving us access to the necessary equipment,¹ and for his assistance during our tests of these devices. The time he took to educate me about nanowires was greatly appreciated, and this chapter (not to mention our nanowire project itself) would not have been possible without him. Thank you also to Dr. Philip Poole and Dr. Robin Williams for helpful and informative discussions.

4.1 PL Measurements of Nanowire Quantum Dots

The first series of tests we performed on nanowire samples involved the observation analysis of their emission spectra. A lensed fibre was focused on individual nanowires and used to collect the PL photons. As a precursor to testing single-photon emission, the sources of the various peaks in the spectra were identified so that the quantum dot emissions could be isolated. Tests of both single- and double-dot nanowires were conducted, as their optical properties provide unique insights into the structure and composition of each type of sample. Understanding the characteristics of these nanostructures will be key to determining their potential usefulness in future quantum technologies.

¹This includes the spectrometer, four single-photon counters, and the HBT setup.

4.1.1 Structure and Layout of the Nanowire Samples

The samples we use in our experiments are photonic nanowires with quantum dots embedded near the base.² As depicted in Figure 4.2a), a gold droplet acts as a catalyst for the epitaxial growth of each nanowire’s InP core, and the fabrication conditions ensure a wurtzite (WZ) crystal structure. The temporary substitution of the P with As produces an InAs(P) quantum dot, and InP core growth resumes with the reintroduction of phosphorus. This process allows complete control of the placement and number of dots in the nanowire. Once the desired core height has been achieved, the growing conditions are altered, causing the wire to begin growing outwards instead. This radial growth forms an InP shell (also with WZ structure), which covers the core and the quantum dot (Fig. 4.2b). The difference in refractive index between the shell and the exterior environment causes it to act as a waveguide, which directs light along the length of the nanowire and prevents significant losses due to radial emission of the quantum dot [15].³ Fig. 4.2b) shows the taper of the waveguide, which is engineered to expand the photon mode from inside the nanowire to outside as it propagates towards the tip [6].⁴ The result is that for a typical taper of 1-2°, the nanowires have a numerical aperture (NA) of ~ 0.5 [16], and their emissions have a Gaussian profile (Fig. 4.2c), which is perfect for coupling into a single-mode fibre [6].

The nanowires in our samples are grown in a bottom-up technique, where there is full control of the nanowire positioning. The gold droplets are positioned in specific configurations, with the width of each droplet dictating the diameter of its respective nanowire core and quantum dot(s) [5, 15]. Taking into account the three processes by which indium reaches the gold catalyst,⁵ it has been shown that the rate of core growth is dependent upon its width, with narrower wires growing faster (and therefore taller) during a given period [18]. Since all wires on a sample are grown simultaneously, it is therefore the relationship between diameter and growing time which dictates the final height of each wire, as shown in Figure 4.3a).

During growth, the gold catalyst is limited in the amount of indium it can contain [75]. The catalyst’s In saturation levels ensure that the height of a wire’s quantum dot is actually *independent* of the core width [75]. When growing the InP layers of the core, the saturation limit is 45%, although the gold typically does not reach saturation during this process. The rate of growth of the InP layers therefore depends on the three mechanisms of

²Section 1.2 provides detailed information on the fabrication techniques and basic properties of these nanowires.

³Our samples are still susceptible to a 50% loss of quantum dot emissions, however, because half of the light escapes out the base of the nanowire, where it is lost to the substrate. The addition of some type of metallic mirror at the nanowire base, such as a layer of gold, has been shown to prevent this type of loss [15], although the wires we used were not equipped with such a reflective surface.

⁴Since the light mode has expanded well beyond the radius of the nanowire by the time it reaches the tip, the reflective properties of the gold particle at the top of the nanowire have little effect on the emission intensity or profile.

⁵These three mechanisms (direct, direct sidewall and indirect sidewall) are described in further detail in Section 1.2.2.




Figure 4.2: a) A nanowire core (red) with an embedded quantum dot (blue) near the base, and capped by a gold droplet (yellow). b) The waveguide shell covering the nanowire core. Nanowire schematics in a) and b) reproduced with permission of Dan Dalacu from [10]. c) High-resolution spatial image of the Gaussian emission profile from a single nanowire quantum dot. Circle indicates nanowire width, which has a diameter of 250 nm at its base.

indium collection in the catalyst. On the other hand, the saturation limit for InAs growth is just 35%. Upon switching to InAs, the indium in the gold will be at the level for InP growth, which is too high. The catalyst's In concentration must drop down to 35% before the growth rate becomes dependent on the external conditions again. During this period of indium concentration decay, which lasts about 10 seconds, the growth rate of the dot is constant. Given the duration of InAs growth for our samples is only 2–3 seconds [11, 15], the dots are always about 5–7 nm thick regardless of the width of their respective nanowire cores.

The width of a dot affects the wavefunctions of its trapped excitons, however, so our samples are fabricated with nanowires of various diameters. This allows us to study a variety of dot configurations. Figure 4.3b) shows the layout of our nanowire chips, which is the same for both the single-dot and double-dot samples. All numbers, letters and boxes in this image are formed of nanowires, thanks to the patterning of the sample's SiO₂ mask (see Section 1.2.2). In the set of boxes outlined by the red square, each box contains a single, well-formed nanowire at its centre. The box columns are labelled by their Dose numbers, which dictate the nanowire diameters. Dose 1 nanowires have a nominal core diameter of approximately 18 ± 2 nm and a height of 5 ± 2 μm , while each consecutive column adds two nanometers to the diameter. The Dose 1 wire has a waveguide shell with a diameter of ~ 250 nm. The lower left image in Fig. 4.3b) shows a close-up of several boxes. Note that the boxes themselves, which have dimensions of 20 μm by 20 μm , are composed of two rows of short, deformed nanowires. These wires typically have low-intensity PL emissions with many broad emission lines, so we do not use them for experiments. The wire at the

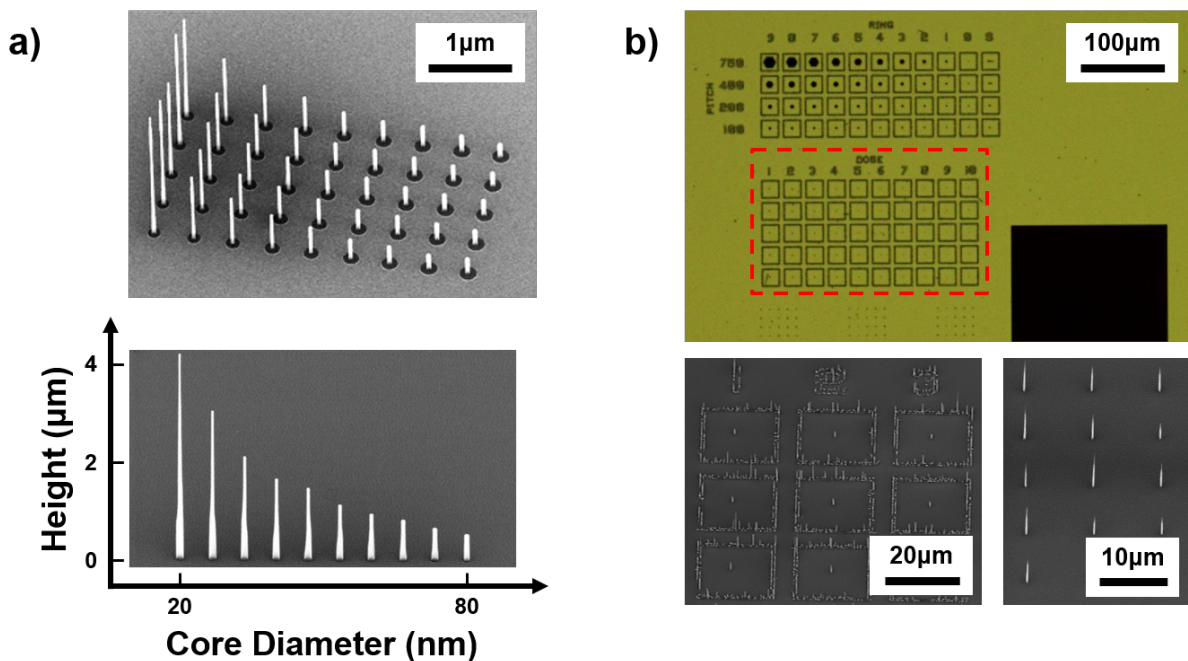


Figure 4.3: a) Relationship between a nanowire’s core diameter and its height. Reproduced with permission of Dan Dalacu from [5] (top) and [18] (bottom). b) Layout of the nanowire samples used in our experiments. Top: Microscope image of the sample, red rectangle highlights boxes containing individual nanowires. Bottom left: Close-up SEM image of some boxes (highlighted by red rectangle in top image), each of which contains a single nanowire at its centre. Bottom right: SEM image of a nanowire array with a 10 μm pitch. Sample images provided by Dan Dalacu.

centre of each box is sufficiently isolated that the lensed fibre can be focused on it without picking up emissions from the short wires forming the box edges.

For the single-dot experiments summarised in the following sections, most tests were done on wires in a large array elsewhere on the sample (lower right image in Fig. 4.3b). This array has a 10 μm pitch, i.e. a separation of 10 μm between the wires, and is made up of Dose 1 wires. As with the boxes, the distance between wires in the array is large enough that we are able to collect light from a single wire at a time, with no background emissions from its neighbours.

The nanowire samples used in our experiments were grown by Dan Dalacu at the NRC using the methods described in Section 1.2.2.

4.1.2 Experimental Setup used for Fibre-Based PL Collection

Observing the PL spectrum of a nanowire quantum dot requires both a pump to excite the wire and a method of collecting the PL photons. In our setup, both excitation and collection are accomplished using a single lensed fibre which is focused on the tip of the nanowire. While the emissions of the nanowire crystal itself may escape radially, all photons generated within the quantum dot will be directed along the length of the wire by the waveguide shell. Those photons which exit the top of the nanowire can then be collected by the focused lens of the fibre.

Figure 4.4 is a schematic diagram of our setup.⁶ A Ti:sapphire laser with tunable wavelength is coupled into a standard, single-mode optical fibre using a fibre coupling kit (Newport, 9131-M). The light is sent into the 1% port of a 980nm, 99-1 fibre beam splitter (OZ Optics, FUSED-12-980-6/125-99/1-3A3A3A-3-1), and from there, it travels to the dilution refrigerator, which houses the nanowire sample. Three piezoelectric positioners are used to focus the lensed fibre on a nanowire and deliver the laser light. Light emitted by the nanowire (or reflected by the sample substrate) is collected by the same fibre, which returns it to the beam splitter. The 99% arm of the beam splitter sends the light to a fibre-coupled spectrometer (Princeton Instruments, SP 2560i), where intensity measurements are made by a CCD camera (Princeton Instruments, PyLoN:100BRX). We use primarily 150 g/mm and 1200 g/mm spectrometer gratings⁷ when acquiring wide (~ 350 nm) and narrow (~ 30 nm) spectral ranges respectively, and the LightField software provided by Princeton Instruments is used to control the spectrometer.




Figure 4.4: Schematic of the setup used to collect PL emission from the nanowire samples. The three ports of the fibre beam splitter are labelled T, 1%, and 99% according to their respective functions.

While separation of the pump and PL light would be best accomplished by a circulator, the only circulator we had available at the time of the experiments exhibited polarisation-dependent behaviour, which introduced noise in the PL intensity. A 99-1 fibre beam splitter

⁶See Chapter 2 for further details about the various components of our experimental setup.

⁷These gratings are usually characterised by their groove density, which is measured in grooves per mm. The units are represented by the symbols: g/mm.

was therefore used to direct the pump laser and collected emissions. Its three ports are labelled T, 1%, and 99%, as shown in Fig. 4.4. The beam splitter has virtually no coupling between the 1% and 99% ports, so any laser light observed in the measured spectrum is entirely due to reflection off of the sample surface. The beam splitter is designed to be unidirectional, with light entering Port T and exiting the other two. Consequently, the laser entering the 1% port is highly attenuated, with a transmission factor of roughly 0.005 to Port T. Taking into account the $\sim 30\%$ coupling of the free-space beam into the fibre, as well as the various types of fibres used to transmit the light to the fridge, the total transmission from the free-space laser to the lensed fibre is approximately 0.05%. (By comparison, 30% of the light collected by the lensed fibre is transmitted to the spectrometer.) We find this attenuation of the laser to be helpful, given that overheating of the sample can alter the emission spectrum and melt the wire [67]. Laser power levels between 1 nW to 1 μ W at the lensed fibre were therefore used in most experiments.⁸

Locating the nanowires with the lensed fibre is accomplished using the methods described in Chapter 3. Our high-resolution imaging system, which requires only the piezoelectric positioners and the lensed fibre in our dilution refrigerator, allows us to systematically scan the sample for nanowires with usable spectra. A site-selective growth technique ensures that the nanowires are sufficiently spaced for the lensed fibre to isolate their individual emissions without capturing background light emitted by their neighbours. This spatial isolation is an important prerequisite for any source of highly-indistinguishable single photons [10], and analysis of the quantum dot emission spectrum can begin once the fibre is focused on a nanowire located by our imaging system.

Unfortunately, focusing of the lensed fibre is complicated by the fact that the ideal focus for pump delivery differs from the focus required for PL collection. In theory, the laser light should be focused on the location of the quantum dot, so as to maximize the number of excitons which are generated in its vicinity or within the dot itself. The ideal focal point for PL collection, however, is the tip of the nanowire, where the PL photons are emitted. Our technique has generally been to focus the fibre at the point where PL collection is greatest for a given pump power. Another option is to increase the pump power while focusing the fibre lens on the nanowire tip to enhance PL collection. A higher excitation power ensures that many excitons are generated within the wire, some of which will scatter through the crystal and become trapped within the dot.

Resonant pumping⁹ ensures that excitons are only generated within the dot. While this allows for greater indistinguishability of the PL photons, polarisation is often used to separate the reflected laser light from the dot emissions [10]. Unfortunately, our cryogenic setup currently does not have room for the additional optical components necessary to accomplish this filtering, and since optical fibres mix polarisation states, room-temperature separation of the pump and PL is also impractical. We therefore use above-band pumping of

⁸These values are estimates based on the various transmission ratios of the fibres, as well as comparisons of our pump power dependence data to free-space tests performed by others (as in Fig. 4.10c).

⁹Resonant pumping involves illuminating the sample at a wavelength equal to the dot emission wavelength.

the nanowires instead, with a pump wavelength of approximately 790 nm. This wavelength can be easily distinguished from the dot emission lines, which, in spite of the above-band excitation, are sufficiently narrow for our purposes.

The intensity of the reflected laser light collected by the fibre is far too bright for the CCD camera, however, and surpasses its photon counting ability. As overly-intense light can cause damage after prolonged exposure, we use a free-space, longpass filter with a cut-on wavelength of 850 nm to remove most of the reflected laser light before it reaches the grating in the spectrometer. The remaining light appears in spectral measurements as a narrow peak of reasonable intensity. The filter attenuates some of the nanowire's emissions at around 830 nm, but this peak is usually bright enough and close enough to the cut-on wavelength that it can still be seen in the spectral data.

4.1.3 Acquiring and Processing Nanowire Emission Spectra

The LightField software made by Princeton Instruments is used to acquire spectral data, with the light intensity, in photon counts per second, measured as a function of wavelength. It also controls the various settings of the spectrometer and CCD camera, such as the wavelength range and the exposure time of the camera per frame.¹⁰ Longer exposure times yield more counts, which is helpful for low-intensity sources. The measured intensity of quantum dot emission lines may vary wildly, from less than 10 counts per second to several tens of thousands of counts/sec. This variability is the result of several combined factors, including the unique properties of the quantum dot, the focus of the lensed fibre, and the pump power level. The noise level is usually about 20 counts/sec regardless for low exposure times,¹¹ so low-intensity PL emissions may be buried in the background. High-intensity emission lines, on the other hand, are well above the noise level, but their peak intensity might vary by a couple thousand counts/sec per frame. Therefore, for every data set where a detailed analysis of the spectrum is desired, we average several frames (usually >20) in order to reduce the background noise and obtain the average intensities of the peaks. In order to achieve the clearest results in the least amount of time, one must navigate the trade-off between exposure time and averaging. The best practice appears to be more averaging for low-intensity dots (to reduce the noise level) and more exposure time for high-intensity dots (to reveal low-probability events).

In addition to averaging, proper removal of the detector's dark counts can help reduce the background noise level. For each acquired spectrum, a dark count 'spectrum' is also obtained (Fig. 4.5a) with the same wavelength range, exposure time and number of averaged frames. If either the exposure time or averaging of the dark counts is not equivalent

¹⁰The word 'frame' refers to an acquired spectrum, especially when data is taken continuously. Consider the analogy of stop-motion frames captured by a video recorder. I will most often use the term 'frame' in the context of multiple spectral data sets acquired under identical conditions, usually for the purpose of averaging.

¹¹High exposure times and high pump powers may both cause the background noise level to increase.




Figure 4.5: a) Dark counts for a 28 s exposure interval. b) Average of 100 dark count frames, each with a 28 s exposure time.

to that of the main data set, there will be a mismatch between their noise levels. Compare, for example, the non-averaged and 100-frames-averaged dark count plots shown in Figure 4.5. If the non-averaged dark count frame were subtracted from a spectrum formed by 100 averaged frames, the resulting noise level would be quite large and might hide low-intensity features.

Aside from averaging, another of our post-processing procedures involves the elimination of spectral peaks caused by the passage of cosmic rays (CR) through the detector, an event which typically occurs at least once every few minutes. These peaks are extremely narrow (usually only one or two data points, even at highest wavelength resolution) and often fairly bright (thousands of counts/sec). They usually occur at varying wavelengths and only appear in a single frame, making them easy to identify. Figure 4.6a) shows a spectrum with three CR peaks. These particular CR lines have rather low intensities compared to most, likely because only part of the event was captured by the frame.

Due to the high intensities of most CR peaks, averaging many frames does not entirely eliminate them. We must therefore rely on two removal methods: discard affected frames or excise the CR peaks by hand.¹² Manual removal is accomplished by changing the offending data points to match the values on either side of them. This procedure is only performed for frames where the CR lines appear in empty sections of the spectrum (i.e. where there is only background noise); in order to avoid artificial alterations to the spectral features being studied, the frame must be deleted if a cosmic ray peak overlaps with a PL emission line. Averaging several data sets eliminates the minor aberration introduced by a manual correction, as seen in Fig. 4.6b). This averaging technique is valid because: 1. CR peaks rarely appear at the same wavelength across multiple frames, 2. most frames do not have

¹²The LightField software does come with a tool designed to automatically remove cosmic rays from each frame, but because our quantum dot emission lines are so narrow, it tends to remove them as well.

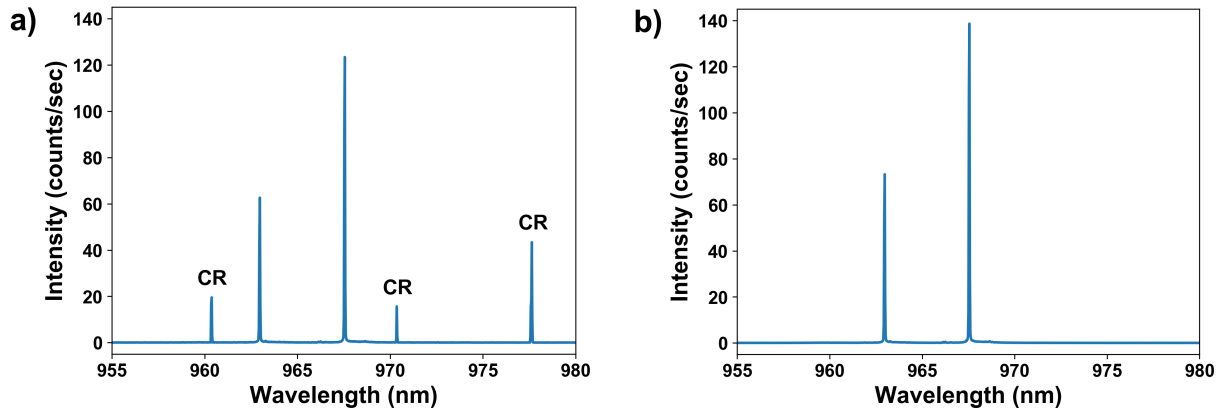


Figure 4.6: a) Quantum dot emission spectrum with three CR peaks. b) Same spectrum with CRs removed. To remove all trace of the CRs, several frames were averaged to produce this plot (hence the slightly brighter PL peaks in this version of the spectrum).

CR peaks (especially for small exposure times), and 3. CR peaks are only manually removed when they appear in regions of background noise alone.

4.1.4 Analysis of Single-Dot Spectra

The first type of nanowire sample we studied was fabricated with a single quantum dot in each wire. Preliminary analysis involved merely the identification of the sources responsible for the various peaks in the emission spectra of different wires. A typical spectrum for a single-dot wire is shown in Figure 4.7a), and the emission lines in this image have been labelled according to their sources. Table 4.1 provides a list of some common emission sources for InP/InAs nanowire quantum dot samples. Note that the wavelengths associated with sources are approximate, and some values will be different for each nanowire.¹³

If we look at Figure 4.7a), the peak with the lowest wavelength occurs at approximately 780 nm. This peak is produced by laser light which is reflected off the sample and recaptured by the fibre. Since we are using non-resonant pumping, the wavelength of the laser (which is tunable) is usually set somewhere between 780 and 810 nm during our experiments. The reflected laser appears as a very narrow peak, and it is heavily filtered to avoid damaging the CCD camera.

The next emission line, at 832 nm, is partially cut off by the laser filter. This peak is at the band gap energy of the nanowire crystal itself. Although both the nanowire and sample substrate are composed of InP, the band gap of the wire differs from that of the substrate due to their respective crystal structures. The nanowire has a WZ structure and emits

¹³The SiO₂ optical mask used for patterning the sample does not produce a peak because its band gap is much larger than the energy of the pump laser we use.

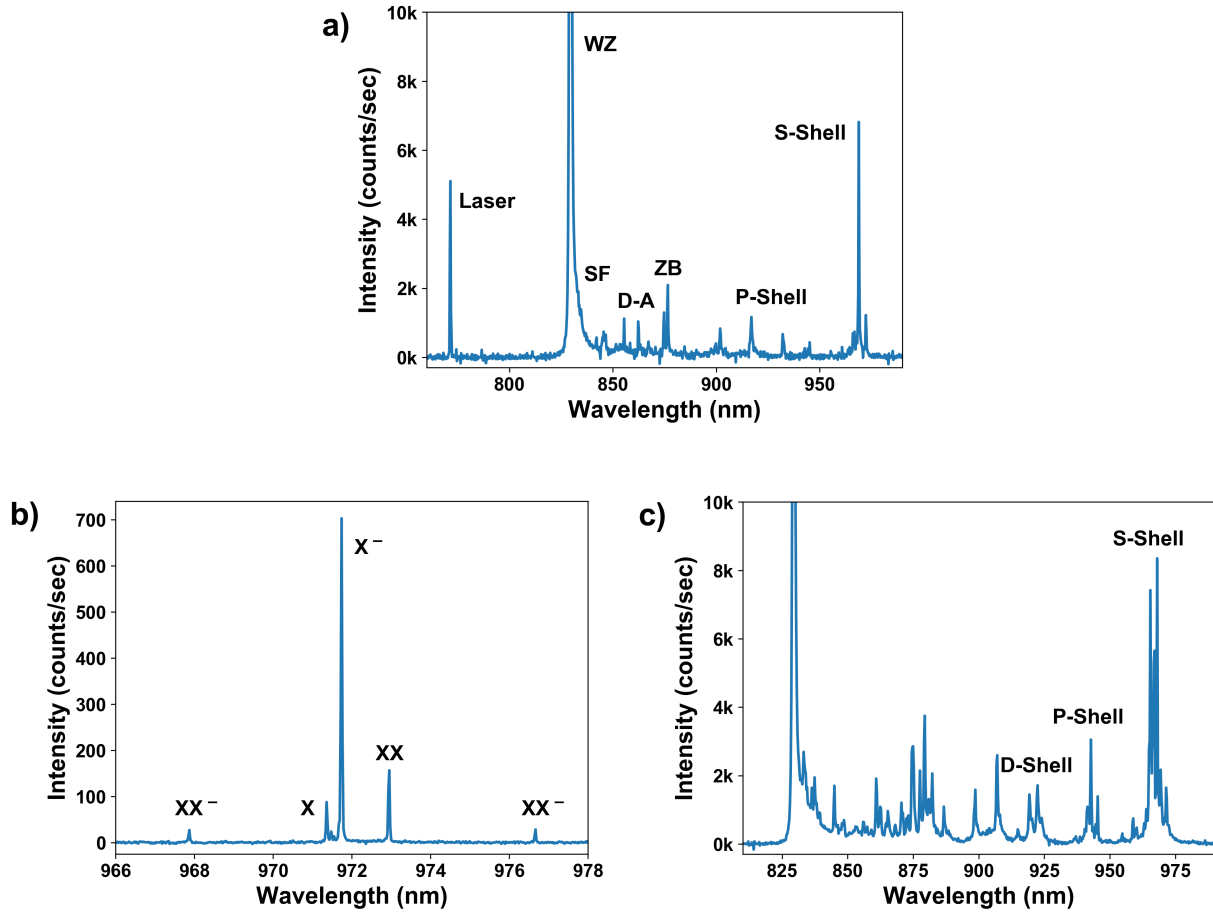


Figure 4.7: a) Example emission spectrum of a single-dot nanowire, with the various PL sources labelled. b) Quantum dot emission spectrum showing peaks produced by different exciton complexes. c) Emission spectrum of a nanowire when a high pump power is used. Note how noisy the spectrum becomes, with many peaks produced by higher-order excitations of the various sources.

photons with a wavelength of approximately 832 nm, while the zincblende (ZB) structure of the substrate emits at around 875 nm [10].

Additional peaks near 840 nm are caused by stacking faults (SF), which are undesired ZB layers in the WZ nanowire. InP has a ZB structure in bulk, but sufficiently thin InP nanowires grown under the appropriate conditions will have a pure WZ structure [11]. As the diameter of the wire increases (>20 nm), ZB insertions begin to appear [10]. These defects act as electron traps, so nearby electrons in the WZ will fall into the ZB layers, while holes remain in the WZ [11]. Stacking faults therefore cause electrons in the ZB conduction band to recombine with holes in the WZ valence band, as shown in Figure 4.8. The width of the ZB layer dictates the emission energy [76], so there are often many

Centre Wavelength (nm)	Source
790	Reflected Laser
832	Wurtzite band gap (nanowire)
840	Stacking faults in nanowire
860	Donor-acceptor band gaps
875	Zincblende band gap (substrate)
950	Quantum dot
> 1000	Large or deformed quantum dot

Table 4.1: Common sources of the peaks in a nanowire emission spectrum and the wavelengths at which they typically emit.

different peaks associated with stacking faults, usually at wavelengths very close to the WZ emission line. By trapping electrons near the dot, stacking faults create an unstable charge environment, which results in a broadening of the quantum dot emission lines (see Section 4.2.3 for further details) [11]. Fortunately, one can reduce the likelihood that stacking faults will be present by ensuring the nanowire core is sufficiently narrow during the fabrication process.¹⁴

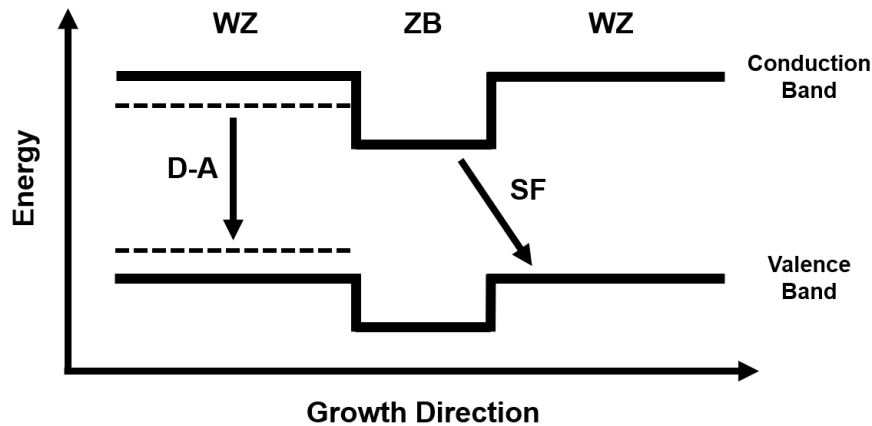


Figure 4.8: Energy band diagram showing emissions from two types of InP nanowire defects: stacking faults (SF) and donor-acceptor impurities (D-A).

¹⁴The nanowire shell generally grows outwards with the same crystal structure as the core layers, so the shell width is not a factor in the formation of SFs.

Another type of nanowire defect, which emits at around 860 nm, is a donor-acceptor (D-A) impurity in the WZ. D-A impurities produce shallower energy levels within the WZ band gap [10], as shown in Figure 4.8. Like SFs, these impurities act as electron traps and introduce noise into the charge environment. Arising during fabrication, they can typically be eliminated by using higher temperatures during core growth.

In our samples, the quantum dot emissions have wavelengths between 900 nm and 1000 nm, where longer wavelengths are generally associated with wider dots (i.e. higher Dose). The QD emissions are recognisable as a small collection of very narrow peaks spread over a range of >10 nm. Each peak is produced by the recombination of a different exciton complex. The s-shell complexes are the neutral exciton (X), the charged exciton (X^-), the biexciton (XX), and the charged biexciton (XX^-).¹⁵ Figure 4.7b) shows an example of a quantum dot spectrum with the various exciton peaks labelled.

As the intensity of the pump increases, small satellite peaks begin to appear all around the main peaks due to changes induced in the charge environment, such as the population of multiple nearby electron traps. Additional clusters of bright, narrow emission lines also begin to appear elsewhere in the spectrum. These are the higher-energy excitations of the quantum dot. Figure 4.7c) highlights the s-shell, p-shell and d-shell emission peaks. The p-shell can generally be identified by its appearance only at high pump powers, the fact that it appears at lower wavelengths than the s-shell, and its tendency to have far more peaks than the s-shell. Even higher energy levels may begin to appear as the laser intensity increases.

Peaks at wavelengths above 1000 nm are produced by wide or deformed quantum dots, such as those forming the letters, numbers or box edges in our samples (see Fig. 4.3b). We do not use these wires for experiments because their emission lines tend to be very broad and low-intensity.

4.1.5 Analysis of Double-Dot Spectra

In addition to single-dot nanowires, we have also studied double-dot samples. Figure 4.9a) shows a bare nanowire core with two dots spaced a few nanometers apart. The dots are separated by an InP spacer, with the width of the core determining the spacer length; narrower wires have longer spacers and vice versa, due to the relationship between the core diameter and growth rate (as described in Section 4.1.1). As the dots in our samples are clad in a waveguide shell and cannot be seen via SEM or transmission electron microscopy (TEM), the separations between the dots in the specific wires we studied are not known exactly. The approximate distances between dots can be obtained by comparing the core diameters and spacer growth time of our samples to those of unclad samples, as well as fitting to a mathematical model of nanowire growth [18, 77]. Based on this model, we

¹⁵The various exciton complexes are described in detail in Section 1.2.3. Section 4.1.6 explains the methods by which each peak in a quantum dot spectrum is identified with a given exciton complex.

estimate that the double-dot wires we examined had inter-dot separation distances of approximately 45 nm, 30 nm, and 20 nm, corresponding to a 10-second growth period for Dose 1, 2, and 3 wires respectively.




Figure 4.9: a) InP nanowire core with two InAs quantum dots and a ~ 5 nm spacer between them. Reproduced with permission of Dan Dalacu from [78]. b) Emission spectrum of a DQD showing the individual s-shells of the two dots. The high pump power also reveals the P-shells. c) Emission spectra of three double-dot wires as a function of dot separation: 20 nm (blue), 30 nm (orange), and 45 nm (red). d) The same three spectra as in c), but with the low-energy s-shells aligned. S-shells are highlighted by light blue boxes.

Figure 4.9b) shows the emission spectrum for a wire with a 30 nm separation between its dots. Note that the s-shell emissions of the two dots have different energies, which is true of all the double-dot wires we tested. In this particular instance, the dot's ground-state emissions are separated by 47 meV. The two dots are labelled the 'high-energy dot' and the 'low-energy dot' based on the relative emission energies of their s-shells.

In theory, as the energy of a PL photon is primarily determined by the geometry

and composition of the dot from which it was emitted, we should expect the two s-shells to overlap. For a given wire in our DQD sample, the dots should have nearly identical dimensions, but the arsenic content of the two dots is always subtly different. The first dot grown in the nanowire will typically have lower arsenic levels than the second, because ambient arsenic left over from the first dot is likely to still be present when the second dot is grown. The further away the dots are, the less arsenic will be present during growth of the second dot. As lower arsenic concentrations result in higher-energy photons, this means that the high-energy dot is likely to be the one which is closer to the nanowire base (i.e. the dot which was grown first), while the low-energy dot is the one closer to the tip [24, 77].¹⁶

An additional factor which is relevant to the ground state emission energies is the interaction between the dots. The dots are tunnel-coupled; free electrons and holes are constantly moving between them. Just as the charge environment affects the emission energy in a single-dot sample, the tunnel coupling (not to mention the presence of nearby excitons in a second dot) has important effects on the double-dot emission energies [79]. The coupling strength evidently depends upon the physical separation distance between the dots, with greater coupling for smaller separations.

The emission spectra for wires with dot separations of 20 nm, 30 nm, and 45 nm are shown in Figures 4.9c) and d). The greatest energy difference between s-shells, 94 meV, is observed for a spacer length of 20 nm. As the spacer length increases (and the coupling between dots decreases), we can see that the energy separation between the s-shells decreases. This phenomenon has been well-documented in stacked DQD systems [77, 79]. For a small spacer length L , the energy separation between the ground state emissions is expected to be proportional to $1/L^3$. The ground-state emissions are never identical, however, even when the dots are only weakly coupled. For the maximum spacer length of 45 nm, we still observe a 38 meV energy separation between the ground state emissions. This energy difference, which should not be present if the dots are identical and uncoupled,¹⁷ is due to the slightly differing geometries and compositions of the two dots, especially the difference in arsenic concentrations [77]. The arsenic content may also affect the energy separation when the dots are grown closer together and strongly coupled; the increased arsenic concentration in the second dot (i.e. the low-energy dot) causes its emissions to be even lower in energy, while the arsenic content, and hence the emission energy, of the first dot is unaffected by the separation distance.

The double-dot samples tend to have additional peaks between the dots' s-shells, as can be seen for the 20 nm separation in Fig. 4.9c). These peaks, which are often broad and low in intensity, are likely caused by tunnelling interactions, such as the recombination of

¹⁶The emission spectrum of a dot can be controlled during fabrication by altering the pressure of the AsH₃ [24]. It has been shown that increasing the arsenic content of a quantum dot can push its ground state emissions from 880 nm to about 1300 nm, although at the cost of emission intensity. Additional adjustments to the dot dimensions can yield emission wavelengths of up to 1550 nm.

¹⁷We will assume that the dots are entirely physically uncoupled for a 45 nm separation, although it may be that a longer spacer is needed to achieve true decoupling.

an electron and a hole which originate in separate dots. The p-shells of the two dots also appear nearby when the pump power is sufficiently high, as in Fig. 4.9b).

4.1.6 Pump Power Dependence and Lensed Fibre Collection Efficiency

Simple relationships exist between the power of the excitation laser and the relative intensities of the quantum dot's spectral peaks [11, 27, 80]. Consequently, a power dependence test can be used to identify the source of each emission line in the dot's spectrum, particularly the bright peaks produced by the various types of exciton complexes.¹⁸ With increasing pump power, more excitons and free carriers are generated in the sample, so certain types of recombination processes are more likely to occur.

In general, the exciton complexes tend to follow a simple trend as a function of pump power. At very low laser powers, there will be few excitons or free electrons/holes present in the wire, so only the X peak will be visible. With slightly greater pump power, the emission line of the negatively charged exciton (X^-) begins to dominate because there are more free electrons available to be trapped. Electrons tend to be trapped more often than holes because our samples are n-type, with donor impurities such as S, Si, and antisite defects (P replacing In). (The holes also have a greater effective mass, so they move more slowly through the nanowire.) Eventually, as the power continues to increase, the XX line becomes prominent as well, as it is now more likely that a new XX will form in the dot before the remaining exciton recombines. Still higher pump powers will of course produce a XX^- . Eventually, the dot's S-shell becomes saturated, and the p-shell and other higher-order processes begin to appear.¹⁹ There is a limit to the pump power which can be applied, however, as atomic intermixing is induced in the quantum dot when high laser intensities are used (i.e. the nanowire begins to melt). This can significantly alter the state of the system and change the emission energies [67].

Figure 4.10a) shows the change in the emission spectrum of a quantum dot as the pump power, P , increases. In this series of spectra, the X^- peak dominates between power levels of $P = 0.4 \mu\text{W}$ and $P = 2.1 \mu\text{W}$. The XX line starts to become more prominent at higher powers, while the intensity of X^- drops relative to X. Smaller peaks also begin to appear around the main three as higher-order transitions become more common.

To identify these three exciton complexes as the sources of these emission lines, we used the exponential relationship between the integrated counts of the peaks (i.e. the total intensity) and the pump power. In Figure 4.10b), the X, X^- , and XX peaks in a) were fit

¹⁸An analysis of power-dependent behaviour is not always sufficient to identify the exciton complexes, so other methods are often used as well. These include cross-correlation measurements [69], polarisation measurements [11], and magnetic field tests [80].

¹⁹While this is the general trend, the presence of electron traps and the bending of the Fermi level near the exterior edge of the nanowire may serve to alter the expected behaviour in some samples. This is why we do not rely on just the relative dot intensities to identify the exciton complexes.

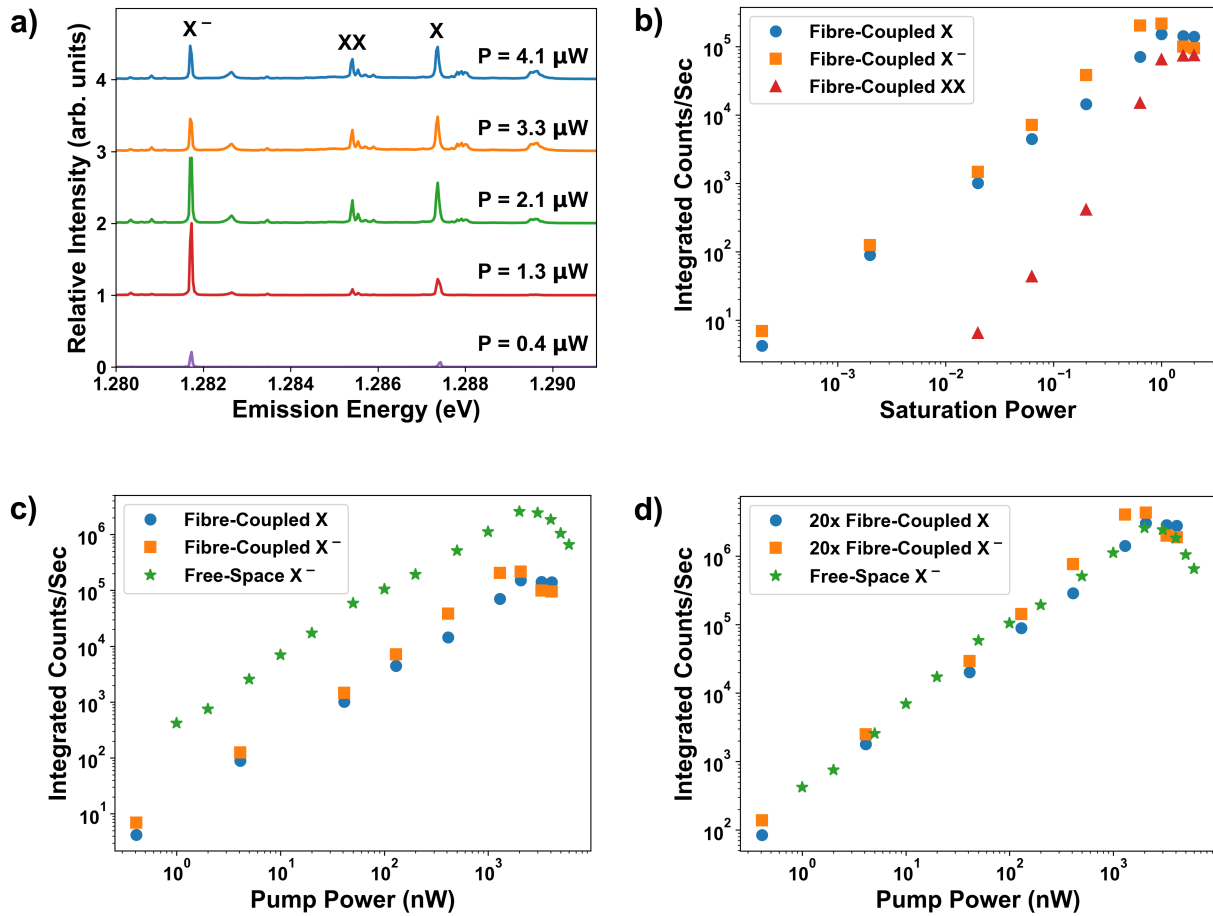


Figure 4.10: a) Quantum dot spectrum as a function of estimated pump power. Power levels are taken from graph c). b) Integrated intensities of Lorentzian fits to the X (circles), X^- (squares), and XX (triangles) peaks in a), plotted as a function of the saturation power. c) Integrated intensities of the X and X^- peaks in a), now with the saturation power scaled to match PL data collected using free-space optics in a 4 K cryostat (stars). d) Integrated intensities of the peaks in a), with intensity multiplied by a factor of 20 so that the results match the free-space data. Free-space data reproduced with permission of Dan Dalacu from [10].

with Lorentzians, and the integrated counts of the fits have been plotted as a function of pump power.²⁰ While the exact behaviour of each exciton complex varies from dot to dot,

²⁰For the integration, the number of counts is integrated as a function of data points, rather than wavelength. This is done for two reasons: the wavelengths assigned to data points by the spectrometer vary non-linearly over the measured spectral range, and it has been shown that when changing the centre wavelength of the spectrometer's grating, emission lines do not always appear at the same wavelength. So if the centre wavelength is changed from 950 nm to 930 nm, for example, a feature which was originally

the integrated counts of X and X⁻ are expected to exhibit a roughly linear relationship to the pump power, while the XX and XX⁻ are approx. quadratic [11, 80]. For the dot under examination in Fig. 4.10b), we can identify the exciton complexes by the slopes of the lines. The log-log slopes of X and X⁻ are quite linear, being about 1.21 counts/nW and 1.28 counts/nW respectively. The X⁻ is distinguished by its slightly larger slope, and the fact that its intensity is greater at higher pump powers. Additionally, the X peak should be slightly broader because it is actually a doublet; the emission line is formed from two peaks which can't be individually resolved due to the resolution of our spectrometer [11].²¹ Finally, the XX emission line is identified by its quadratic dependence on pump power; its log-log slope is 2.24 counts/nW.

The dot intensity eventually reaches a limit at high pump power, where there is a maximum probability that the dot will be filled by a given exciton complex; this point is called saturation [10]. In the case of Fig. 4.10b), all three peaks reached saturation at approximately the same pump power. The exact power of the laser at the sample is unknown, but the laser intensity was attenuated using filters of known optical density. The data is therefore plotted as a function of attenuated power relative to the saturation power, P_{sat} .

By comparing our data to that of samples tested in a free-space setup, we can (very roughly) estimate both the laser power at the sample and the collection efficiency of our lensed fibre. Figure 4.10c) shows the integrated counts of the X⁻ emission line from an identically-made sample.²² The spectrum of this dot was measured using free-space optics in a 4 K cryostat, where the PL photons were collected using an objective with a NA which matches or exceeds that of the nanowires (to ensure ideal collection efficiency). Assuming that the behaviour of the free-space dot is similar to that of the sample in our fibre setup, we can set the saturation power of the fibre to be equal to that of the free-space data.²³ With $P_{\text{sat}} = 2 \mu\text{W}$, the two data sets are plotted on the same axes in Figure 4.10c). This graph of the fibre data therefore represents the measured intensity as a function of the 'actual' pump power at the nanowire, with approx. 2-3 million integrated counts/sec at saturation.

Note that the log-log slopes of both the fibre and free-space data sets are nearly identical, but there are fewer counts in the fibre data. Figure 4.10d), however, shows how the two data sets overlap almost perfectly if we multiply the fibre counts by a factor of 20. This

peaking at 935.5 might now appear to peak at 935.3 nm, and its apparent width may be different because of the non-linearity in assigned wavelength values. The number of data points therefore provides a constant and absolute bin width for integration.

²¹The two peaks are associated with the recombination of the $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$ states, which are non-degenerate due to fine structure splitting (see Section 1.2.3). In spite of the low spectrometer resolution, they could still be distinguished using polarisation analysis [11], but we were not setup for this type of measurement. The X⁻ states, $|\uparrow\downarrow\uparrow\rangle$ and $|\uparrow\downarrow\downarrow\rangle$, are degenerate in the absence of a magnetic field.

²²Data was provided by Dan Dalacu.

²³This assumption may not be valid; the saturation power could actually be very different for the two dots. This approximation, however, should give us a somewhat better understanding of our fibre-based system.

result indicates that the fibre setup allows us to measure 5% of the PL emissions compared to the free-space setup.²⁴

We have two methods of estimating the collection efficiency of the lensed fibre itself. With the first (which is likely less accurate), we shall compare the fibre data to the free-space data. The throughput between the lensed fibre and spectrometer is 30%, while, due to its many mirrors, the free-space setup has a throughput of just 10% between the front of the objective and the spectrometer. We consider the first lens of the free-space objective to collect 100% of the nanowire emissions. Therefore, by comparing the free-space and fibre setups, and assuming identical quantum dot brightness in both cases, we estimate our fibre lens to have a collection efficiency of:

$$\alpha_F = 5\% \frac{100\% \cdot 10\%}{30\%} \approx 2\% \quad (4.1)$$

This value is low, but hardly surprising. The lensed fibres were originally selected and purchased with the intent of delivering light to a lateral quantum dot for photon-to-spin qubit transfer, not collecting nanowire PL emissions. Ideally, a system with wide-scale practical use should have a collection efficiency which approaches 100%.²⁵ The poor coupling between the fibre and nanowire may be partially due to the focus/alignment of the fibre, but the most probable cause is the lensed fibre having an NA of about 0.21.²⁶ The fibre therefore collects light over a total angular range of 24° , while the nanowire emits with a much wider divergence angle of about 60° . Much of the nanowire's emissions are lost due to this difference in NA, especially if the fibre is not aligned properly or is tilted with respect to the sample.

An alternative (and perhaps more accurate) method of estimating the collection efficiency of the lensed fibre is to calculate the maximum number of photons emitted by the quantum dot itself, multiplied by all the potential losses in the system, and compare to the data. For example, in Section 4.3.2, we will show that the average radiative lifetime of the charged exciton is $T_1 = 2.14$ ns. If we assume this to be the standard lifetime for all exciton complexes, then the total number of photons emitted per second by the recombinations of X, X⁻ and XX is about 470 million. (We are ignoring the negligible contributions of XX⁻.) This value presumes a continuous wave (CW) laser²⁷ pumping at saturation. Compare this

²⁴This assumes that the dots used in the two experiments are equally bright, and that the fibre was properly focused/aligned.

²⁵Note that we are only looking at the ability of the lensed fibre to collect the light emitted from the top of the nanowire. We have not considered here the 50% loss of photons to the substrate (an issue which could be resolved with the use of a metallic mirror), nor the coupling of the quantum dot emissions to the waveguide mode. We also omitted the spectrometer and CCD from this calculation, since both the free-space and fibre setups use them.

²⁶The actual NA of the fibre lens is not provided by the manufacturer, so this is a calculated value based on the $2.2 \mu\text{m}$ lens radius, the $8 \mu\text{m}$ focal length, and the $0.5 \mu\text{m}$ beam waist radius.

²⁷The data in Fig. 4.10 was actually obtained using a pulsed laser, but due to poor mode-locking and re-excitation of the dot, the pumping was effectively continuous. Section 4.3.3 will discuss this issue in more detail.

470 MHz total emission with the data in Figure 4.10b), which sums to just 310 kHz at saturation for the three exciton complexes. By accounting for known losses, we should be able to determine the collection efficiency of the fibre, which is the only unknown. We start with the approximately 95% coupling of quantum dot emissions into the nanowire’s HE₁₁ fundamental mode [10], and the loss of 50% of the photons to the substrate through the nanowire base. The optical fibres between the fridge and spectrometer were measured to have a 30% transmission ratio. The spectrometer itself has an 80% throughput (three mirrors and one grating), while the CCD camera has a detection efficiency of 50% [26]. These five factors add up to a throughput of 5.7%. The theoretical collection efficiency of the fibre is therefore:

$$\alpha_F = \frac{0.31 \text{ MHz}}{0.057 \cdot 470 \text{ MHz}} \approx 0.01 \quad (4.2)$$

This result of $\alpha_F = 1\%$ is in line with the value predicted by the first method. Unfortunately, a 1–2% collection efficiency for the lensed fibre is not sufficient for wide-scale use of this system. A common measure of the efficiency of a single-photon source is [10]:

$$\eta = \alpha\beta_{\text{HE}_{11}}\eta_{\text{IQE}} \quad (4.3)$$

where η_{IQE} is the internal quantum efficiency of the dot (which is nearly 100% for high-quality devices like ours), $\beta_{\text{HE}_{11}} = 95\%$ is the coupling of photons into the HE₁₁ mode of the waveguide, and α is the coupling of these photons to the external system, including the 50% loss of photons to the substrate. The value η therefore indicates the ability of the optical system to collect the photons emitted by the quantum dot. For our system, $\alpha = 0.5\alpha_F = 0.5\%$, which means that our source efficiency is $\eta = 0.48\%$. As might be for a free-space objective with nearly 100% collection efficiency, other groups have found source efficiencies as high as 43% for the same type of sample [69], improving upon our system by a factor of 100. Overall, efficiencies of at least ten percent are obtained for most free-space collection systems [10].

Further effort is evidently required to create a practical fibre-based system for single-photon collection. To improve collection efficiency, we are currently constructing and conducting preliminary tests on a fibre-based system which includes a gradient-index (GRIN) lens. The GRIN lens is designed such that it couples to the NA of the optical fibre on one side and to the NA of the nanowire on the other. By matching (or exceeding) the NA of the nanowire, the theoretical collection efficiency of the GRIN lens system should approach 100%, although the actual collection efficiency will likely be a slightly lower. (Preliminary tests indicate that it just barely improves upon the lensed fibre at the moment.) The GRIN lens will also have a much larger focal length than our lensed fibre (approx. 100–200 μm vs. 5 μm), reducing the likelihood of accidental sample damage during focusing procedures.

4.2 Indistinguishable Photons

It is of great importance that the photons used in various optical quantum computing applications be completely indistinguishable in all their properties, including wavelength [10, 13]. We have therefore conducted a preliminary investigation into the degree of photon indistinguishability which is achievable with fibre-coupled nanowire quantum dots in a dilution refrigerator. There are several factors which can affect photon indistinguishability, but our experiments focused primarily on the role of phonons in emission line broadening.

4.2.1 Lifetime-Limited Emission Linewidths

The majority of PL photons are emitted from a quantum dot with the bare recombination energy of an exciton, possibly modulated by the presence of an additional charge or exciton within the dot. For a given exciton complex, these photons will form a bright peak in the emission spectrum, called the zero-phonon line (ZPL). The width of the ZPL follows directly from the Heisenberg uncertainty principle, $\Delta E_i \Delta t_i \geq \hbar$, where ΔE_i is the uncertainty in the energy of the i^{th} energy level, and Δt_i is the uncertainty in the duration of occupation of the i^{th} energy level. For the purely radiative decay from excited state $|1\rangle$ to ground state $|0\rangle$, Δt_i is the lifetime, τ_1 , of the system in the excited state. Starting from the uncertainty principle, the *minimum* uncertainty in the energy of a photon emitted by such a process is therefore:

$$\Delta E_1 = \frac{\hbar}{\Delta t_1} \quad (4.4)$$

which means:

$$\begin{aligned} h\Delta\nu_1 &= \frac{\hbar}{\Delta t_1} \\ \Delta\nu_1 &= \frac{1}{2\pi\Delta t_1} \\ &= \frac{1}{2\pi\tau_1} \end{aligned} \quad (4.5)$$

As the smallest spread $\Delta\nu_1$ of the emission line is limited by the lifetime τ_1 of the excited state, this is called a lifetime-limited emission line.²⁸ Radiative lifetimes of excitons in quantum dots tend to be on the order of nanoseconds [81], corresponding to a linewidth

²⁸Technically, the width of an emission line is actually given by $\Delta\nu = \Delta\nu_1 + \Delta\nu_0 = \frac{1}{2\pi\tau_1} + \frac{1}{2\pi\tau_0}$. If we assume $|0\rangle$ to be the ground state, where no further decay is possible, then $\frac{1}{\tau_0} = 0$, and we arrive at the result of Equation 4.5.

(FWHM) of about 160 MHz, or $0.66 \mu\text{eV}$. An emission line produced by exciton recombination alone is as narrow as possible and has a Lorentzian shape. Any additional interactions with the environment serve only to broaden the emission line from its lifetime-limited state.

For an exciton with a 1 ns lifetime, the corresponding spectral linewidth of the ZPL is approximately $0.66 \mu\text{eV}$, or 0.5 pm (FWHM) for a central wavelength of 950 nm. Even with a 1200 g/mm grating, the resolution of our spectrometer is roughly $33 \mu\text{eV}$, or 24 pm at 950 nm, which is far larger than the width of a lifetime-limited emission line [11]. By fitting Lorentzians to the observed emission lines, we can obtain values which are closer to the true widths. The accuracy of these fits is somewhat questionable though, as we are fitting Lorentzians to perhaps four or five data points in the cases of the narrowest peaks observed. Other techniques, such as Fabry-Pérot interferometry (see Section 4.2.4), are therefore required to properly measure the linewidth.

4.2.2 Emission Line Broadening due to Phonon Interactions

Indium arsenide, which forms the dots in our nanowire samples, is a direct band gap semiconductor, so excitons within this material generally release just a single photon when they recombine. Occasionally, however, there occur more complicated recombination processes involving the scattering of phonons, specifically longitudinal acoustic phonons [12, 13, 82]. Thermal phonons present in a nanowire may be absorbed by a quantum dot at the same time as a trapped exciton recombines. The additional energy provided by the phonon results in the emitted photon having a slightly higher energy than usual. Conversely, a phonon may be emitted during exciton recombination, resulting in a lower-energy photon. These phonon interactions occur during time scales on the order of picoseconds, and they tend to have greater effects on excitons which are spatially smaller [13].

Up to 20% of photons emitted by the quantum dot may not appear in the ZPL due to these phonon scattering processes [10]. They instead form small lobes on either side of the main peak; these are the phonon sidebands, which are visible in the spectra of Figure 4.11. As the temperature decreases, the density of thermal phonons in the sample decreases, so the absorption sideband begins to disappear. The possibility of emitting both a phonon and a photon can never be entirely eliminated, however, so the emission sideband is still present at low temperatures [82].

The emission spectrum of a nanowire quantum dot at 4 K is shown in Figure 4.11a). Lorentzians have been fit to the ZPLs to accentuate the phonon sidebands and allow for comparison with purely radiative emission processes. The FWHMs of the fits are 12.5 pm (low-energy peak) and 29.0 pm (high-energy peak), indicating these emission lines may be nearly lifetime-limited, although spectrometer's resolution prevents the measurement of their true widths. The absorption sideband of each peak, which is slightly higher in energy than the ZPL, is largely absent at temperatures of a few Kelvin. The emission sideband very clearly remains, although it is also affected by the temperature. According to simulations which use the interaction Hamiltonian to describe an exciton in a quantum dot coupling to

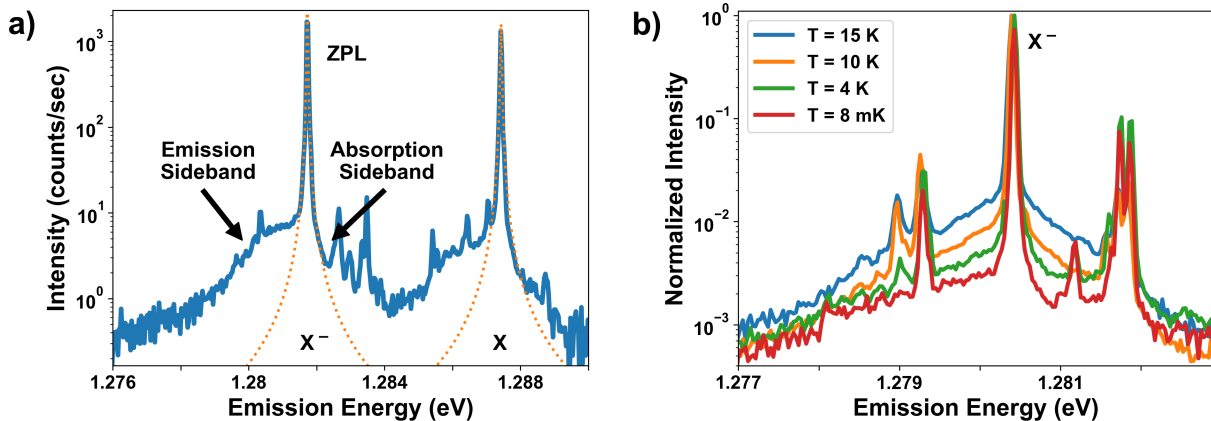


Figure 4.11: a) Example emission spectrum (blue) of a nanowire quantum dot at 4 K showing asymmetric phonon sidebands. (Note the logarithmic scale on the y-axis.) Lorentzian fits of the ZPLs (orange) have been included to demonstrate the broadening effects of the phonons. b) Temperature dependence of the phonon sidebands, showing greater asymmetry at lower temperatures.

external phonons, we should expect the emission sideband to be somewhat reduced as the temperature drops [13]. Experimental verification of the phonon sidebands' temperature-dependent intensities yielded the results of Figure 4.11b). The sideband intensities can be seen to decrease relative to the ZPL as the temperature drops, indicating that more of the photons are emitted in the ZPL instead.²⁹ The sideband asymmetry also gradually increases as the ambient phonons in the sample begin to vanish. We are therefore able to observe the results of phonon interactions in our system and reduce the number of interactions by shifting to lower temperatures. With the reduction in the phonon sideband intensity, more of the emitted photons are in the ZPL, which improves the overall indistinguishability.

Instead of changing the temperature, the sidebands can also be eliminated by using bandpass filters with appropriately steep cut-on/cut-off slopes. One still risks a reduction in emitter efficiency, however, as we shall see in Section 4.3.1. A third method of reducing the sideband intensities is to couple the quantum dot emissions into a microcavity prior to collection with the fibre. Although we have not tested this technique ourselves, it has been demonstrated that the cavity can actually assist in shifting the sideband emissions back into the ZPL, thereby achieving far greater indistinguishability [83]. Future solid-state sources of single photons may therefore be incorporated into optical cavities as a matter of course in order to improve indistinguishability.

²⁹To avoid accidental destruction of the nanowire, the fibre collecting the dot emissions had to be pulled away from the sample each time the temperature was changed (see Section 3.3.4). Each time the fibre was moved back into position, the focus and excitation power was necessarily different. The four data sets have therefore been normalized to their respective ZPL intensities, and we are analysing the change in the sideband intensities relative to each other and the ZPL at each temperature.

Finally, in addition to inelastic phonon scattering processes, another potential cause of reduced indistinguishability is the elastic scattering of phonons off the exciton, which can lead to pure dephasing through virtual transitions to higher-energy states [13, 84, 85]. Pure dephasing is the overlap of the wavefunctions of different states, for example the ground and first excited states, without an actual change in the state populations. The resulting time-dependent fluctuations in the wavefunction phases cause a loss of coherence and a broadening of the ZPL. As this dephasing mechanism depends upon the number of ambient phonons, it is temperature-dependent. It only begins to have its greatest effects on the ZPL linewidth at temperatures above 10 K [81, 85], which is partially why it is not visible in the data in Figure 4.11b). The resolution of our spectrometer, being far wider than the lifetime-limited width of the ZPL, is another likely reason.

4.2.3 Spectral Wandering due to a Fluctuating Charge Environment

As discussed in Section 1.2.3, the charge environment of the quantum dot heavily affects the energy of emitted photons. For example, the presence of an additional electron or hole, as in a charged exciton, can cause the emission energy to vary by several meV. Therefore, if the electric field around the dot is not stable in time, the result is a wandering of the emission energy and a broadening of the emission line.




Figure 4.12: Width of the X^- ZPL as a function of temperature, taken from two different single-dot nanowires. As the temperature drops, so too does the ZPL width. The spectrometer resolution prevents accurate readings of the ZPL width below about $35 \mu\text{eV}$.

Higher pump powers, higher temperatures and above-band excitation can all produce free charges which will move about within the nanowire. The presence of additional charges and excitons in the nanowire is one of the major causes of spectral wandering, and such

a fluctuating charge environment can cause the ZPL to wander over a period of several milliseconds [13], resulting in a wider measured emission line (unless one has a detector with millisecond detection resolution). Figure 4.12 shows the width of the X^- ZPL as a function of temperature. This data was obtained by fitting Lorentzians to the ZPL of the charged exciton for two different single-dot nanowires. The ZPL narrows as the temperature decreases, as the number of free electrons decreases and electron traps are filled. It appears, however, that the resolution of our spectrometer ($\sim 35 \mu\text{eV}$) prevents detection of this effect at temperatures below ~ 4 K. It should be noted that phonons also have an important role to play in emission line broadening at higher temperatures, as their interactions within the dot tend to cause the excitons to decay faster [4], thereby increasing the linewidths in accordance with Equation 4.5.

To reduce the number of free electrons and phonons moving through the nanowire, our samples are studied at temperatures of a few Kelvin or even a few millikelvin. The relationship between pump power and ZPL linewidth (called power broadening) has been well documented [69], and we attempt to resolve this issue by using the lowest possible pump power which still produces visible peaks. Although we are not set up for it, resonant pumping also reduces spectral wandering by exciting the quantum dot directly without generating thousands of additional excitons in the nanowire.

Impurities and defects in the wires can act as charge traps which collect free electrons/holes and alter the electrostatic environment on longer time scales. Some of the small satellite peaks which appear around the main emission lines at higher pump powers indicate the population of electron traps near the dot. It is for this reason that we attempt to work only with wires which have clean spectra devoid of the extra peaks associated with wire defects. Stacking faults alone can broaden an emission line by a factor of 10-100 [10, 11], although these may be easily eliminated by increasing the growth temperature during fabrication.

4.2.4 Quantifying Photon Indistinguishability

There are several experiments which can be conducted to quantify the degree of photon indistinguishability. Unfortunately, we were not able to perform these tests on our samples due to a lack of appropriate equipment and time. For the sake of completeness, however, I will briefly review a few of these experiments here.

The most commonly-used test of indistinguishability is Hong-Ou-Mandel (HOM) interferometry. The HOM effect is a two-photon interference phenomenon which is achieved with the use of a 50/50 beam splitter. In particular, this type of experiment requires a beam splitter which produces a phase shift of π when light is reflected off one side, but not the other.³⁰ Two detectors record the transmission/reflection of photons incident on the splitter. The result of the phase shift induced by one side of the beam splitter is a

³⁰One beam splitter configuration which can achieve this is a single glass plate with a semi-transparent dielectric along one side. When light travelling through the air reflects off the dielectric, a phase shift of

destructive interference effect: if a photon is simultaneously directed at each side of the splitter, both photons will always exit on the same side, with a 50/50 chance of detecting them both on one side or the other. For this to be the case, however, the two photons must be identical in every respect, including polarisation, arrival time, and wavelength. The number of coincidence counts of the detectors can therefore be used as a measure of indistinguishability. While we were not able to perform this test, HOM interferometry has been conducted on these nanowire samples by another group, and they were able to demonstrate a decently high degree of indistinguishability (83% HOM visibility with a lot of post-selection, though only 4% without post-processing) [69].

To measure the coherence length of the PL photons,³¹ we attempted to perform Mach—Zehnder interferometry. In this experiment, a beam splitter divides the light and sends half it on a path which includes a delay line. The two beams are then merged at a second beam splitter, and the resulting interference fringes are observed as a function of delay length. We should expect to see the fringe amplitudes decay as the delay line exceeds the coherence length, with a large coherence length indicating a high degree of photon indistinguishability. We were not able to accomplish this measurement due to a number of factors. The delay line equipment available to us had only a manual dial to control the path difference, the device provided no indication of the actual path difference, and the resolution of the dial was too poor to see the fringes. There were also continuous fluctuations in the output intensity, even when the path difference was not being changed. It is suspected that very slight vibrations and temperature fluctuations in the delay line and beam splitters (which were fibre splitters) were causing the delay length to change subtly on very short time scales. To properly perform an interferometry experiment such as this, the interferometer would need to be placed in a temperature-controlled box with some form of vibrational dampening. The transmission of the light via fibre also complicated the issue, as waveplates were needed to compensate for the fibre’s effects on the polarisation. If proper interference is to occur, the light in each path must be identical in polarisation when it reaches the second splitter, a requirement which was made excessively difficult by our use of fibre beam splitters.³² A free-space version of the experiment would likely be more successful. Past interferometry measurements of these nanowire devices, conducted by other groups, have demonstrated coherence times as low as 1.2 ns at 300 mK [69].

A final measure of photon indistinguishability is obtained by measuring the true ZPL

π is produced because the dielectric has a higher index of refraction than air. If the light is transmitted through the dielectric, there is no phase shift. Likewise, light which enters the glass side of the beam splitter does not experience a phase shift if it is transmitted through the dielectric. Importantly, light incident on the glass side of the dielectric is also *not* phase shifted if it is reflected. This is because the dielectric is a substance of lower refractive index than the glass through which the light is travelling prior to reflection.

³¹i.e. the distance over which they remain coherent, as measured by the ability of two beams to produce interference fringes as the difference in path length increases

³²Optical fibres alter the polarisation of the light they transmit, but we chose to use fibre splitters so as to avoid translating back and forth between fibre and free-space, which would have significantly reduced the intensity of the light reaching the photodiode. Had our delay line device actually functioned as desired, we might have switched to free-space to resolve the polarisation issue.

linewidths of the emission peaks. Due to the timing jitter and the poor resolution of our interferometer, the linewidths observed in our spectral data are not perfectly accurate. A Fabry-Pérot interferometer (also called an etalon) is an optical cavity which consists of two partially-reflective surfaces. Light only passes through the cavity when it is in resonance, so by measuring the intensity of transmitted light as a function of mirror separation, we can determine the true emission linewidth. We did attempt Fabry-Pérot interferometry on the PL emissions of our nanowire dots, but due to time constraints and difficulty aligning the interferometer, we were unable to complete this measurement.

4.3 A Fibre-Coupled Source of Single Photons

One of the main goals of our work with nanowire quantum dots has been to demonstrate that we can use our system as a fibre-coupled source of single photons. These particular samples have previously been shown to emit individual photons [6, 11, 15, 24], so our intent was to determine whether high single-photon counts could still be detected with direct coupling of a nanowire to a lensed fibre. A Hanbury Brown-Twiss (HBT) auto-correlation experiment was therefore used to verify single-photon emission in our system. If successful, our setup could form the basis for a future source of on-demand, fibre-coupled single photons.

4.3.1 Preparing for Hanbury Brown-Twiss Interferometry

The HBT experiment consists of a source of photons, a beam splitter (BS), and two single-photon detectors (D1 and D2). Photons from the source are sent to the beam splitter, which directs them to either D1 or D2. The correlation between photon arrival times at the two detectors is measured and used to determine the statistical likelihood of single-photon emission. If D1 and D2 consistently make simultaneous detections, this indicates multi-photon emission. On the other hand, if simultaneous detections never occur, the emitter must be a source of single photons. Figure 4.13 contains a schematic of the HBT experiment. A more detailed description of the experiment and the associated mathematics is provided in Section 1.2.4.

The first step in demonstrating single-photon emission requires the isolation of a single peak in the quantum dot spectrum, with all other light filtered out [10]. This is to eliminate photons produced by uncorrelated sources, such as those from the sample substrate or the WZ structure of the nanowire [13, 27]. By design, our samples are fabricated such that the exciton complexes in the quantum dot have non-overlapping emission lines at cryogenic temperatures [10], making spectral isolation a relatively simple affair.

In the case of the double-dot spectrum in Figure 4.14a) (only the low-energy dot's emission lines are shown), the X^- peak at 942.7 nm was selected for tests, as it was the brightest. Figure 4.14b) shows the setup used to isolate the peak, where a fibre-to-fibre




Figure 4.13: Setup used for a Hanbury Brown-Twiss autocorrelation experiment, with a nanowire as the photon source. If photons are never detected at D1 and D2 simultaneously, then the source must be emitting a single photon at a time.

U-bench (Thorlabs, FBP-B-FC) provided a free-space break in the fibre for two bandpass filters (Iridian Spectral Technologies, custom notch filter). The centre wavelength of each filter is 950 nm when the beam is normal to its surface, but adjusting the angle of incidence alters the notch wavelength. The filter bandwidth is approximately 2 nm (FWHM), which is too wide for a single filter to perfectly isolate the 0.3 nm-wide emission line. Consequently, a second filter was also used. By manually twisting the posts on which they were mounted, one filter was adjusted so that it cut off the higher-wavelength satellite peaks in the vicinity of the 942.7 nm line, while the other filter eliminated the lower-wavelength peaks.³³ Figures 4.14c) and d) show the effects of the two filters on the spectrum. Angling the filters in opposite directions helped to reduce beam displacement, so only minimal realignment of the U-bench was required after the filters were installed.

Unfortunately, the transmission profile of these filters does not have a sharp edge on the high-wavelength side (i.e. the cut-off slope is too shallow). The intensity of the desired peak is significantly reduced if we attempt to perfectly align the filter cut-off with the high-wavelength side of the ZPL. To maintain sufficient intensity, we initially left a slight shoulder on the X^- peak, as shown in Figure 4.14c). Through careful adjustments to the alignment of the U-bench and the filters, we were eventually able to perfectly isolate the peak, although the emission intensity was reduced by a factor of 2.5 (Fig. 4.14d).³⁴ With just the one peak present in the spectrum, we were then able to perform auto-correlation experiments to verify single-photon emission. Future tests of this all-fibre single-photon

³³This procedure was used only for the all-fibre HBT experiment with the double-dot samples. When performing the experiment with a free-space BS and a single-dot wire, we were able to use a single bandpass filter because the dot in question had emission lines which were sufficiently separated and bright at low pump power. (Low pump powers are used to prevent the appearance of small satellite peaks around the main emission lines.) Higher pump powers were required for the all-fibre tests (possibly due to the particular quantum dot under observation), which resulted in the presence of the small satellite peaks.

³⁴The full spectrum of the dot in Fig. 4.14a) was acquired with a lower pump power than the data in c) and d). This was done so that we could acquire a clean spectrum devoid of many satellite peaks. With the filters in place, however, the satellite peaks were no longer an issue, and a greater pump intensity was required in order to get the counts needed to perform the HBT experiment.

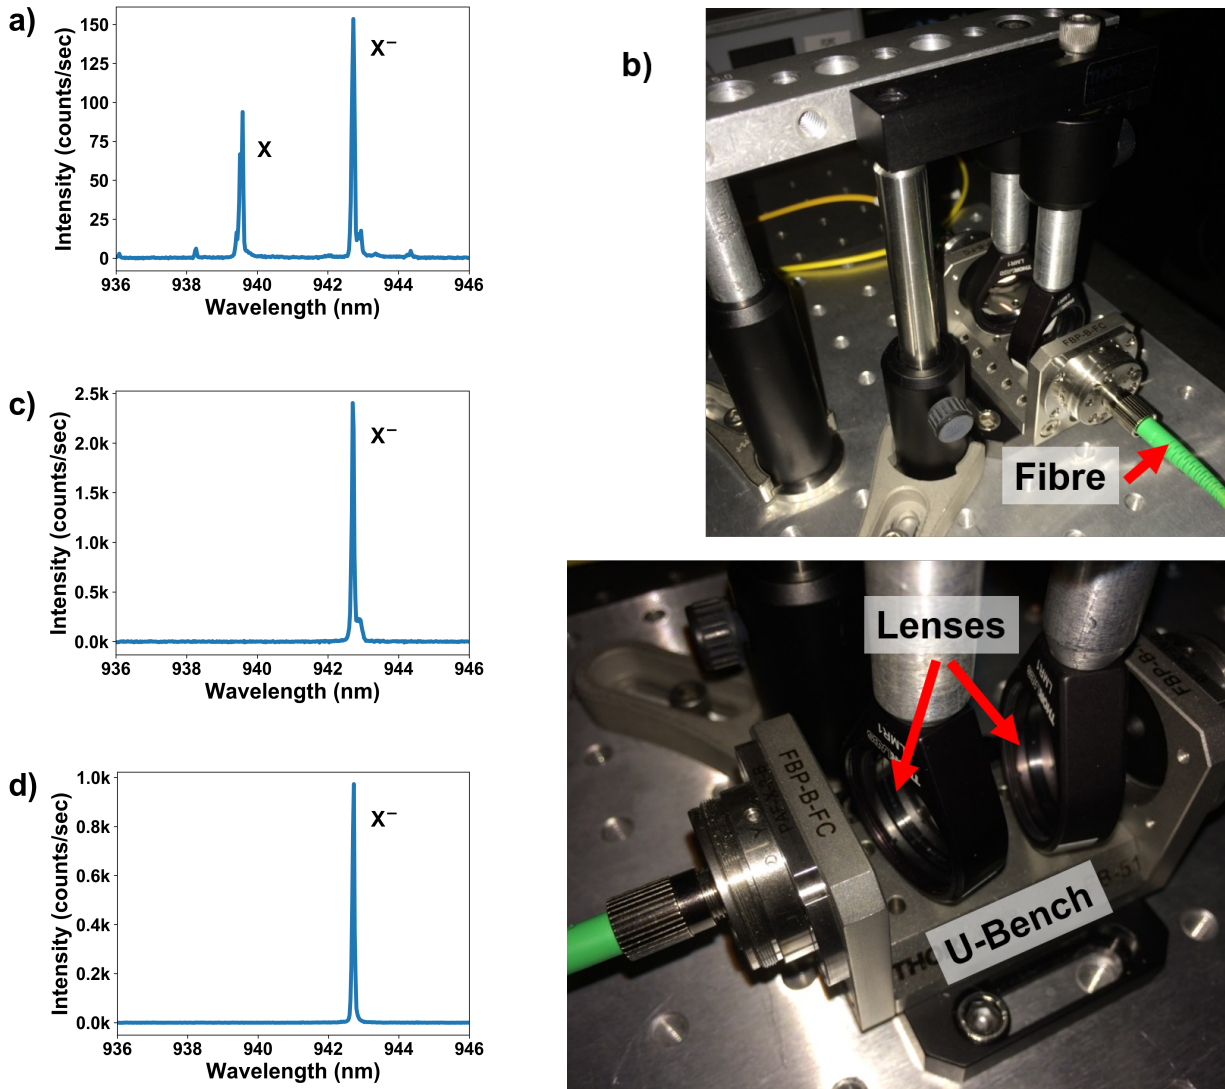


Figure 4.14: a) Unfiltered spectrum of the low-energy dot's ground-state emissions. b) U-bench setup with two bandpass filters to isolate the X^- emission line at 943 nm. c) First attempt at isolating X^- peak: a shoulder remains on the high-wavelength side. d) Second attempt at isolating X^- peak: ZPL almost perfectly isolated.

source will likely employ a bandpass fibre optic filter instead of free-space filters, which were the only type we had available at the time of our experiments.

4.3.2 Continuous Wave Hanbury Brown-Twiss Measurement

CW auto-correlation measurements were performed using a fibre-coupled, handheld laser diode source (OZ Optics, FOSS-21-3A-5/125-810-S-1) with a wavelength of 810 nm. To

reduce the likelihood of multi-photon emission, the laser intensity, which cannot be tuned on this source, was adjusted by pulling the fibre slightly out of the mating sleeve. While not an overly-precise method of intensity control, it achieves the desired result by reducing the coupling of the laser into the fibre mode. The laser light reaching the lensed fibre in the fridge had an estimated power of about 100–200 nW. The CW laser was used to excite a double-dot nanowire, where only the X^- emission line, the brightest peak in the spectrum, was isolated for use (as described in the previous section). This was done to reduce the chances of detecting uncorrelated emissions from the surrounding nanowire material. PL photons were collected by the lensed fibre, and a 50/50 fibre beam-splitter (OZ Optics, FUSED-12-980-6/125-50/50-3A3A3A-3-1) directed the nanowire emissions to two fibre-coupled single-photon counting modules (Excelitas Technologies, SPCM CD3700H), which were used as our single-photon detectors. A driver for single-photon detectors (Single Quantum, Argos 410) was used for correlating the timing of detections made by these Si avalanche photodiodes (APDs).

Figure 4.15a) shows the resulting HBT data for the X^- emission peak. A delay line is used to shift the ‘zero delay’ to 25 ns, allowing us to see the symmetry of the plot.³⁵ The number of counts at zero time delay is very low, indicating single-photon emission, but a quantitative description of the probability of multi-photon emission requires a calculation of the $g^{(2)}(0)$ value.

The empirical equation for the second-order correlation function is given by [26, 86]:

$$g^{(2)}(\tau) = 1 - a \exp\left(-\frac{|\tau|}{T_1}\right) \quad (4.6)$$

where T_1 is the average lifetime of the exciton complex before recombination, and $a = 1/n$ (n is the average number of photons in the photon stream at any given time). This function does not take into account the timing jitter of the single-photon detectors, which will serve to increase the observed $g^{(2)}(0)$ from its actual value. The experimental data is in fact a convolution of Equation 4.6 and the Gaussian-shaped timing resolution of the APDs:

$$\begin{aligned} g_{\text{Exp}}^{(2)}(\tau) &= g^{(2)}(\tau) * \Phi(\tau) \\ &= \left[1 - \frac{a}{2} \exp\left(-\frac{|\tau|}{T_1}\right)\right] * \left[\frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{\tau^2}{2\sigma^2}\right)\right] \\ &= \int_{-\infty}^{\infty} g^{(2)}(\tau - t)\Phi(t)dt \end{aligned} \quad (4.7)$$

In this equation, $\Phi(\tau)$ represents the Gaussian timing jitter of the APDs, and σ is its standard deviation. The APD specifications state the FWHM of the timing resolution to

³⁵See Appendix E for further details about the delay line and an explanation of why we observe ‘negative’ time delay in these plots.

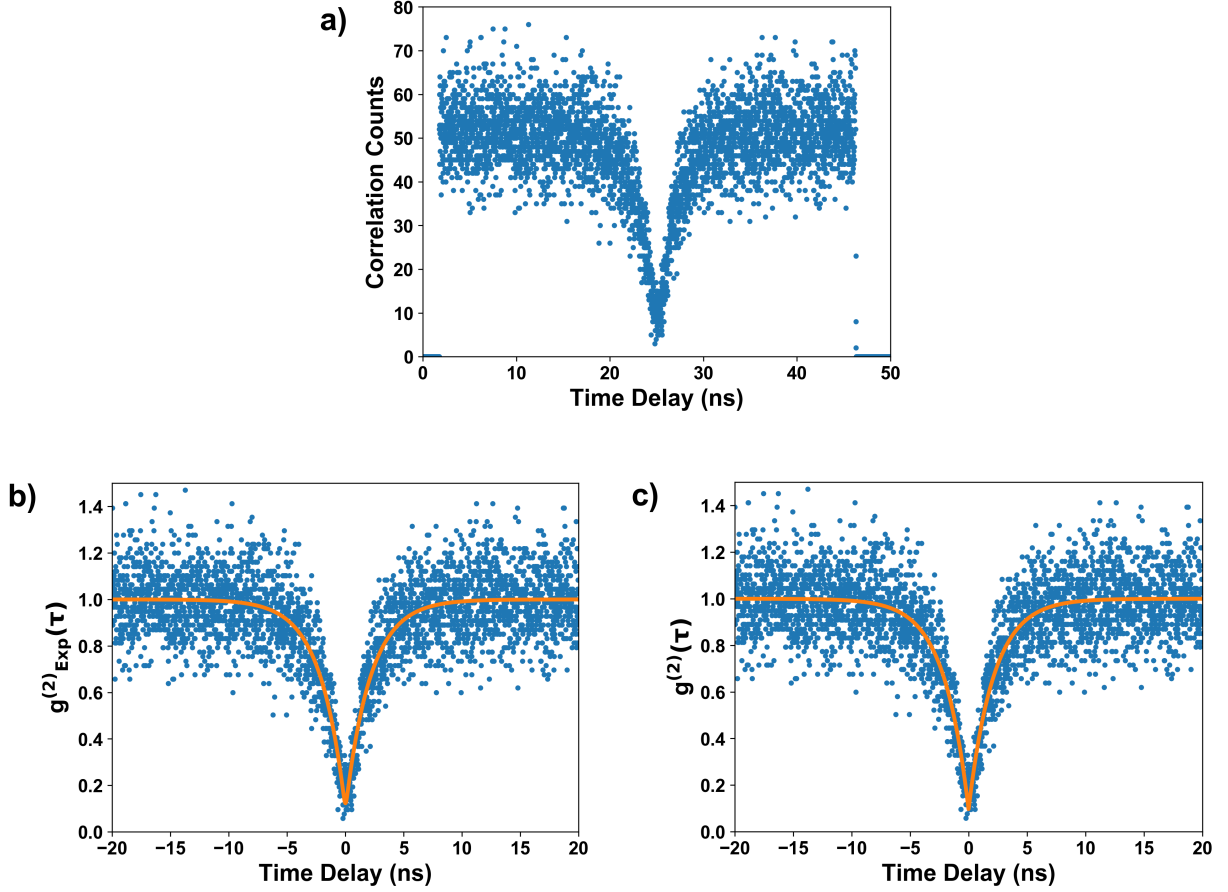


Figure 4.15: a) Experimental HBT data for a double-dot nanowire sample, with a CW pump and a 15-hour exposure period. b) HBT data with a fit to Equation 4.8. This fit is the convolution of the actual second-order correlation function and the Gaussian timing jitter of the detectors. c) Experimental HBT data with the deconvoluted $g^{(2)}(\tau)$ fit. The fit gives $g^{(2)}(0) = 0.095 \pm 0.027$, which proves single-photon emission in the quantum dot's X^- emission line.

be approx. 200 ps, which corresponds to $\sigma = \text{FWHM}/2\sqrt{2 \ln 2} = 84.9$ ps. If the data is not centred about zero time delay, as in Fig. 4.15a), τ can be replaced with $\tau - t_0$ in the fitting equation. Evaluating the convolution integral, we arrive at the following after several lines of math:

$$g_{\text{Exp}}^{(2)}(\tau) = 1 - \frac{a}{2} \left\{ \exp\left(\frac{\sigma^2 + 2T_1\tau}{2T_1^2}\right) \left[1 - \text{erf}\left(\frac{\sigma^2 + T_1\tau}{\sqrt{2}\sigma T_1}\right)\right] + \exp\left(\frac{\sigma^2 - 2T_1\tau}{2T_1^2}\right) \left[1 - \text{erf}\left(\frac{\sigma^2 - T_1\tau}{\sqrt{2}\sigma T_1}\right)\right] \right\} \quad (4.8)$$

By fitting this equation to the data, as in Fig. 4.15b), we can extract the parameters a and T_1 , allowing us to calculate the real $g^{(2)}(0)$. From the fit, we obtain $a = 0.90 \pm 0.01$. This means that the average number of photons in the stream at any given time is $n = 1.11$, indicating single-photon emission and multi-photon emission probability of just $\sim 10\%$. We also find that the X^- lifetime is $T_1 = 2.14 \pm 0.05$ ns, which is similar to the values others have observed for these samples [10, 69]. Using these two parameters, the deconvoluted second-order correlation function (Eq. 4.6) is plotted with the data in Figure 4.15c), where it can be seen to approach marginally closer to zero correlation counts than the convoluted fit in b). We calculate $g^{(2)}(0) = 0.095 \pm 0.027$ when using the fitted parameters. This $g^{(2)}(0)$ value is below the single-photon benchmark of 0.1, which is necessary for the single-photon source to be more experimentally useful than a highly-attenuated laser [28]. Most quantum light sources based on nanowires have managed to achieve $g^{(2)}(0)$ values of 0.1–0.3 using free-space collection-optics, but only a few varieties have been proven to reach $g^{(2)}(0) < 0.1$ [10]. The results shown here demonstrate that our InP/InAs samples can be used as part of a practical fibre-coupled source of single photons.

Other CW tests of these samples using free-space collection optics have demonstrated $g^{(2)}(0)$ values as low as 0.003 [10, 11]. We may partially attribute our significantly larger $g^{(2)}(0)$ to background counts, which we did not subtract from these calculations. Leakage of background light into the detector causes unwanted correlation counts, raising the probability of multi-photon detections, but subtracting background counts from a CW correlation plot requires additional data sets to be taken. Removal of correlation counts from the background alone, without nanowire emissions, do not significantly effect the $g^{(2)}(0)$, but we did not measure the instances of correlation between background and PL photons. The actual $g^{(2)}(0)$ of our dot's X^- emission line, without background noise, may therefore be marginally better than the value calculated here.³⁶

The stability of our setup in the dilution refrigerator allows us to perform HBT measurements over long periods, which is necessary due to our poor PL collection efficiency. The data in Figure 4.15 was acquired over a period of 15 hours. Another data set, acquired with greater excitation power, achieved twice the number of correlation counts after just 4.5 hours, but at the cost of the $g^{(2)}(0)$ value, which rose to nearly 0.5.

4.3.3 Pulsed Hanbury Brown-Twiss Measurement

The ability to emit photons on demand would be a major feature of any practical single-photon source. We can imagine, for example, that a short laser pulse could generate a single exciton within the quantum dot, yielding a single photon upon recombination. In our experiments, we attempted to determine whether our system could be used as a fibre-coupled source of on-demand single photons, with a laser pulse acting as a trigger for the emission of each photon.

³⁶See [86] for details on the proper normalization process of a CW auto-correlation data set; we merely normalized using the $g_{\text{Exp}}^{(2)}(\tau)$ fit.

HBT measurements can be done with either a CW laser or a pulsed laser, both of which yield a value for the $g^{(2)}(0)$. The pulsed experiment provides slightly more information, however, especially about background counts and the likelihood of re-excitation [10]. Re-excitation is the process whereby the dot is excited for a second time during a single pulse of the laser, resulting in two consecutive PL photons being emitted for a single pulse.

Pulsed measurements of the nanowire quantum dots were made using a Ti:sapphire laser, which has a pulse period of 12.4 ns. In theory, the pulse width should be approximately 1.2 ps, but due to poor mode-locking of the laser at the time of our experiments, the actual pulse width was likely closer to 100 ps.³⁷ Unfortunately, this issue affected our ability to prove that we could achieve on-demand single-photon emission. The data shown in this section was acquired at two different times. The first instance, shown in Figure 4.16a), involved a single-dot wire and two free-space APDs (Excelitas Technologies, SPCM CD3492H). The second instance (Fig. 4.16b) occurred several months later, when we used a double-dot sample and fibre-coupled APDs (the same ones as were used for the CW data presented in the previous section). The ‘zero delay’ point is offset in both plots through the use of a delay line (see Appendix E).

The pulsed HBT data shows multiple peaks, each produced by a laser pulse. At the centre is the ‘zero delay’ point, which is flat in the case of single-photon emission. The appearance of a central peak indicates multi-photon emission, the probability of which is correlated to the peak amplitude.³⁸ Comparing the two data sets in Figure 4.16, we observe that while the plot in b) is flat at the centre (albeit raised from 0 due to high background counts relative to PL counts), the data in a) has a small central peak. There is still a dip down to near-zero correlation counts at the very centre (more visible in Fig. 4.16c), but a small lobe appears on either side of this dip. This double-peak feature arises due to re-excitation of the quantum dot [10, 14]. When the laser pulses are too wide,³⁹ the first photon is emitted (individually) near the start of the pulse, but some brief period later, the laser power is still strong enough that a second exciton may be generated within the dot. The photon produced by this second exciton is emitted before the next laser pulse arrives, thereby causing the double-peak feature. The dip to near-zero counts at the centre, however, does prove that we are still achieving single-photon emission.

The data presented in these plots is convoluted with the response function of the APDs, but the equation describing this scenario is far more complicated than in the CW case. We therefore use a computer model to extract the probability of single-photon emission. This Monte Carlo model simulates the creation of excitons and biexcitons by consecutive laser pulses, then records the expected correlation counts based on the given radiative lifetimes and detection probabilities [26].⁴⁰ Figure 4.16c) shows the single-dot data (from

³⁷The mode-locking problem was due to misalignment of the internal optics of the laser and was not resolved until after we had completed our nanowire HBT tests.

³⁸Information about pulsed HBT plots is provided in Section 1.2.4.

³⁹Improper mode-locking of the Ti:sapph during our experiments caused its pulses to be far wider than the ideal 1.2 ps.

⁴⁰Further details of the simulation program can be found in [26] and its supporting material.

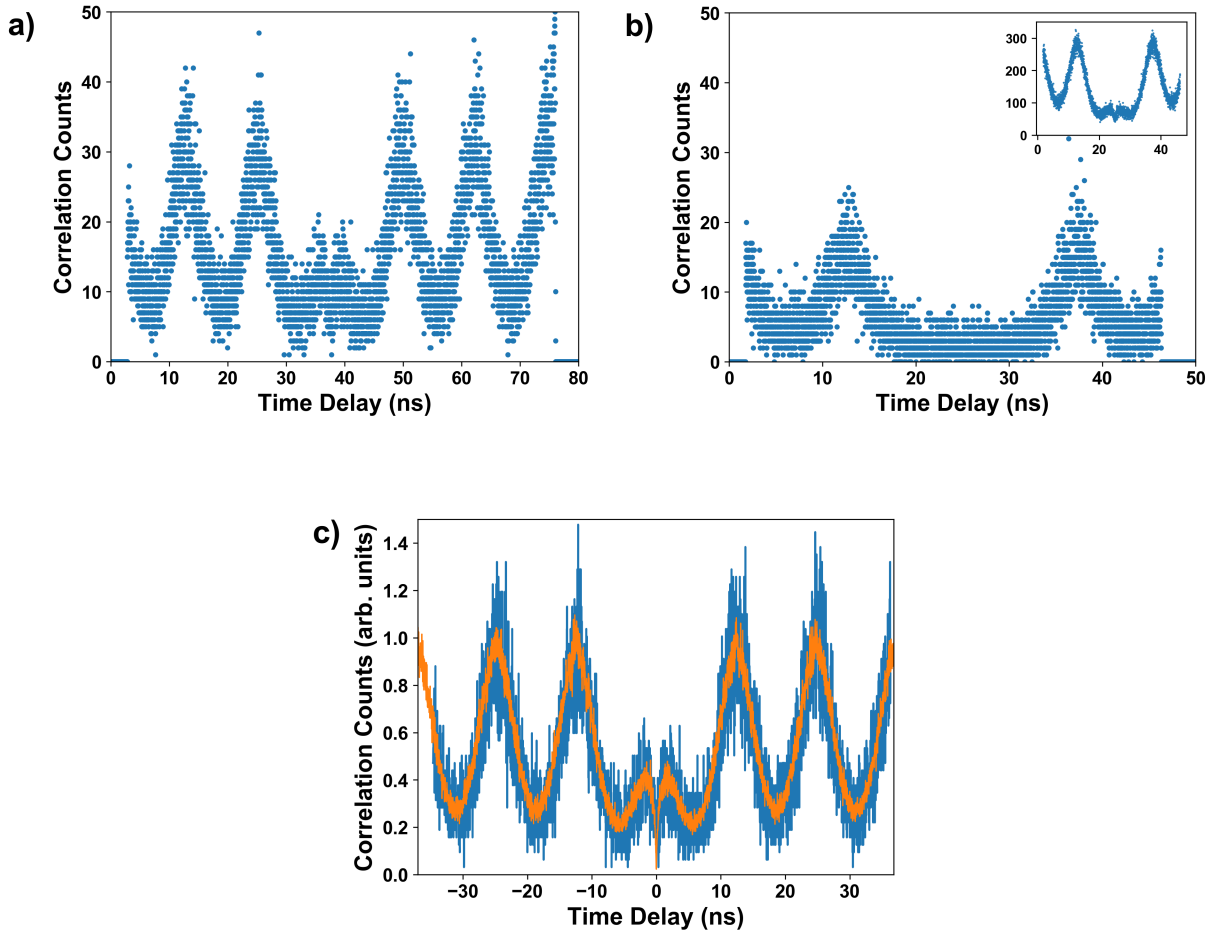


Figure 4.16: a) HBT data for a single-dot nanowire using a pulsed laser and pumping at saturation (15.5-hour exposure period). Note the re-excitation feature at the centre. b) HBT data for a double-dot nanowire using a pulsed laser and pumping well below saturation (12-hour exposure period). No observable re-excitation of the dot. Inset: Same setup with laser pumping at saturation. c) Experimental HBT data (blue) from a), with Monte Carlo simulation (orange) overlaid. Simulation reveals $g_{\text{Exp}}^{(2)}(\Delta\tau_p) = 0.49$.

Fig. 4.16a) with the simulated correlation counts overlaid in orange. The simulated results were obtained by manually adjusting the model parameters, such as excitation probability and number of uncorrelated sources. The program does not have a fitting procedure incorporated, so the simulated results were matched to the data by eye.

The multi-photon emission probability for a pulsed measurement is determined by a slightly different method than in the CW case; rather than using the value at $g^{(2)}(0)$, we use $g^{(2)}(\Delta\tau_p)$, which is the ratio of the area under the central peak to the (average) area under one of the other peaks [10]. From the simulation, we find that $g_{\text{Exp}}^{(2)}(\Delta\tau_p) \approx 0.49$, which is quite high. Pulsed HBT tests conducted by other groups on these same nanowire devices

have shown $g^{(2)}(\Delta\tau_p)$ values as low as 0.008 [69] and 0.004 [10] without dot re-excitation. We expect that most of this discrepancy is due to our use of a high pump power and the overly wide pulses, which can both lead to a greater probability of dot re-excitation.⁴¹

There are not sufficient counts in Fig. 4.16b) to realistically model the data, and its purpose here is merely to illustrate (approximately) what the data would look like without re-excitation of the dot. In an effort to reduce the probability of producing multiple photons per laser pulse, the pump intensity was lowered significantly for this data set (until only 3×10^3 photons/sec were reaching each APD).⁴² The laser was still not mode-locking properly, and as shown by the figure inset, the volcano feature also appeared for this wire when pumped at higher powers. The $g_{\text{Exp}}^{(2)}(\Delta\tau_p)$ for the data in the inset is 0.39, which is only slightly better than in the single-dot case of Fig. 4.16a).

Although the number of correlation counts in Fig. 4.16b) was too low for a reliable $g^{(2)}(\Delta\tau_p)$ calculation, we can clearly see that the central peak is entirely absent, indicating a high degree of single-photon emission. If we also consider the fact that past free-space measurements of these samples have demonstrated one photon per pulse, this suggests that on-demand, fibre-coupled single photons should be possible with our setup, provided the laser pulses are sufficiently narrow.

4.3.4 Exciton Lifetime Measurement

The final test conducted using the APDs was a measurement of the charged exciton lifetime. While a fit to the CW HBT data revealed a lifetime of $T_1 = 2.14 \pm 0.05$ ns, we decided to perform a more direct measurement as well.

As with the HBT tests, the Ti:sapphire laser, in pulsed mode, was used to pump the nanowire, and all but a single peak (X^-) was filtered out of the emission spectrum. For this experiment, the 50/50 beam splitter was removed and the dot emissions were directed to just one of the fibre-coupled APDs, specifically D2. D1 was triggered using the photodiode within the Ti:Sapph, as it provides an electrical signal which is synchronized with the laser pulses. With D1 triggered by a laser pulse, the computer then recorded the delay until a photon emitted by the quantum dot was detected at D2. After many pulses, a statistical lifetime measurement of the exciton can be built up, as shown in Figure 4.17a).

The probability of exciton recombination rises very quickly at the onset of the pulse, then decays slowly after the pulse stops.⁴³ When pumping very hard (as we were doing in

⁴¹Some background correlation counts may also be affecting the results. In a pulsed measurement, background illumination has the effect of lifting the minimum values up from zero, but due to the degree of re-excitation and the overly-wide pump pulses, we cannot accurately subtract background counts from this measurement.

⁴²In addition to lowering the pump power, the probability of re-excitation can also be reduced through resonant excitation [14]. Pumping at the emission energy of the dot eliminates the additional free carriers in the nanowire, thereby preventing them from relaxing into the dot and recombining.

⁴³See Appendix E for an explanation of why we observe symmetric peaks in the pulsed HBT measurement, even though the probability of exciton decay is asymmetric in time.

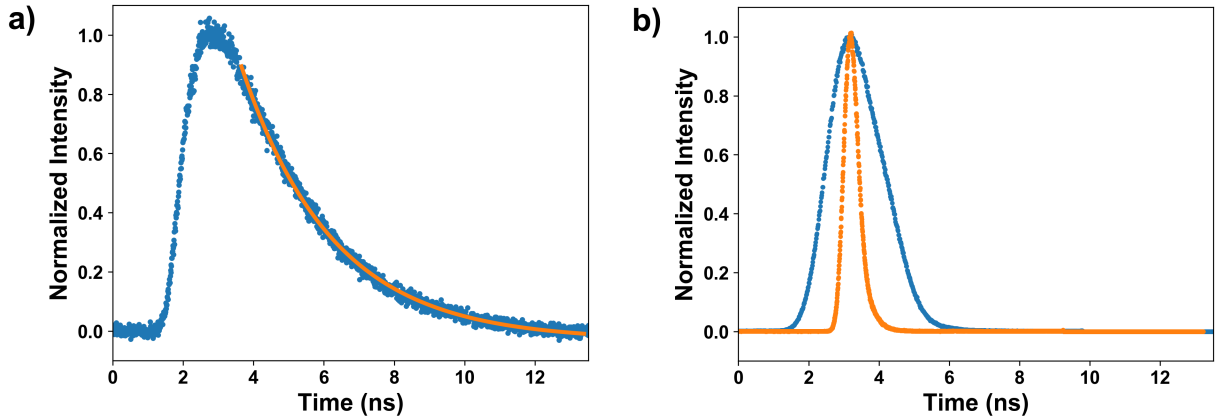


Figure 4.17: a) Lifetime measurement of charged exciton in a nanowire quantum dot. The fit to an exponential decay yields a lifetime of $\tau_d = 2.60$ ns. b) Lifetime measurement of pulsed laser with poor mode-locking (blue) and proper mode-locking (orange).

these experiments), many of the quantum dot's higher energy levels are filled by trapped electrons/holes [4]. As everything except the X^- emission line was filtered out, recombination is only observed once all the extra excitons in the dot decay, which happens fairly slowly. This is why the peak in the lifetime plot does not appear sharply at the onset of the pulse, but is instead smooth and slightly offset from the pulse start [75]. It is followed by the exponential decay of the exciton. Although we do not observe it in this data, the tail-end of the lifetime is occasionally dominated by a second exponential decay characterised by a different lifetime (usually longer). This second decay corresponds to the spin-flip lifetime, when a spin-down (spin-up) electron and a spin-down (spin-up) hole are both in the same dot, and they must wait for the electron to flip its spin before recombining.

From the data in Fig. 4.17a), we can extract the average lifetime of the charged exciton. The decay of the exciton is exponential, given by the simple equation:

$$f(t) = A \exp\left(-\frac{t}{T_1}\right) \quad (4.9)$$

where T_1 is the exciton lifetime. As with the HBT experiment, the data in this plot is a convolution of Equation 4.9 and the Gaussian timing response of the APD, which has a FWHM of 200 ps. The timing jitter is much smaller than the decay period, however, so it is generally disregarded when calculating T_1 [75]. From a fit of the data, which is shown in Fig. 4.17a), the lifetime of the charged exciton in this quantum dot is 2.60 ns. This is very close to the decay time of 2.14 ns which we obtained by fitting the CW HBT data in Section 4.3.2 (the same dot and emission line were used to obtain that data as well). The discrepancy between our two values for the X^- lifetime could be partially due to the overly-wide pump pulses, as the laser was not mode-locking properly while we

were acquiring the lifetime data. Re-excitation of the dot could have increased the apparent decay time through secondary photon emissions. In any case, other groups have tested these same nanowire samples with free-space optics and observed similar X^- lifetimes of 1.6 ns [69] and 2.2 ns [24]. Most photonic nanowires with quantum light sources, including those made using different fabrication techniques and semiconductor materials, have radiative lifetimes ranging from 0.3 ns to 3 ns [10].

The Ti:sapph pulses are Lorentzians and should have a width of approximately 1.2 ps. While we did not directly measure the difference in pulse widths between the laser’s mode-locked and non-mode-locked states, we did make ‘lifetime measurements’ of the laser pulses. The setup was nearly identical to that of the exciton lifetime measurement: the pulsed laser beam (after heavy filtering) was sent to D2 directly, and D1 was triggered electrically by the signal from the laser photodiode. The result is shown in Fig. 4.17b), where there is a very clear difference between the mode-locked (orange trace) and non-mode-locked (blue trace) states. The laser produces much wider peaks when it is not mode locking properly, which explains the appearance of the re-excitation feature in the pulsed HBT measurements. It should be noted, however, that this data is provided only to demonstrate the extreme difference between the two states of the laser; we cannot extract the actual pulse widths by this method. In addition to the timing jitter of the detectors, the mechanics of the laser itself result in an uncertainty in the timing of the pulses, so the time delay between the pulse which triggers D1 and the pulse which is measured by D2 varies.⁴⁴ Therefore, the mode-locked pulse is observed to be much wider in the plot than the true width of 1.2 ps.⁴⁵ The ‘lifetime’ of a laser pulse is therefore not an overly useful or common measurement, but in this case, it does serve to highlight the difference between the two states of the laser and explain the re-excitation of our dot. With the laser properly mode-locked, however, we should be able to achieve on-demand, fibre-coupled single photons.

4.4 Conclusions and Future Work

A fibre-based source of individual photons is a highly desirable technology in the field of quantum communication, and variants on our system of optically coupling individual photonic nanowires to a lensed fibre may have many potential applications in the near future. Our system has sufficient source efficiency (0.48%) to permit analysis of the PL spectra of both single-dot and double-dot samples. By investigating the relationships relating the pump power and temperature to the ZPL width and phonon sideband amplitudes, we have demonstrated the ability to achieve a high degree of photon indistinguishability. HBT tests have proven that we can use our setup as a fibre-coupled source of single photons, with a measured $g^{(2)}(0)$ of 0.095 ± 0.027 . There are also indications that single-photon emission could be triggered on demand with the use of narrow laser pulses.

⁴⁴The electrical signal transmitted to D1 is too slow to catch the pulse which produced it.

⁴⁵It is likely that this relative variation in emission time also contributes to the longer lifetime observed in the pulsed measurements, compared to CW.

Many of these tests would not have been possible, but for the stability of our cryogenic system. Due to the low collection efficiency of our lensed fibre (1–2%), the HBT auto-correlation experiments were conducted over many hours, and high-quality images of the phonon sidebands were often obtained through overnight acquisition of many spectral frames.⁴⁶ Measurements such as these could not be conducted with a helium flow cryostat, as the temperature fluctuations in that type of system tend to occur over a period of minutes. These fluctuations are generally large enough to alter the position of the sample and significantly reduce the coupling efficiency to the nanowire. Our system, however, has demonstrated the ability to remain stably focused on a single nanowire for over three weeks without the need for realignment. This has enabled us to perform long-term measurements without interruption, a requirement made necessary by our poor collection efficiency.

Other groups still use helium flow cryostats for their nanowire experiments, however, as they can achieve a much better collection efficiency using free-space PL collection. A useful single-photon source should certainly have a high source efficiency, ideally nearing 100%. To improve upon the lensed fibre’s collection efficiency, preliminary tests of a fibre-coupled GRIN lens are being conducted. By matching the NA of the fibre to that of the nanowires, the GRIN lens should have a much greater efficiency, making this system much more practical for single-photon collection and delivery.

Further measurements which could help demonstrate the potential of this system include Mach-Zehnder interferometry to measure the photon coherence length, Fabry-Pérot interferometry to measure the ZPL linewidth, and HOM interferometry to quantitatively determine the degree of photon indistinguishability. If some method could be found to make it work with our system, resonant pumping would improve indistinguishability and lower the probability of multi-photon emission by eliminating uncorrelated sources. With the Ti:sapph mode-locking issue resolved, a new pulsed HBT measurement could be used to show that a single photon is produced for each laser pulse, with no re-excitation. This would prove that our system can be used for on-demand triggering of fibre-coupled single photons.

While there is still more work to be done, the basic principles of our system may have many applications in the fields of quantum communication and quantum computing. A fibre-based source of single photons could be very useful for quantum key distribution, where individual flying qubits are transmitted over long distances. Although we did not explore entanglement in our experiments, it has been demonstrated that nanowire quantum dots can also act as sources of entangled photon pairs via biexciton-exciton cascade [14]. With a reliable, fibre-coupled source of entangled photons, the prospect of realising long-distance entanglement distribution and quantum teleportation becomes far more tenable, and we move one step closer to our goal of creating photon-to-spin hybrid quantum repeaters.

⁴⁶Fig. 4.11a), for example, is the averaged result of 1800 frames, each obtained with 22 seconds of CCD exposure.

Chapter 5

Full Polarisation Control of Fibre-Delivered Light in a Dilution Refrigerator

The ability of optical fibres to transmit light over long distances with minimal signal loss (~ 0.2 dB/km) has made them a common telecommunications technology and a prime candidate for use as a means of photon delivery in quantum communication applications. It is a well-known drawback of standard optical fibres, however, that the birefringence of their cores prevents them from preserving the polarisation of the light they transmit. Compensation is therefore required to accurately deliver the desired photon polarisation states through a fibre [42, 44–46].

Adjustment of the polarisation can be performed using polarisation controllers [42, 44–46, 49], which twist, bend, squeeze or stretch the fibre, thereby artificially altering its birefringence [51]. Electronic versions of these devices employ feedback and optimization systems to manipulate the fibre and produce the appropriate polarisation at the fibre output. Compensation for a fibre’s effects on polarisation is often achieved by transmitting an ancillary signal with a known polarisation. The second signal provides a measure of the change in fibre birefringence and acts as a reference when attempting to correct the output polarisation. Some of the schemes which are used to perform these corrections include temporal multiplexing [40, 45, 46], wavelength multiplexing [42, 44, 45] and two-fibre systems, where the reference pulses are transmitted along a secondary channel neighbouring the primary [45, 49]. While the overall efficiency and success of these schemes depends upon such factors as the fibre’s length and external environment, they are generally ill-suited for the extreme conditions of cryogenic experiments.

Much of the literature on transmitting polarised light via optical fibre deals exclusively with linear polarisation states at the fibre output [42, 44–46]. These systems, which often use polarisers to ensure linear output, are not concerned with the phase between the x and y components of the electric field. In this chapter, we will explore a phase-preserving

system for delivering any polarisation state, including elliptical and circular states, to a dilution refrigerator via a standard, 30-metre, single mode optical fibre. A compact optical setup installed in the refrigerator allows us to accurately identify the polarisation of light reaching the sample while cryogenic temperatures are maintained. A mathematical model of the system is generated by sending known polarisations into the fibre and observing the corresponding output states. Waveplates (half-wave and quarter-wave) are used to compensate for the polarisation-altering effects of the fibre so that we may deliver high-fidelity, arbitrary polarisation states from room temperature to the bottom of the refrigerator while operating at milli-Kelvin temperatures. The cryogen-free dilution refrigerator used in these experiments maintains a stable temperature gradient along the fibre, which resolves the issue of temperature fluctuations experienced by wet dilution refrigerators due to changing helium bath levels. Future versions of this experiment will employ the same setup to deliver single polarised photons to a sample in the dilution refrigerator as part of the photon-to-spin project.

I would like to thank Dr. Louis Gaudreau for the initial conception and development of the system presented in this chapter, and for his assistance and advice as we worked on realising it. Thank you also to Dr. Marek Korkusinski, who developed the genetic fitting algorithm, Piotr Zawadzki for assisting with the design and construction of our cryogenic optical systems, and Dr. Alex Bogan for devising the proof that our setup can uniquely identify any polarisation (Appendix F). Many thanks to Dr. Sergei Studenikin, Dr. Robin Williams, and Dr. Andrew Sachrajda for helpful discussions about this project. The contents of this chapter are to be published in a peer-reviewed journal in the near future, largely without further modification. While I have personally written this chapter and will appear as first author on the associated paper, corrections and suggestions were made by all those mentioned previously, who will also appear as co-authors on the paper.

5.1 Details on the Polarisation Control Setup

The setup used to compensate for the polarisation-altering effects of a fibre is depicted schematically in Figure 5.1 and described in detail in Section 2.4. In our optics lab, a beam from a Ti:Sapph laser with a 780 nm wavelength is directed through a linear polariser (labelled P1), a half-wave plate (HWP), and a quarter-wave plate (QWP), the latter two of which are mounted in computer-controlled rotators. We label these waveplates HWP and QWP1 in our setup. A collimator is used to couple the beam into a fibre. Following the collimator, a 50/50 fibre beam splitter and photodiode (labelled PD1) are used to monitor the intensity of the light making it into the fibre. The rest of the beam is transmitted through a 30-metre, single-mode optical fibre, which takes the light from the optics lab to the cryogenics lab and down into a dilution refrigerator. At the mixing chamber (MC) of the fridge, which has a base temperature of ~ 10 mK, nanostructure samples are mounted on the cryogenic SPIN setup. Three piezoelectric positioners can direct a lensed fibre to either focus the light on the sample or to direct it at the three optical components used

for polarisation identification: a QWP (labelled QWP2) in a piezoelectric rotator, a linear polariser (P2), and a photodiode (PD2).




Figure 5.1: A schematic of the full setup used to deliver polarised light to the dilution refrigerator by fibre, including the three polarisation identification components in the fridge. The curved arrows on the waveplates represent the fact that they are mounted in rotators, while the horizontal arrows on the linear polarisers indicate their polarisation direction (i.e. horizontal).

As mirrors induce a slight change in polarisation, there are no mirrors placed between P1, HWP and QWP1. Mirrors are used after the waveplates, however, to direct the beam into the collimator, which couples the free-space light into the fibre. We account for these mirrors and the collimator in our simulation by considering them to be part of the fibre; our model simulates the fibre (and other components, like the fibre lens) as a long series of rotated retarding plates.

Intensity fluctuations within the fibre (Fig. 5.2a) may be caused by noise in the power output of the laser or by a wobbling of the beam path as the waveplates rotate. This noise reduces the accuracy of our characterisation of the fibre. To minimize beam wobble, the rotators are aligned with tilt stages to adjust pitch and yaw. With a mirror in the place of the waveplate, the tilt of the rotator is adjusted until the reflected beam is centered along the beam path during a full rotation of the mirror.¹ The tilt of the mirror itself is controlled using a kinematic mirror mount attached to the rotator, which is adjusted until the reflected beam completely overlaps with the incident beam during rotation. While this procedure significantly reduces beam wobble, we are still left with the issue of laser power fluctuations. Preliminary tests of the system employed a fibre-coupled, 980 nm, benchtop laser source (Thorlabs, S1FC980), which has a much more stable power output than our Ti:Sapph. The wavelengths required for the photon-to-spin experiments are much closer to that of the Ti:Sapph (~ 800 nm), however, so we needed to develop some method of compensating for the noise.

To resolve the noise issue, a fibre beam splitter and additional photodiode (PD1) were

¹The mirror is typically not perfectly vertically-aligned within the rotator, so the reflected beam traces out a circle during one full rotation. The tilt stage is used to align the centre of that circle with the beam path.

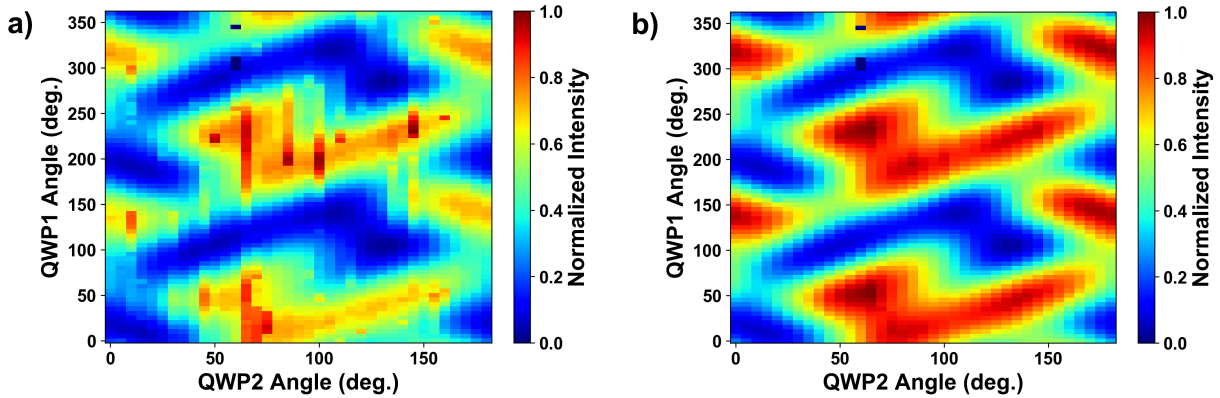


Figure 5.2: a) Raw data showing noise due to laser intensity fluctuations. b) The same data set after post-processing to eliminate intensity noise.

installed in front of the 30m fibre, as shown in Figure 5.1. This setup allows for monitoring of the light intensity entering the fibre without affecting the polarisation of the light reaching the sample space. For each intensity measurement we make in the fridge, a PD1 measurement is also made of the light intensity entering the fibre system. We calculate the ratio of each PD1 data point to the average intensity, and we then divide the corresponding polarisation datum by that ratio. The result is a rescaling of the polarisation data with respect to the average and a correction of all noise originating from the room-temperature setup in front of the fibre. Figure 5.2 compares the raw data with a corrected version. We note that before post-processing, the average signal-to-noise ratio can be as low as 6, while it is about 75 afterwards. Hence, we conclude that almost all noise in the system appears to originate from the optics lab components in front of the fibre. The fibre itself and the refrigerator setup appear to introduce very little noise into the polarisation, in spite of the small vibrations caused by the pulse-tube refrigerator [87].

The fibre used in our setup is a standard, single mode fibre with no built-in polarisation-maintaining stress rods. While there are polarisation-maintaining (PM) fibres available, they only preserve the two perpendicular linear polarisations which are aligned with the core's fast and slow axes, i.e. the horizontal and vertical states. All other polarisations, however, will be projected onto these two modes, each of which travels at a different speed. The result is that one cannot use PM fibres to send diagonal linear, circular or elliptical polarisations [41, 47]. (Spun optical fibres preserve circular polarisation, but there are no fibres which will preserve all polarisation states simultaneously.)

The optical anisotropy in the fibre core causes one component of the electric field to travel faster than the other, resulting in a phase shift in the polarisation. The difference between the arrival times of light travelling in these two modes, called the differential group delay (DGD), is typically on the order of a picosecond or under for a 100km-fibre, which is not measurable by modern detectors [45]. Current detector technology allows for

temporal resolution on the order of a few hundred picoseconds, which is much larger than most DGD values [44]. Therefore, if the system described in this chapter were applied to transmit single polarised photons, there would be no timing issues associated with the two components of the electric field travelling at different speeds.

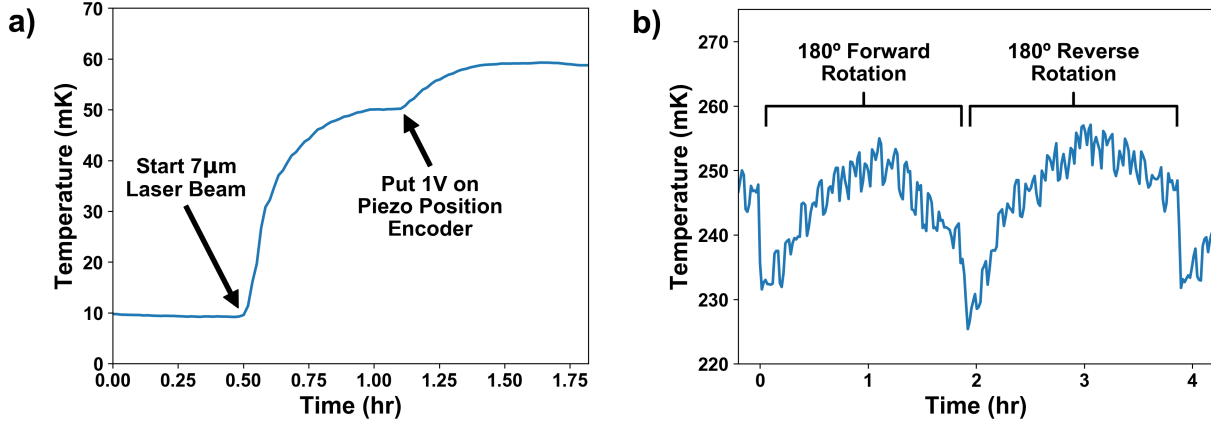


Figure 5.3: Fridge heating during polarisation tests. a) Temperature increasing from 10 mK due to laser beam and voltage applied to piezo position encoder. b) Heat dissipation in MC due to rotation of QWP2.

The intensity of light exiting the lensed fibre in the refrigerator was approximately 7 μ W in our experiments, as measured by PD2 at low temperatures. When operating the refrigerator at low temperatures, this laser power had the effect of increasing the mixing chamber temperature from base temperature up to 50 mK. Additionally, the MC temperature increased from 50 mK to 60 mK by the application of 1V to the piezoelectric rotator’s resistive encoder, which was necessary for readout of the rotator’s position. The temperature changes corresponding to these heat inputs are shown in Figure 5.3a). Further heating was caused by the rotator’s movement during experiments, requiring us to navigate the trade-off between heat dissipation and rotator speed/step size. Figure 5.3b) demonstrates the heating caused by two polarisation identification tests, which raised the temperature to 250 mK due to heat dissipation by the rotator.

5.2 Preliminary Characterisation of Optical Components

Accurate simulations of our system and proper characterisation of the fibre requires that each optical component be examined beforehand to verify that it is performing as expected. While they may theoretically act in specific ways, no optical element is perfect, and the specifics of our setup can affect their behaviour. This is especially true of those components

whose effects are dependent upon the wavelength or angle of incidence (AOI) of the light. Preliminary tests of the waveplates, polarisers, and a gradient-index (GRIN) lens were conducted to verify proper operation and determine the parameters for our simulations.

5.2.1 Determining Waveplate Retardances

In the first section of our room-temperature setup, a free-space laser beam travels through a polariser, a HWP, and a QWP. The waveplate retardances are dependent upon both the wavelength and the AOI of the laser beam, and they therefore differ slightly from the ideal values of 180° (HWP) and 90° (QWP). For accurate modelling and fitting of the experimental results, we must reduce the number of unknown parameters by determining the true waveplate behaviour in advance.

The first step in characterising the waveplates is to align two linear polarisers, namely P1 and a second polariser. This second polariser can be either another room-temperature polariser or the fridge polariser, i.e. P2.² The light passing through P1 is always defined to be the ‘horizontal’ polarisation state to simplify our calculations. With a laser beam passing through the two polarisers, the second polariser is rotated until the transmitted intensity is maximized, as measured by a photodiode. When the intensity is greatest, the polarisers are aligned and most of the light is transmitted.




Figure 5.4: a) With a HWP placed between P1 and P2, the measured intensity (blue dots) of transmitted light as a function of HWP angle is plotted. Fitting the data to the model (black line) reveals a 36.8° misalignment between the polarisers. b) Transmitted intensity as a function of HWP angle for a HWP inserted between two polarisers. Dashed line indicates the measured intensity with the waveplate absent, solid line is a fit to the model.

²Testing P2 requires the cryogenic SPIN setup to be disconnected from the fridge and specially mounted on a breadboard.

In the case of aligning P1 and the fridge polariser P2, the latter of which is fixed in its orientation, a (pre-characterised) HWP can be inserted between the polarisers. The HWP is used to rotate the polarisation of the light exiting P1, and a photodiode measures how much of the light is then able to pass through P2. Figure 5.4a) shows the measured intensity of light reaching the photodiode as a function of the HWP orientation. (The angles on the x-axis are those inscribed on the HWP’s manual rotator.) A fit of the intensity data is produced using the equations of Section 1.5 and the Jones matrices of Table 1.2. The result of the fitting procedure indicates that P2 is offset by 36.8° from P1, so a rotation of P1 by 36.8° aligns the polarisers.

To characterise the behaviour of a HWP, the waveplate is placed between two aligned polarisers, and the transmitted intensity is measured as a function of HWP angle. As shown in Figure 5.4b), the rotation of the HWP causes a rotation of the polarisation, which oscillates between parallel to the second polariser and perpendicular to it. The data is normalized using the intensity measured with the HWP removed. The fast axis of the HWP is not necessarily aligned with the ‘ 0° ’ setting of the computer-controlled rotator within which it has been mounted, so the values along the plot’s x-axis are those of the rotator. In order to accurately simulate the system, the offset angle between the waveplate’s fast axis and the rotator’s ‘ 0° ’ position needs to be determined. If we define the orientation of the first polariser to be horizontal, we can use Jones calculus to obtain a fit of the data.³ From the fit, we find that the offset angle is 44.3° , and the retardance is 185.4° . This imperfection in the retardance accounts for the fact that the normalised intensity data never quite reach zero or one, which is what would normally be expected of perfectly linear polarisation.

As with the HWP, QWP1 is also tested by placing it between the two polarisers. As the effect of a QWP depends upon the orientation of the incident linear polarisation, the HWP is placed in front of QWP1. The transmitted intensity is then measured as a function of both the HWP and QWP1 angles, as shown in Figure 5.5. A fit to the model, employing the known HWP retardance and offset, indicates that QWP1 has a 21.8° offset and a 90.72° retardance.

Testing QWP2 is more difficult, as it is integrated into the cryogenic SPIN setup in the fridge and its removal would alter its AOI-dependent retardance. The entire SPIN setup is therefore removed and mounted on a breadboard. The symmetry of the optical components means that they are unaffected by a 180° rotation, so the SPIN setup is flipped upside-down and fastened to the table by its copper ring. (See Section 2.1.2 for a detailed description of the SPIN setup.) As described above, P1 and a HWP are also mounted on the breadboard and the two polarisers are aligned (during which QWP2 is moved out of the way). The transmitted intensity is then measured as a function of QWP2 angle for several HWP orientations. While this is a valid technique for determining the waveplate

³The rotator’s ‘ 0° ’ setting is horizontal, so in this scenario, the rotation matrices are written as $R(\theta+\chi)$. The angle θ is the theoretical rotation from horizontal, and the constant χ accounts for the actual offset between ‘ 0° ’ on the rotator and the orientation of P1.

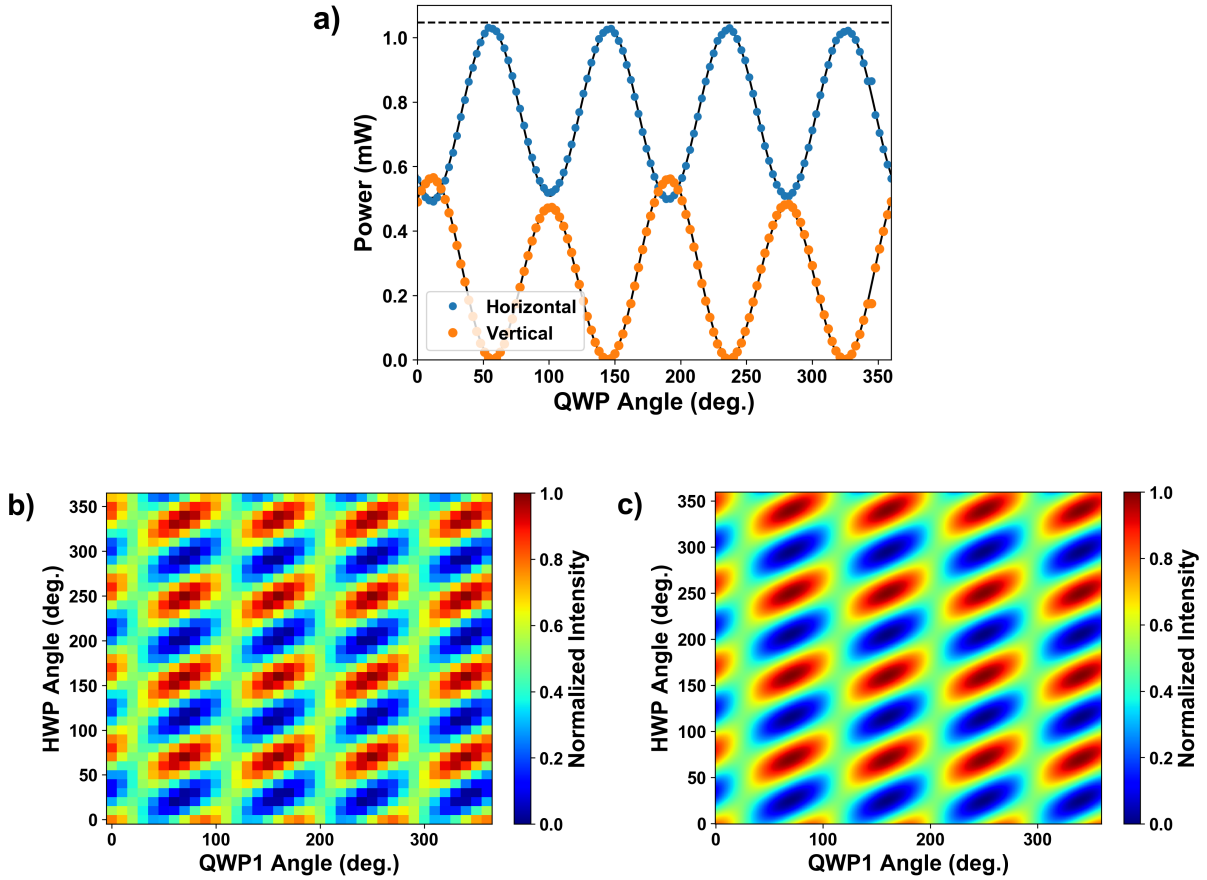


Figure 5.5: Characterising a QWP. a) Transmitted intensity as a function of QWP angle for horizontal (blue) and vertical (orange) input polarisations. Dashed line is the background level with the waveplates absent. b) Low-resolution plot of the measured intensity as a function of HWP and QWP orientations. c) Simulation of the data using the actual waveplate retardances and offsets. Note how well it matches with the data.

parameters, the initial values found using this method (68.0 offset and 92.5 retardance) did not produce overly accurate simulations of the data. As will be discussed below, they were instead fit using a genetic algorithm which was simultaneously attempting to retrieve the parameters describing the fibre. The resulting values of a 49.73° QWP2 offset and a 95.43° retardance allowed for more accurate modelling of the system.

Despite careful attention to detail, some amount of human error will lead to imperfect characterisation of these components. This is especially the case for the room-temperature polarisers, which determine the transmitted intensities used to calculate the waveplate retardances. These polarisers are mounted in manual rotators, which do not allow for a precision of much more than 0.5° . With P1 and P2 in series, the rotating HWP between

them was used to show that there was a residual 0.11° offset even after alignment was attempted. This offset cannot be removed because of the difficulty in accurately rotating a ~ 2 -inch-wide rotator 0.11° by hand. This implies that some amount of the infidelity in our fibre-delivered states is due to error in our characterisation of the waveplate retardances/offsets, although this is expected to be only a very slight contribution.

As for the waveplates themselves, the minor deviations from the ideal retardances⁴ are taken into account by our simulations and do not have any significant effect on our ability to deliver the appropriate states through the fibre or identify the states arriving in the fridge. Due to their wavelength- and AOI-dependence, the characterisation of the waveplates must be redone each time a different wavelength is used or the setup is altered. The alignment of the lensed fibre in the fridge does not appear to have any major effects on the retardance of QWP2.

5.2.2 The Depolarising Effects of a Gradient-Index Lens

As the focal distance of the fibre is only $8\ \mu\text{m}$, the beam begins to diverge long before it passes through the cryogenic optical components, i.e. QWP2, P2 and PD2. This is shown schematically in Figure 5.1. In our initial setup, a gradient-index (GRIN) lens was placed within the hole in the circuit board to collimate the light before it reached QWP2. The expectation was that this would result in less scattering of the laser and more accurate results when using low light intensities.

During initial polarisation tests in the fridge, however, we discovered an unexpectedly large degree of depolarisation, one which was not observed when using a fibre of similar length during preliminary tabletop tests. Individual tests of all cryogenic components proved that the GRIN lens was altering the polarisation and partially depolarising the light. While Mueller calculus can account for depolarisation, it was decided that the GRIN lens would be removed altogether. Despite the diverging beam, the absence of the GRIN lens does not appear to have led to scattering of light inside the fridge, nor to a reduction in our ability to monitor the light intensity as a function of QWP2 angle. In the future, no focusing will be employed beyond that which is achieved by the fibre lens. With one less polarisation-altering component, the quality of the experimental results can only improve, as our goal is to show that it is possible to compensate for a *fibre's* effects on polarisation.

5.3 Modelling the Optical Fibre's Effects on Polarisation

To deliver particular polarisation states via fibre and know in advance the waveplate orientations which will compensate for the fibre's effects, a complete characterisation of the

⁴These are assumed to be mostly real deviations and not merely the product of imperfect characterisation.

fibre’s birefringent core is required. This is accomplished mathematically using the Jones and Mueller calculi described in Section 1.5.⁵ The first polariser, which we define to be oriented horizontally, gives an input state of $E_{\text{in}} = (1, 0)^T$ every time, so the only components which we need to account for are the HWP, QWP1, QWP2, P2, and, of course, the fibre F. Based on the rules of Jones calculus, we can represent the entire optical system as:

$$E_{\text{out}} = \text{P2} \cdot \text{QWP2} \cdot \text{F} \cdot \text{QWP1} \cdot \text{HWP} \cdot E_{\text{in}} \quad (5.1)$$

where each variable represents the Jones matrix for the corresponding optical element, the equations for which are summarised in Table 1.2. Rotation matrices account for the orientations of the various components. When determining the parameters of the fibre matrix F , the measured offsets and retardances of the other components are fed into the equation so that F remains the only unknown. As derived in Section 1.5.3, the fibre matrix is given by Equation 1.43, which we restate here:

$$F = R(-\theta_a)D(\phi)R(\theta_b) \quad (5.2)$$

This is in fact a generic equation for any optical element which alters polarisation, including fibres, mirrors, and lenses, and this single matrix may represent several of these components in series. Solving for the three parameters θ_a , θ_b , and ϕ of the fibre matrix is merely a matter of analysing the output state with respect to the input state. Our system is designed to identify the output state of the fibre, while the input state is determined by the waveplates in front of the fibre.

The actual measurements we make with the cold photodiode PD2 are not of the electric field’s x- and y-components, but rather of the intensity of the light transmitted through P2. Therefore, after calculating the state E_{out} , the expected intensity is given by the complex square of E_{out} . These calculated intensities can then be plotted as functions of the waveplate angles. Such simulations help reveal the accuracy with which the fibre has been characterised and can be used to predict the orientations of the HWP and QWP1 which will produce the desired polarisations at the fibre output.

5.4 Delivery of Polarised Light

With the preliminary work, such as identifying waveplate retardances, now complete, we can begin to deliver polarisation states to the fridge. Our standard procedure for delivering polarised light by fibre is comprised of the following four steps:

1. **Measure Fibre Effects:** Send known polarisation states into the fibre and measure the output states.

⁵We will use Jones vectors and matrices in this section to simplify things, but in reality we sometimes use Mueller calculus to account for depolarisation.

2. **Fit Fibre Parameters:** Fit the data to the fibre model using a genetic algorithm and extract the three fibre parameters $(\theta_a, \theta_b, \phi)$.
3. **Calculate Waveplate Orientations:** Calculate the waveplate orientations required to deliver the desired states.
4. **Verify Delivery of States:** Verify that the desired states have indeed been delivered to the fridge to prove that the fibre has been accurately characterised.

In the following sections, we will elaborate upon each of these four steps to provide a detailed description of our process.

5.4.1 Step 1: Measure Fibre Effects

The fibre's effects on polarisation are characterised by sending light through it and comparing the known input states with the resulting output states. The input states are prepared using a set orientation of the HWP while rotating QWP1. The output states are then identified by rotating QWP2 and using PD2 to measure the amount of light which passes through the cryogenic polariser P2. This combination in the fridge of a quarter-wave plate, a linear polariser and a photodiode is sufficient to uniquely identify every polarisation state (see Appendix F for a proof of this claim). This procedure is repeated for several HWP angles, each one resulting in a map of the measured light intensity as a function of the QWP1 and QWP2 angles. Figure 5.6a) depicts two such maps for two different angles of the HWP. Accurate characterisation of the fibre's effects on polarisation depends on the quality of these initial maps, so higher resolutions and additional data sets result in better fibre characterisation.

There is a trade-off between the accuracy of the fitted fibre matrix and the time required to produce these maps. In our system, the speed at which we can produce each data set is limited by the speed of QWP2's piezoelectric rotator, which dissipates heat in the refrigerator as it rotates. At room temperature, we are able to produce 4000-data-point maps in 3 hours. For identical resolution at low temperatures, each data set requires approximately 7 hours, and the MC temperature rises to approximately 140 mK during those runs.

5.4.2 Step 2: Fit Fibre Parameters

To obtain the fibre parameters, the polarisation maps are fit to the model using a custom-made genetic algorithm program, which was written for us by Dr. Marek Korkusinski of the NRC. The algorithm simulates the optical system using Jones calculus and the fibre model discussed in Section 5.3. Starting with several randomly-selected seed values for the unknown parameters, the program employs mutations and recombinations over many

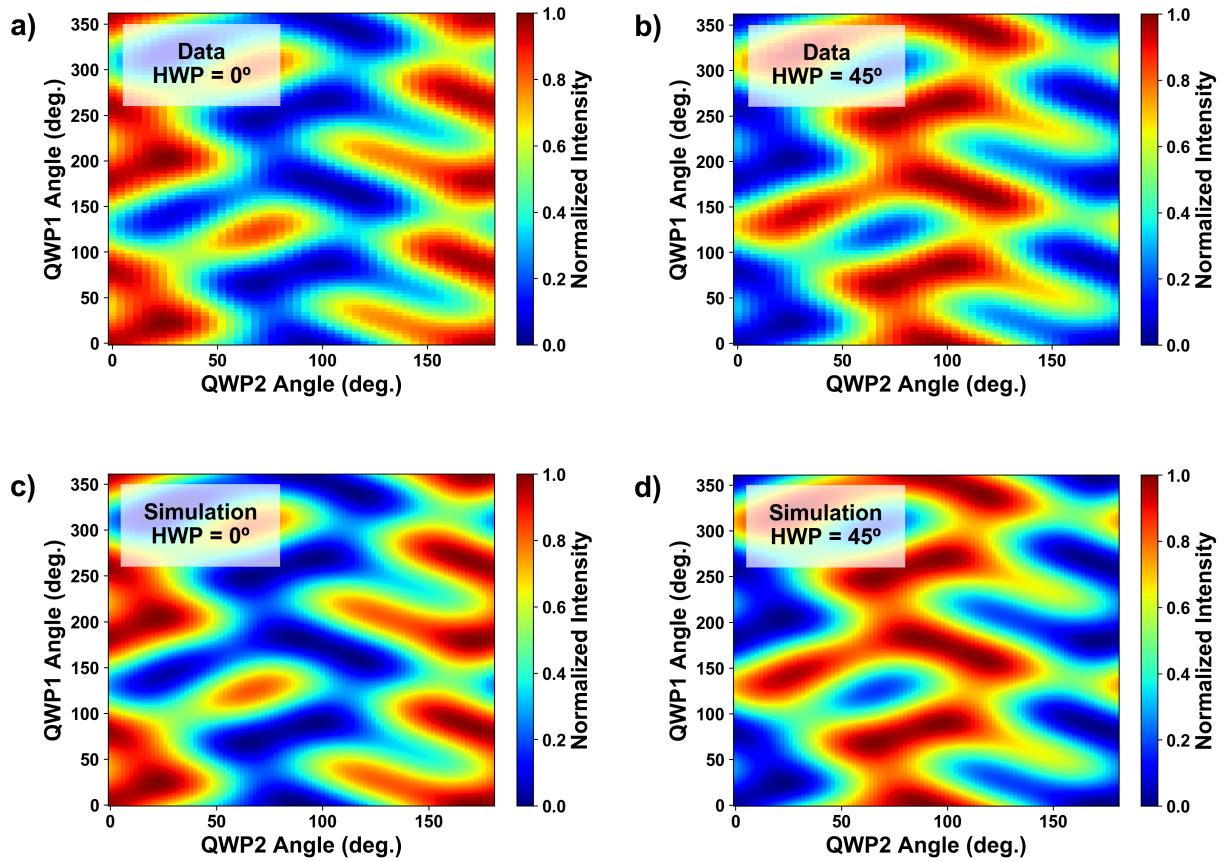


Figure 5.6: Polarisation maps produced for HWP angles of a) 0° and b) 45° . Simulated versions of the data are shown in c) and d). These simulations were made after extracting the fibre parameters using a genetic algorithm.

generations to evolve each solution and find the values which allow the closest fit of the simulation to the data. For greatest accuracy of the results, multiple data sets (each with a different HWP angle) are loaded into the algorithm program simultaneously. Depending on the amount of data being fit, the program usually requires only 200-300 genetic iterations to solve for the fibre parameters with sufficient accuracy.

The genetic algorithm program has ten input parameters which are set by the user: the three fibre parameters, the retardances of the three waveplates, the offset angles between the three waveplates' fast axes and the ' 0° ' setting of their respective rotators, and the offset angle between the first linear polariser and the second. In each execution of the program, the user is able to choose which of these parameters to fix at specific values and which are to be fit by the algorithm. The final results begin to decrease in accuracy if too many input parameters are left floating. The three fibre values are typically the only variables which are fit, while the other parameters are measured directly and set as fixed parameters in the fitting procedure. It should be noted that the program is unable to simultaneously fit

both the fibre parameters and the offset of the second polariser, so the orientation of P2 must be measured directly.

A preliminary test of the fitted fibre parameters is to use them in a simulation and compare the simulated polarisation maps to the actual data. Figure 5.6 shows a comparison of low-temperature data and the corresponding simulations after fitting. The fitted fibre values for those data sets were: $\theta_a = 96.96^\circ$, $\theta_b = 101.5^\circ$, and $\phi = 106.9^\circ$. The retardance of QWP2 was also fit in this case, just to verify that it was not significantly affected by the cryogenic temperatures. The fitting algorithm returned a retardance of 95.4° , which is close enough to the ideal value of 90° that we are still able to use the waveplate for accurate identification of polarisation states in the refrigerator.

The reliability of the polarisation maps and fitted fibre parameters depends on the fibre birefringence remaining constant. Consequently, all polarisation measurements are made only when the refrigerator has reached a stable temperature, and new polarisation maps must be acquired and fit if there is any significant change in the state of the fibre. Minor changes in birefringence, such as those due to small vibrations, do not appear to significantly affect the results for fibres of the length used in our experiments (30 m).

5.4.3 Steps 3 & 4: Calculate Waveplate Orientations and Verify Delivery of States

Each polarisation requires a unique combination of HWP and QWP1 angles in order to recreate the desired state in the fridge. Note that we are not using these waveplates to create the inverse of the fibre matrix, F^{-1} , but merely to take the fibre's birefringence into account in order to transform the input state (horizontal) into the desired state at the fibre output.⁶ The waveplate angles are calculated by inserting the fibre parameters into the model (Section 5.3) and then fitting/solving for the HWP and QWP1 orientations which produce the desired state at the fibre output.⁷

Once the waveplates have been properly oriented, the polarised light is sent through the fibre to the dilution refrigerator. Measuring the light intensity with PD2 as a function of the QWP2 angle produces a signature which is unique to each polarisation (see Appendix F for uniqueness proof), allowing for identification of the states exiting the lensed fibre. Figure 5.7 shows simulations of six such signatures. These six 'standard' polarisations (horizontal, vertical, diagonal, anti-diagonal, right circular and left circular) form the basis in which polarisation experiments are conducted.

The measured signatures of the six standard polarisations are displayed in Figure 5.8. The simulations in these graphs (black lines) were produced using the fitted QWP2 re-

⁶We will prove in Section 5.6 that it is not possible to create F^{-1} with just two waveplates. We are therefore just transforming the input state, not fully compensating for the fibre. Future versions of this experiment will allow for complete compensation by using additional waveplates.

⁷Note that we are calculating the state at the *fibre output*, not the light reaching PD2, as in Equation 5.1.

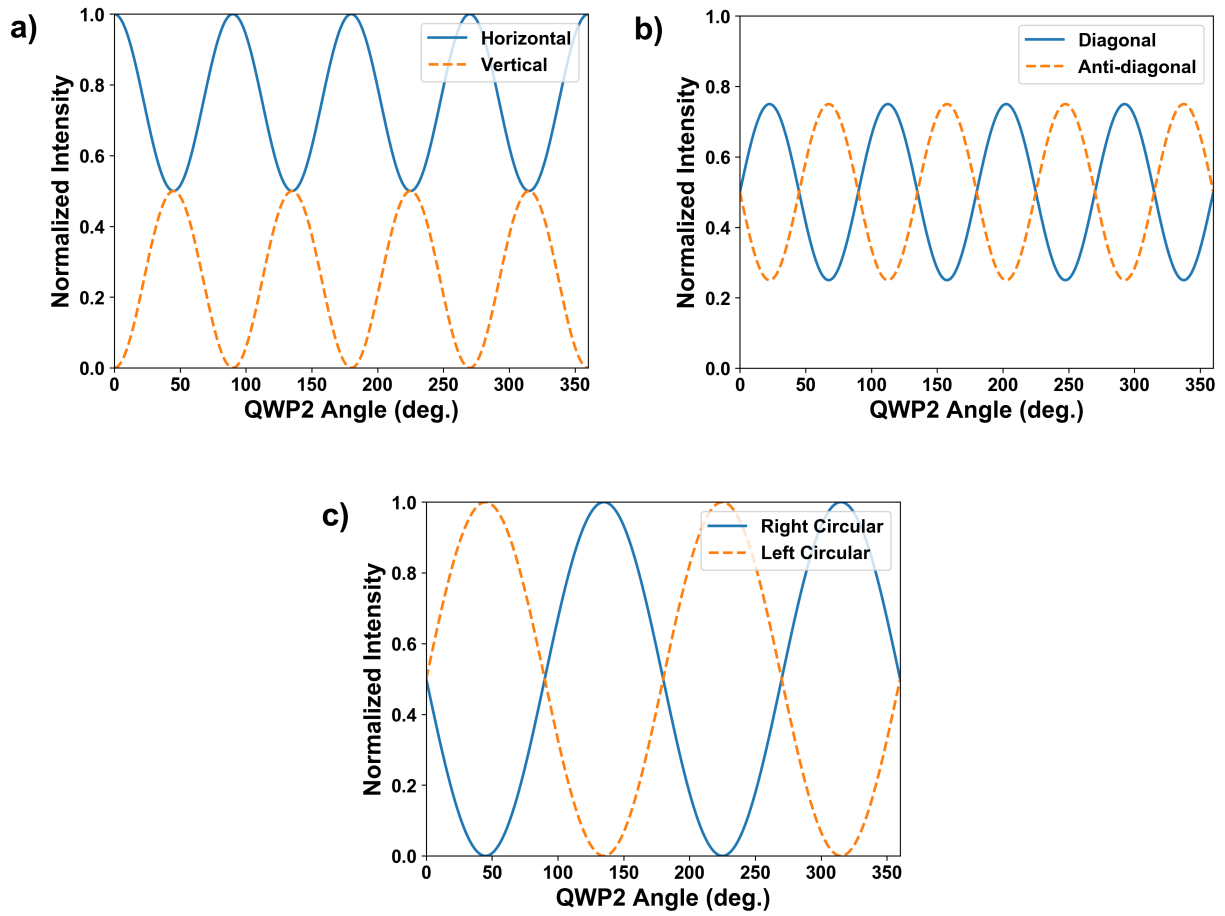


Figure 5.7: Simulations of the signatures generated by the QWP-polariser-photodiode combination for a) horizontal and vertical, b) diagonal and anti-diagonal, and c) right and left circular polarisations. All were simulated using ideal waveplate retardances.

tardance, whereas the simulations shown in Figure 5.7 use the ideal retardance. As each signature repeats after a 180° -rotation of QWP2, data was only taken over that range to reduce heat dissipation in the refrigerator. We used the fibre splitter and PD1 to eliminate noise fluctuations (see Section 5.1), and all six data sets were normalized to the maximum value of the horizontal data set.

For each of these polarisations, we were able to successfully deliver the desired state to the refrigerator at low temperatures using the unique set of waveplate angles calculated for it. Deviations of the data from the simulations in these plots are mostly attributed to slight depolarisation (note how the vertical and circular states don't quite touch the 'zero intensity' level) and issues with the step size/position of the QWP2 rotator. The latter in particular is responsible for the stretched appearance of the data, which is perhaps most noticeable in the circular states (although depolarisation is also contributing there). The

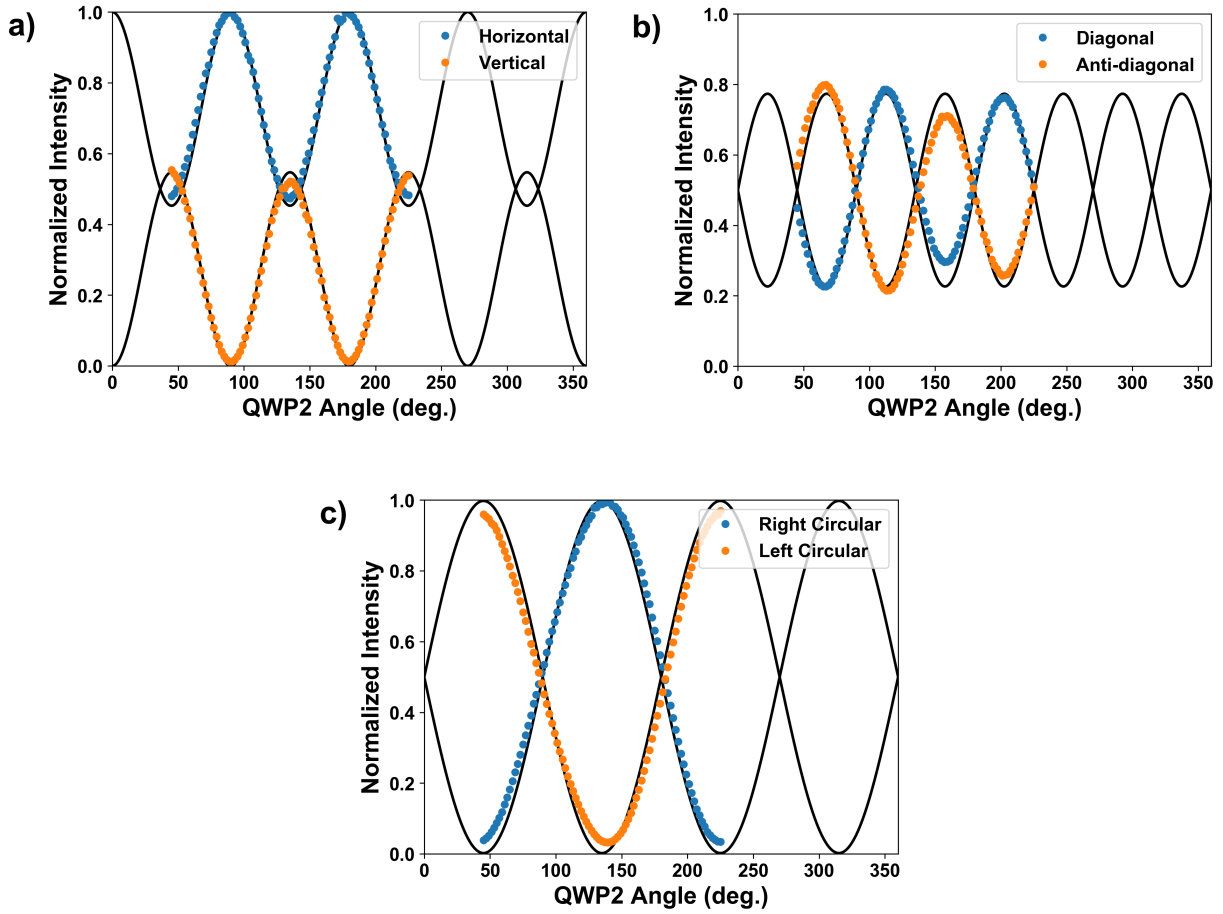


Figure 5.8: Measured signatures generated by the QWP-polariser-photodiode combination for a) horizontal and vertical, b) diagonal and anti-diagonal, and c) right and left circular polarisations. All were measured at low temperatures and overlaid on the simulated signatures (using the real waveplate retardances).

piezoelectric rotator uses an imperfect, home-made calibration file (see Section 3.2.1) which results in some non-linear behaviour, e.g. occasionally rotating by 1.2° when instructed to move by 1° . This can lead to some distortion of the data, but overall it appears that we have been able to deliver the desired states with decent accuracy.

To quantify the quality of the output states, we calculate their fidelities. For pure quantum states, the fidelity of a measured state $|\psi\rangle$ with respect to the expected state $|\phi\rangle$ is defined as:

$$F = |\langle\phi|\psi\rangle|^2 \quad (5.3)$$

Note that although we are dealing with a beam of light, which is generally described

in classical terms, the lack of interaction between the individual photons in the beam allows us to apply this quantum definition of fidelity to their collective polarisation. To account for depolarisation, we apply a generalized definition of fidelity. The fidelity of a given polarisation with density matrix ρ with respect to the ideal state σ is defined as:

$$F = \left(\text{Tr} \sqrt{\sqrt{\sigma} \rho \sqrt{\sigma}} \right)^2 \quad (5.4)$$

As described in Section 1.5.4, one can easily derive the fidelity equations for the six standard polarisation states. These equations are summarized in Table 5.1, as are the calculated fidelities for these states measured at both room temperature and low temperatures. A low-temperature fidelity of greater than 0.96 was achieved for all polarisations, with uncertainties of 0.007 or less.

State	Fidelity Equation	Room-Temp. Fidelity	Low-Temp. Fidelity
Horizontal	$\frac{1+S_1}{2}$	0.993 ± 0.001	0.994 ± 0.004
Vertical	$\frac{1-S_1}{2}$	0.978 ± 0.004	0.978 ± 0.003
Diagonal	$\frac{1+S_2}{2}$	0.986 ± 0.005	0.961 ± 0.007
Anti-Diagonal	$\frac{1-S_2}{2}$	0.996 ± 0.005	0.964 ± 0.006
Right	$\frac{1+S_3}{2}$	0.924 ± 0.003	0.983 ± 0.002
Left	$\frac{1-S_3}{2}$	0.922 ± 0.003	0.968 ± 0.002

Table 5.1: Equations for the fidelities of the six standard polarisation states, as well as the measured fidelities for polarised light transmitted through a fibre at both room temperature and low temperature.

In addition to the six common polarisation states, three elliptical states were also chosen at random to verify that we could deliver arbitrary polarisations to the refrigerator. These states were selected by randomly picking HWP and QWP1 orientations (without a fibre), but here we define them using Stokes vectors, $S^\mu = (S_0, S_1, S_2, S_3)^T$. As we saw in Section 1.5.1, S_0 is the total (often normalized) light intensity, and S_1 , S_2 , and S_3 are the

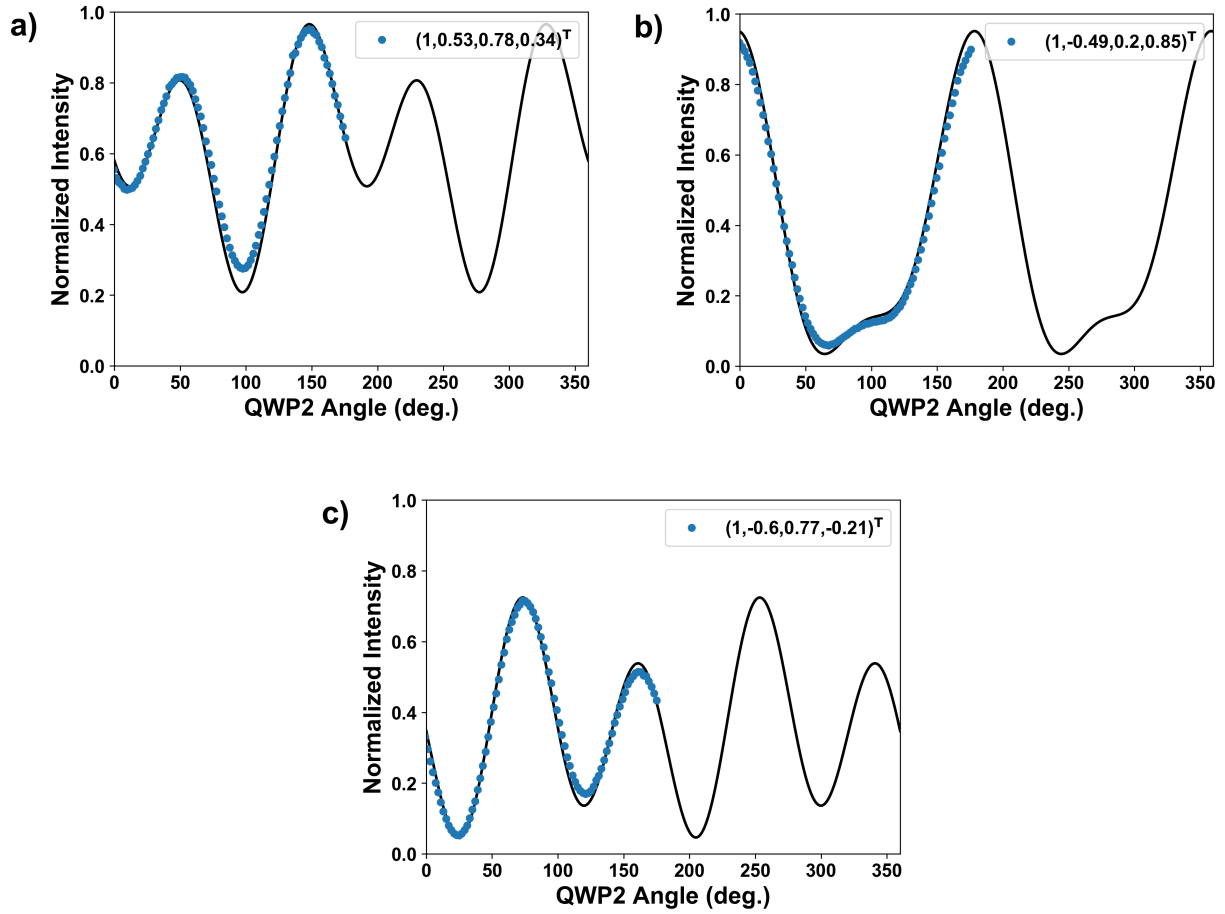


Figure 5.9: Measured signatures generated by the QWP-polariser-photodiode combination for three randomly-selected elliptical states. Written as Stokes vectors, these states are: a) $(1, 0.53, 0.78, 0.34)^T$, b) $(1, -0.49, 0.2, 0.85)^T$, and c) $(1, -0.6, 0.77, -0.21)^T$. All were measured at low temperatures and are displayed with the simulated signatures (using the real waveplate retardances).

proportion of the light which is polarised in the horizontal/vertical, diagonal/anti-diagonal, and right/left circular bases respectively. Written as Stokes vectors, the randomly-chosen elliptical states are: $(1, 0.53, 0.78, 0.34)^T$, $(1, -0.49, -0.2, 0.85)^T$, and $(1, -0.6, 0.77, -0.21)^T$. After performing calculations of the HWP and QWP1 orientations required to produce these states, we were able to generate them at the fibre output, as shown in Figure 5.9. The measurement of these *randomly-selected* elliptical states, in addition to the six standard states, appears to prove that our system is able to deliver any arbitrary polarisation to the refrigerator by fibre. To our knowledge, this is the first reported instance of arbitrary polarisation being delivered in a controlled manner to a dilution refrigerator by optical fibre.

The low-temperature verification tests were conducted at a temperature of 240 mK due to heating caused by QWP2’s piezoelectric rotator. Once the polarisation identification is complete, the refrigerator returns to base temperature as the rotator is no longer needed and the laser intensity can be reduced. Now that we have confirmed that the system works, future tests may be conducted using slower QWP2 rotation to reduce heating and avoid any resulting changes in fibre birefringence. (Note however that the initial polarisation maps were produced at temperatures of approximately 140 mK, while the polarisation states were tested at 240 mK. This would seem to indicate that a temperature change of around 100 mK has only a minimal impact on the state of the fibre.)

5.5 Stability of Fibre Compensation System

As a measure of the system’s stability, the fidelity of the horizontal state was monitored over a period of approximately 35 hours at low temperatures, as shown in Figure 5.10a). The monitoring process involved continuous acquisition of the polarisation signature, followed by fitting of each plot to determine the state. Any change in the fibre’s birefringence would be observable as a change in the fidelity of the polarisation.

During our observation period, the fidelity of the horizontal state remained above 0.955 and was very stable. Changes in the mixing chamber temperature which occurred during the monitoring period (Fig. 5.10b) caused a shift in the polarisation from very nearly horizontal, i.e. $(1, 1, 0, 0)^T$, to the slightly elliptical state $(1, 0.97, 0.05, 0.2)^T$. This new state remained highly stable over the next 24 hours.

While the accidental temperature spike up to nearly 1 K should certainly be expected to have altered the system, the initial change in the state occurred when the monitoring procedure stopped briefly and the temperature decreased from about 240 mK to 60 mK. While the resulting drop in fidelity was relatively minor, this indicates that a temperature change of just 200 mK may be sufficient to alter the polarisation. This is surprising, however, given that the temperature had dropped down to 60 mK at several points over the days prior to this monitoring period, after the fibre had already been characterised. This may therefore be a phenomenon worth investigating, to determine under what conditions the fibre needs to be re-characterised. Nonetheless, performing the fibre characterisation tests more slowly in the future will lower the temperature and should help preserve the polarisation when the QWP2 rotator stops moving.

It should also be noted that the waveplate angles used to produce the horizontal state were calculated using polarisation maps obtained 2.5 days before we started monitoring. Our system was therefore quite stable over a period of at least 4 days, with no indication of changes at the conclusion of the monitoring period. This degree of stability is likely due to the relatively short length of fibre in our setup (30 m), and the controlled environment of our laboratories. As we saw in Section 1.6, multi-kilometer-long fibres which are exposed to environmental influences such as vibrations and temperature fluctuations experience

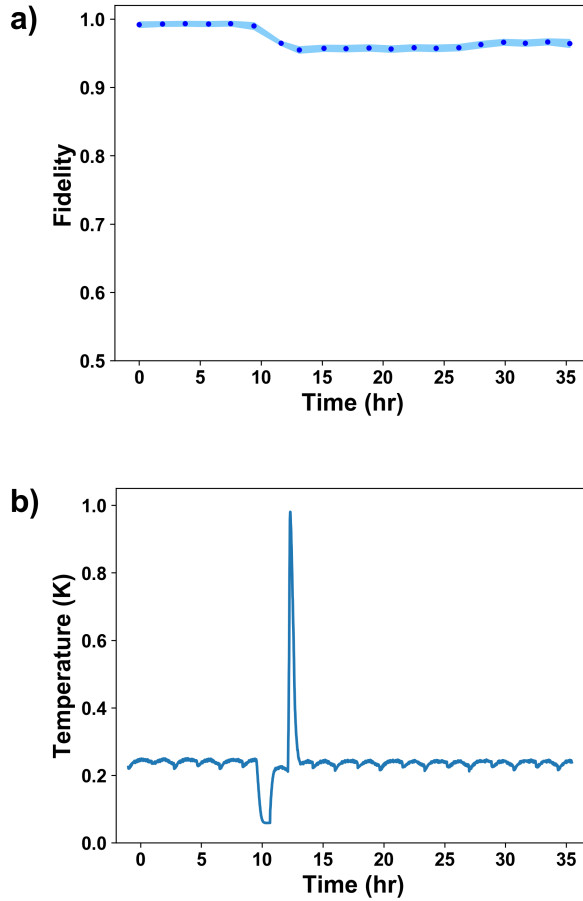


Figure 5.10: a) Fidelity of the horizontal state monitored over 35 hours. The uncertainty of the fidelity (light blue) arises from the standard deviation in the fit of the fibre output state. b) Mixing chamber temperature during monitoring period, including a brief temperature increase to 1K at hour 10.

far greater polarisation instability and may require more frequent fibre characterisation procedures [40, 46, 49]. The stability of our system, however, allows us to conduct low-temperature experiments with polarised light over a period of at least several days without the need to re-characterise the fibre.

5.6 Conclusions and Future Work

In summary, we have achieved successful fibre-based delivery of polarisation states to a dilution refrigerator at low temperatures, and all six standard states were delivered to the fridge with fidelities of >0.96 . Our ability to transmit several randomly-selected elliptical states, in addition to the standard six, indicates that our setup allows for delivery of

any polarisation state to the refrigerator. To achieve this, we have created a compact polarisation readout scheme suitable for installation in cryostats, and demonstrated that our system is very stable over a period of at least four days. We are currently attempting to determine whether we can patent this system, after which we will submit for publication a paper we have written on these results.

This system is designed to correct the unitary transformation of the polarisation by the fibre, but does not prevent depolarisation. Fortunately, very little depolarisation was observed during our experiments, although longer fibres (e.g. tens of km in length) may experience this issue.

Since photons propagating in parallel tend not to interact with each other, the system presented here should function identically when single photons are used instead of a laser beam. The single photon will experience the same effects as each photon in the beam and should therefore have the same state at the fibre output as would be expected for the photons in a beam. After characterising the fibre using a laser beam, future experiments will involve the use of a single photon source for delivery of individual polarised photons at cryogenic temperatures, which is a major component of the photon-to-spin experiment. For the moment, however, we have at least demonstrated that it is possible to transmit polarised beams via fibre.

Finally, it should be noted that this method of delivering polarised states via optical fibre always uses horizontally-polarised light as its input.⁸ It also requires advance knowledge of the states which are to be delivered to the fridge, and without this information, it is not possible to properly set the waveplates so as to compensate for the fibre. In many quantum computing and quantum communication applications, however, qubits are held in superposition states, which would be destroyed by measurement. We would therefore be unable to properly set the waveplates without destroying the qubit.

For such systems, it would be more convenient to have a single waveplate setting which could completely compensate for the fibre, regardless of the input state. Mathematically, the combined waveplates in such a system would form the matrix F^{-1} , the inverse of the fibre matrix F . Simple math, however, shows that this cannot be achieved with just two waveplates in front of the fibre. Recall that the fibre is represented by:

$$F = R(-\theta_a)D(\phi)R(\theta_b) \tag{5.5}$$

and the waveplates are represented by:

$$QWP1 \cdot HWP = R(-\theta_2)D(\pi/2)R(\theta_2)R(-\theta_1)D(\pi)R(\theta_1) \tag{5.6}$$

⁸It need not necessarily be horizontal; any polarisation will do. The point is that we are always transforming the input state into the desired output state, rather than preserving the input state at the fibre output.

We know that we must have at least one HWP and one QWP in order to produce all desired states, so these are the minimum requirements. Noting that $R(-\theta)R(\theta) = \mathbb{I}$, suppose that we set $\theta_2 = \theta_b$ and $\theta_1 = \theta_a$. In doing this, it becomes immediately evident that multiplying the fibre and waveplate matrices together is unlikely to yield unity in every instance:

$$\begin{aligned} F \cdot QWP1 \cdot HWP &= R(-\theta_a)D(\phi)R(\theta_b)R(-\theta_b)D(\pi/2)R(\theta_b)R(-\theta_a)D(\pi)R(\theta_a) \\ &= R(-\theta_a)D(\phi)D(\pi/2)R(\theta_b)R(-\theta_a)D(\pi)R(\theta_a) \end{aligned} \quad (5.7)$$

Our goal is to achieve $D(\phi)D(\pi/2)R(\theta_b)R(-\theta_a)D(\pi) = \mathbb{I}$, but this is unlikely because all of the variables are already fixed by the fibre.⁹ It is therefore impossible for just two waveplates to compensate for the fibre and deliver any unknown input polarisation (i.e. $QWP1 \cdot HWP \neq F^{-1}$ in general). Further calculations reveal that at least four waveplates are required in order to generate F^{-1} . Preliminary simulations suggest that the addition of these two extra waveplates would enable our setup to send any state through the fibre with a single waveplate configuration and without requiring foreknowledge of the states being transmitted. Experimental tests of this hypothesis will likely commence shortly.

⁹A quick calculation shows that we would need $\theta_a = \theta_b$ and $\phi = -3\pi/2$ to achieve the identity matrix. (Some other strange combination of values might work also, but the point is that these variables are set by the fibre's birefringence and are unlikely to have the perfect values required for $F \cdot QWP1 \cdot HWP = \mathbb{I}$.)

Chapter 6

Conclusions and Outlook

While waiting for functioning photon-to-spin DQD samples over the past two years, a lot of work has been put into the NRC's SPIN lab to ensure its components are operational when samples arrive. As discussed in this thesis, much of the work has been focused on the optical systems, which are now ready for use as DQD samples become available. In my capacity as a student at this lab, I have assisted in the creation of a sample imaging system for delivering light to nanostructures, the testing of nanowire quantum dots as fibre-coupled sources of single photons, and the demonstration that polarised light can be delivered by fibre to a dilution refrigerator at cryogenic temperatures. I have also worked on the construction and preliminary testing of the RF-QPC for charge detection of photo-generated electrons in quantum dots.

Conclusions

In writing custom driver and experiment software for the piezoelectric positioners in the fridge, I have developed a system capable of imaging nanostructure samples with sub-micron resolution. This has required a lot of trial-and-error tests to understand and adapt to the behaviour and limitations of the positioners, such as the variability and hysteresis in their step sizes. Sample imaging can now be achieved at cryogenic temperatures in the dilution refrigerator using photocurrent measurements, reflectance measurements, or measurements of photoluminescence (PL) emissions. These procedures allow one to observe the sample topography and position the lensed fibre at the appropriate location for focused light delivery. Recent optical scans of DQD samples provided by Osaka University suggest that we can now use this system to locate the 400nm-wide hole in the optical mask. This is an important requirement for photon delivery in the photon-to-spin project, and further tests will be conducted to verify that we can actually focus the fibre on the hole.

This new imaging procedure has been used extensively for tests of nanowire quantum dot devices, which were grown by Dr. Dan Dalacu of NRC. By imaging these samples, we were able to locate specific nanowires and focus the lensed fibre on them, collecting

emissions from just one individual wire at a time. The collection efficiency of the fibre is low, being just 1–2%, but we are still able to achieve sufficient PL intensities for analysis of the nanowire emission spectra. Tests of both single-dot and double-dot nanowires revealed that greater photon indistinguishability could be achieved at lower temperatures and pump powers. A Hong-Ou-Mandel interferometry experiment or direct measurement of emission linewidths via Fabry-Pérot interferometry would provide a more quantitative measure of photon indistinguishability, so we may pursue these tests in the future. In the meantime, we have managed to demonstrate that nanowire quantum dots can be used as sources of fibre-coupled single photons, with a measured $g^{(2)}(0)$ of 0.095 ± 0.027 . Pulsed measurements suggest that we could use this system as an on-demand single photon source, and while we have not demonstrated this due to laser mode-locking issues, others have proven it to be true using these same samples and free-space collection methods. The measured $g^{(2)}(0)$ value surpasses the threshold required for a single photon source to have practical applications, although the fibre collection efficiency likely needs to be increased before such a prospect can be considered realistic. To improve the coupling between the fibre and wire, Dr. Dalacu is now testing a GRIN lens in a 4 K cryostat. This system may therefore be used as our fibre-coupled, single-photon source for tests of the photon-to-spin conversion interface.

One of the largest obstacles to delivering these photons by fibre, however, is the polarisation-altering effects of the fibre’s birefringent core. In testing photon-to-spin conversion, we will need some method of reliably delivering specific polarisation states to the quantum dot sample so we can verify that the appropriate spin states have been generated. Using standard optical components and a custom-made genetic algorithm, we have developed a system which characterises the behaviour of light in the fibre and compensates for the fibre’s birefringent effects. Tests of this setup have shown that we can deliver the desired polarisation states to the fridge at both room temperature and cryogenic temperatures. We have shown that the standard states (horizontal, vertical, diagonal, anti-diagonal, right circular, and left circular) can be delivered to the sample holder at low temperatures with fidelities of $>0.96 \pm 0.007$. To prove that any arbitrary polarisation can be delivered with this setup, three randomly-selected elliptical states were also successfully transmitted to the fridge by fibre. Monitoring this system has shown it to be highly stable over a period of at least four days, although temperature changes of a few hundred mK can cause the strain on the fibre to change and the fidelity to drop slightly. The results of our polarisation tests will be submitted for publication after our application to patent the system has been reviewed.

With these tests completed, the SPIN Lab’s optical systems are now ready for photon-to-spin tests.

Next Steps

With the development of the SPIN Lab’s optical systems now largely complete, we can begin to focus on testing new lateral DQD samples provided by Professor Oiwa of Osaka

University. The next steps include verifying that these samples work as expected and that we can detect photo-generated electrons in the quantum dots. We would also ideally like to demonstrate that this system can indeed enable the transfer of fibre-delivered photonic qubits to electron spin states. Such work would require a functioning RF-QPC charge sensor, a method of transmitting superposition states to the dilution refrigerator via fibre, and a source of single photons. As we have seen in this thesis, preliminary testing of all of these components has now been concluded, and we are ready to begin putting them together for a single purpose.

In the future, as the photon-to-spin system is developed, we may begin testing hole-trapping samples instead, which have more-promising properties (see Section 1.7.3). It may also be that time-bin encoding of photonic qubits is easier to implement in our system, rather than relying upon polarisation states in fibres. We may therefore test this option as well.

Once the proof-of-concept has been demonstrated for the photon-to-spin conversion interface, the next stage could be to experimentally verify whether this system can actually be used as part of a viable quantum repeater network [2]. To do so, we would need to prove that our process for transmitting polarised photons by fibre will still work for long fibres stretched out over many kilometers.¹ Next, entanglement swapping would need to be proven, with full Bell state measurements performed on the photo-generated electron spins. The potential for complete Bell state measurements on spin qubits in lateral quantum dots is still theoretical [56], so this could be another important area of research. The final step in creating the photon-to-spin hybrid repeater would be to perform quantum teleportation of superposition states over long distances (>100 km). These goals are still a few years distant, but I remain excited at the prospect of working towards them.

¹The unitary transformations performed by the fibre on the photons would likely not destroy entanglement, but they would redefine the definitions of states $|0\rangle$ and $|1\rangle$. As foreknowledge of the entangled state is required for successful quantum teleportation, however, we will need to use our system to compensate for the effects of the fibre. The future version of the system (described in the conclusions of Chapter 5) will be required, as it should completely reverse the effects of the fibre.

References

- [1] C. Gobby, Z. L. Yuan, and A. J. Shields. Quantum key distribution over 122 km of standard telecom fiber. *Appl. Phys. Lett.*, 84(19):3762–3764, 2004.
- [2] L. Gaudreau, A. Bogan, M. Korkusinski, S. Studenikin, D. G. Austing, and A. S. Sachrajda. Entanglement distribution schemes employing coherent photon-to-spin conversion in semiconductor quantum dot circuits. *Semicond. Sci. Technol.*, 32(093001), 2017.
- [3] M. Ciorga, A. S. Sachrajda, P. Hawrylak, C. Gould, P. Zawadzki, S. Jullian, Y. Feng, and Z. Wasilewski. Addition spectrum of a lateral dot from Coulomb and spin-blockade spectroscopy. *Phys. Rev. B*, 61(24), 2000.
- [4] John Davies. *The Physics of Low-Dimensional Semiconductors: An Introduction*. Cambridge University Press, New York, New York, 1998.
- [5] D. Dalacu, K. Mnaymneh, X. Wu, J. Lapointe, G. C. Aers, P. J. Poole, and R. L. Williams. Selective-area vapor-liquid-solid growth of tunable InAsP quantum dots in nanowires. *Appl. Phys. Lett.*, 98(251101), 2011.
- [6] G. Bulgarini, M. E. Reimer, M. B. Bavinck, K. D. Jöns, D. Dalacu, P. J. Poole, E. P. A. M. Bakkers, and V. Zwiller. Nanowire waveguides launching single photons in a gaussian mode for ideal fiber coupling. *Nano Lett.*, 14(7):4102–4106, 2014.
- [7] M. Dahl, L. Namazi, R. R. Zamani, and K. A. Dick. Sb incorporation in wurtzite and zinc blende InAs_{1-x}Sb_x branches on InAs template nanowires. *Small*, 14(1703785), 2018.
- [8] P. Staudinger, M. Sistani, J. Greil, E. Bertagnolli, and A. Lugstein. Ultrascaled germanium nanowires for highly sensitive photodetection at the quantum ballistic limit. *Nano Lett.*, 18(8):5030–5035, 2018.
- [9] G. Piret, M.-T. Perez, and C. N. Prinz. Support of neuronal growth over glial growth and guidance of optic nerve axons by vertical nanowire arrays. *ACS Appl. Mater. Interfaces*, 7(34):18944–18948, 2015.

- [10] D. Dalacu, P. J. Poole, and R. L. Williams. Nanowire-based sources of non-classical light. *Nanotechnology*, 30(232001), 2019.
- [11] D. Dalacu, K. Mnaymneh, J. Lapointe, X. Wu, P. J. Poole, G. Bulgarini, V. Zwiller, and M. E. Reimer. Ultraclean emission from InAsP quantum dots in defect-free wurtzite InP nanowires. *Nano Lett.*, 12(11):5919–5923, 2012.
- [12] P. Kaer, N. Gregersen, and J. Mørk. The role of phonon scattering in the indistinguishability of photons emitted from semiconductor cavity QED systems. *New J. Phys.*, 15(035027), 2013.
- [13] P. Lodahl, S. Mahmoodian, and S. Stobbe. Interfacing single photons and single quantum dots with photonic nanostructures. *Rev. Mod. Phys.*, 87:347–400, 2015.
- [14] K. D. Jöns, L. Schweickert, M. A. M. Versteegh, D. Dalacu, P. J. Poole, A. Gulinatti, A. Giudice, V. Zwiller, and M. E. Reimer. Bright nanoscale source of deterministic entangled photon pairs violating Bell’s inequality. *Sci. Rep.*, 7(1700), 2017.
- [15] M. E. Reimer, G. Bulgarini, N. Akopian, M. Hocevar, M. B. Bavinck, M. A. Verheijen, E. P. A. M. Bakkers, L. P. Kouwenhoven, and V. Zwiller. Bright single-photon sources in bottom-up tailored nanowires. *Nat. Commun.*, 3(737), 2012.
- [16] N. Gregersen, T. R. Nielsen, J. Claudon, J.-M. Gérard, and Jesper Mørk. Controlling the emission profile of a nanowire with a conical taper. *Opt. Lett.*, 33(15):1693–1695, 2008.
- [17] M. Friedl, K. Cervený, P. Weigele, G. Tütüncüoğlu, S. Martí-Sánchez, C. Huang, T. Patlatiuk, H. Potts, Z. Sun, M. O. Hill, L. Güniat, W. Kim, M. Zamani, V. G. Dubrovskii, J. Arbiol, L. J. Lauhon, D. M. Zumbühl, and A. Fontcuberta i Morral. Template-assisted scalable nanowire networks. *Nano Lett.*, 18(4):2666–2671, 2018.
- [18] D. Dalacu, A. Kam, D. G. Austing, X. Wu, J. Lapointe, G. C. Aers, and P. J. Poole. Selective-area vapour–liquid–solid growth of InP nanowires. *Nanotechnology*, 20(39):395602, 2009.
- [19] Philip J. Poole. Personal communication, March 2019.
- [20] L. E. Jensen, M. T. Björk, S. Jeppesen, A. I. Persson, B. J. Ohlsson, and L. Samuelson. Role of surface diffusion in chemical beam epitaxy of InAs nanowires. *Nano Lett.*, 4(10):1961–1964, 2004.
- [21] M. Cardona P. Yu. *Fundamentals of Semiconductors: Physics and Materials Properties, Fourth Edition*. Springer, New York, New York, 2010.
- [22] D. Bimberg, M. Grundmann, and NN Ledentsov. *Quantum Dot Heterostructures*. John Wiley & Sons, Chichester, UK, 1999.

- [23] E. Poem, Y. Kodriano, C. Tradonsky, N. H. Lindner, B. D. Gerardot, P. M. Petroff, and D. Gershoni. Accessing the dark exciton with light. *Nat. Phys.*, 6(12):993–997, 2010.
- [24] S. Haffouz, K. D. Zeuner, D. Dalacu, P. J. Poole, J. Lapointe, D. Poitras, K. Mnaymneh, M. Couillard X. Wu, M. Korkusinski, E. Schöll, K. D. Jöns, V. Zwiller, and R. L. Williams. Bright single InAsP quantum dots at telecom wavelengths in position-controlled InP nanowires: The role of the photonic waveguide. *Nano Lett.*, 18(5):3047–3052, 2018.
- [25] M. Bayer, G. Ortner, O. Stern, A. Kuther, A. A. Gorbunov, A. Forchel, P. Hawrylak, S. Fafard, K. Hinzer, T. L. Reinecke, S. N. Walck, J. P. Reithmaier, F. Klopff, and F. Schäfer. Fine structure of neutral and charged excitons in self-assembled In(Ga)As/(Al)GaAs quantum dots. *Phys. Rev. B*, 65(195315), 2002.
- [26] K. Mnaymneh, D. Dalacu, J. McKee, J. Lapointe, S. Haffouz, J. F. Weber, D B. Northeast, P. J. Poole, G. C. Aers, and R. L. Williams. On-chip integration of single photon sources via evanescent coupling of tapered nanowires to SiN waveguides. *Adv. Quantum Technol.*, 1900021, 2019.
- [27] J. Heinrich, A. Huggenberger, T. Heindel, S. Reitzenstein, S. Höfling, L. Worschech, and A. Forchel. Single photon emission from positioned GaAs/AlGaAs photonic nanowires. *Appl. Phys. Lett.*, 96(211117), 2010.
- [28] I. Aharonovich, D. Englund, and M. Toth. Solid-state single-photon emitters. *Nat. Photonics*, 10:631–641, 2016.
- [29] J. Mason. *RF-QPC Charge Detector and S-T₊ Qubit in a Lateral Double Quantum Dot Device*. PhD thesis, University of Waterloo, 2016.
- [30] R. J. Schoelkopf, P. Wahlgren, A. A. Kozhevnikov, P. Delsing, and D. E. Prober. The radio-frequency single-electron transistor (RF-SET): A fast and ultrasensitive electrometer. *Science*, 280(5367):1238–1242, 1998.
- [31] T. Müller, T. Choi, S. Hellmüller, K. Ensslin, T. Ihn, and S. Schön. A circuit analysis of an in situ tunable radio-frequency quantum point contact. *Rev. Sci. Instrum.*, 84(083902), 2013.
- [32] D. J. Reilly, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Fast single-charge sensing with a rf quantum point contact. *Appl. Phys. Lett.*, 91(162101), 2007.
- [33] M. C. Cassidy, A. S. Dzurak, R. G. Clark, K. D. Petersson, I. Farrer, D. A. Ritchie, and C. G. Smith. Single shot charge detection using a radio-frequency quantum point contact. *Appl. Phys. Lett.*, 91(222104), 2007.
- [34] M. F. Gonzalez-Zalba, S. Barraud, A. J. Ferguson, and A. C. Betz. Probing the limits of gate-based charge sensing. *Nat. Commun.*, 6(6084), 2015.

- [35] Kenneth L. Kaiser. *Electromagnetic Compatibility Handbook*. CRC Press, 2004. ISBN 0-8493-2087-9.
- [36] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, New York, New York, 2014.
- [37] Robert Mann. *An Introduction to Particle Physics and the Standard Model*. CRC Press, Boca Raton, Florida, 2010.
- [38] Eugene Hecht. *Optics, Fourth Edition*. Pearson Education, San Francisco, California, 2002.
- [39] Kailash K. Sharma. *Optics: Principles and Applications*. Academic Press, 2006. ISBN 978-0-12-370611-9.
- [40] Y.-Y. Ding, H. Chen, S. Wang, D.-Y. He, Z.-Q. Yin, W. Chen, Z. Zhou, G.-C. Guo, and Z.-F. Han. Polarization variations in installed fibers and their influence on quantum key distribution systems. *Opt. Express*, 25(22):27923–27936, 2017.
- [41] S. Manhas, J. Vizet, S. Deby, J.-C. Vanel, P. Boito, M. Verdier, A. De Martino, and D. Pagnoux. Demonstration of full 4x4 Mueller polarimetry through an optical fiber for endoscopic applications. *Opt. Express*, 23(3):3047–3054, 2015.
- [42] A. N. Pinto, A. J. Almeida, N. A. Silva, N. J. Muga, and L. M. Martins. Optical quantum communications: an experimental approach. In *Proceedings of SPIE*, volume 8001, 80011M, 2011.
- [43] O. Gamel and D. F. V. James. Measures of quantum state purity and classical degree of polarization. *Phys. Rev. A*, 86(033830), 2012.
- [44] G. B. Xavier, G. V. de Faria, G. P. Temporão, and J. P. von der Weid. Full polarization control for fiber optical quantum communication systems using polarization encoding. *Opt. Express*, 16(3):1867–1873, 2008.
- [45] G. B. Xavier, N. Walenta, G. V. de Faria, G. P. Temporão, N. Gisin, H. Zbinden, and J. P. von der Weid. Experimental polarization encoded quantum key distribution over optical fibres with real-time continuous birefringence compensation. *New J. Phys.*, 11(045015), 2009.
- [46] J. Chen, G. Wu, Y. Li, E. Wu, and H. Zeng. Active polarization stabilization in optical fibers suitable for quantum key distribution. *Opt. Express*, 15(26):17928–17936, 2007.
- [47] N. Gisin, J.-P. von der Weid, and J.-P. Pellaux. Polarization mode dispersion of short and long single-mode fibers. *J. Light. Technol.*, 9(7):821–827, 1991.
- [48] F. Treviño Martínez, D. Tentori, C. Ayala-Díaz, and F. J. Mendieta-Jiménez. Birefringence assessment of single-mode optical fibers. *Opt. Express*, 13(7):2556–2563, 2005.

- [49] J. F. Dynes, I. Choi, A. W. Sharpe, A. R. Dixon, Z. L. Yuan, M. Fujiwara, M. Sasaki, and A. J. Shields. Stability of high bit rate quantum key distribution on installed fiber. *Opt. Express*, 20(15):16339–16347, 2012.
- [50] M. J. Nelson, C. J. Collins, and C. C. Speake. A cryogenic optical feedthrough using polarization maintaining fibers. *Rev. Sci. Instrum.*, 87(033111), 2016.
- [51] OZ Optics. *DTS0011 Electrically Driven Polarisation Controller-Scrambler*, October 2016. EPC data sheet.
- [52] A. Pozas-Kerstjens and E. Martín-Martínez. Entanglement harvesting from the electromagnetic vacuum with hydrogenlike atoms. *Phys. Rev. D*, 94(064074), 2016.
- [53] David H. McIntyre. *Quantum Mechanics*. Pearson Education, San Francisco, California, 2012.
- [54] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390:575–579, 1997.
- [55] N. Lütkenhaus, J. Calsamiglia, and K.-A. Suominen. Bell measurements for teleportation. *Phys. Rev. A*, 59(5):3295–3300, 1999.
- [56] N. Yokoshi, H. Imamura, and H. Kosaka. Proposal of a full Bell state analyzer for spin qubits in a double quantum dot. *Phys. Rev. B*, 81(161305), 2010.
- [57] K. Kuroyama, M. Larsson, S. Matsuo, T. Fujita, S. R. Valentin, A. Ludwig, A. D. Wieck, A. Oiwa, and S. Tarucha. Single electron-photon pair creation from a single polarization-entangled photon pair. *Sci. Rep.*, 7(16968), 2017.
- [58] J. I. Colless, A. C. Mahoney, J. M. Hornibrook, A. C. Doherty, H. Lu, A. C. Gossard, and D. J. Reilly. Dispersive readout of a few-electron double quantum dot with fast rf gate sensors. *Phys. Rev. Lett.*, 110(046805), 2013.
- [59] J. M. Hornibrook, J. I. Colless, A. C. Mahoney, X. G. Croot, S. Blanvillain, H. Lu, A. C. Gossard, and D. J. Reilly. Frequency multiplexing for readout of spin qubits. *Appl. Phys. Lett.*, 104(103108), 2014.
- [60] T. Müller, B. Küng, S. Hellmüller, P. Studerus, K. Ensslin, T. Ihn, M. Reinwald, and W. Wegscheider. An in situ tunable radio-frequency quantum point contact. *Appl. Phys. Lett.*, 97(202104), 2010.
- [61] M/A-COM Technology Solutions. *MA46410 thru MA46480 Series*. GaAs Hyperabrupt Varactor Diode technical specifications sheet, Rev. V7.
- [62] Bart Van Zeghbroeck. *Principles of Semiconductor Devices and Heterojunctions*. Pearson Education, London, United Kingdom, 2013.

- [63] Attocube Systems. *User Manual - Premium Line: Positioners & Scanners*, March 2015. Version 2.0.
- [64] Attocube Systems. *Closed loop, linear, horizontal stepper positioner with resistive encoder*, April 2013. ANPx101/RES technical specifications sheet.
- [65] Attocube Systems. *Closed loop, linear, vertical stepper positioner with resistive encoder*, April 2013. ANPz101/RES technical specifications sheet.
- [66] Attocube Systems. *Closed loop, rotary stepper positioner with vertical rotation axis and resistive encoder*, February 2014. ANR240/RES technical specifications sheet.
- [67] A. Fiset-Cyr, D. Dalacu, S. Haffouz, P. J. Poole, J. Lapointe, G. C. Aers, and R. L. Williams. In-situ tuning of individual position-controlled nanowire quantum dots via laser-induced intermixing. *Appl. Phys. Lett.*, 113(053105), 2018.
- [68] Y. Yamamoto and K. Semba. *Principles and Methods of Quantum Information Technologies, Lecture Notes in Physics 911*. Springer, Tokyo, Japan, 2015.
- [69] M. E. Reimer, G. Bulgarini, A. Fognini, R. W. Heeres, B. J. Witek, M. A. M. Versteegh, A. Rubino, T. Braun, M. Kamp, S. Höfling, D. Dalacu, J. Lapointe, P. J. Poole, and V. Zwiller. Overcoming power broadening of the quantum dot emission in a pure wurtzite nanowire. *Phys. Rev. B*, 93(195316), 2016.
- [70] T. M. Babinec, B. J. M. Hausmann, M. Khan, Y. Zhang, J. R. Maze, P. R. Hemmer, and Marko Lončar. A diamond nanowire single-photon source. *Nat. Nanotechnol.*, 5:195–199, 2010.
- [71] O. Gazzano, S. Michaelis de Vasconcellos, C. Arnold, A. Nowak, E. Galopin, I. Sagnes, L. Lanco, A. Lemaître, and P. Senellart. Bright solid-state sources of indistinguishable single photons. *Nat. Commun.*, 4(1425), 2013.
- [72] L. Hanschke, K. A. Fischer, S. Appel, D. Lukin, J. Wierzbowski, S. Sun, R. Trivedi, J. Vučković, J. J. Finley, and K. Müller. Quantum dot single-photon sources with ultra-low multi-photon probability. *npj Quantum Inf.*, 4(43), 2018.
- [73] J. Claudon, J. Bleuse, N. Singh Malik, M. Bazin, P. Jaffrennou, N. Gregersen, C. Sauvan, P. Lalanne, and J.-M. Gérard. A highly efficient single-photon source based on a quantum dot in a photonic nanowire. *Nat. Photonics*, 4(3):174–177, 2010.
- [74] S. Sun, H. Kim, G. S. Solomon, and E. Waks. Cavity-enhanced optical readout of a single solid-state spin. *Phys. Rev. Appl.*, 9(054013), 2018.
- [75] Dan Dalacu. Personal communication, August 2019.
- [76] J.-M. Jancu, K. Gauthron, L. Largeau, G. Patriarche, J.-C. Harmand, and P. Voisin. Type II heterostructures formed by zinc-blende inclusions in InP and GaAs wurtzite nanowires. *Appl. Phys. Lett.*, 97(041910), 2010.

- [77] C. Carlson, D. Dalacu, C. Gustin, S. Haffouz, X. Wu, J. Lapointe, R. L. Williams, P. J. Poole, and S. Hughes. Theory and experiments of coherent photon coupling in semiconductor nanowire waveguides with quantum dot molecules. *Phys. Rev. B*, 99(085311), 2019.
- [78] M. Khoshnegar, T. Huber, A. Predojević, D. Dalacu, M. Prilmüller, J. Lapointe, X. Wu, P. Tamarat, B. Lounis, P. Poole, G. Weihs, and H. Majedi. A solid state source of photon triplets based on quantum dot molecules. *Nat. Commun.*, 8(15716), 2017.
- [79] M. Bayer, P. Hawrylak, K. Hinzer, S. Fafard, M. Korkusinski, Z. R. Wasilewski, O. Stern, and A. Forchel. Coupling and entangling of quantum states in quantum dot molecules. *Science*, 291(5503):451–453, 2000.
- [80] B. J. Witek, R. W. Heeres, U. Perinetti, E. P. A. M. Bakkers, L. P. Kouwenhoven, and V. Zwiller. Measurement of the g-factor tensor in a quantum dot and disentanglement of exciton spins. *Phys. Rev. B*, 84(195305), 2011.
- [81] P. Borri, W. Langbein, S. Schneider, U. Woggon, R. L. Sellin, D. Ouyang, and D. Bimberg. Ultralong dephasing time in InGaAs quantum dots. *Phys. Rev. Lett.*, 87(157401), 2001.
- [82] J. Iles-Smith, D. P. S. McCutcheon, A. Nazir, and J. Mørk. Phonon scattering inhibits simultaneous near-unity efficiency and indistinguishability in semiconductor single-photon sources. *Nat. Photonics*, 11:521–526, 2017.
- [83] T. Grange, N. Somaschi, C. Antón, L. De Santis, G. Coppola, V. Giesz, A. Lemaître, I. Sagnes, A. Auffèves, and P. Senellart. Reducing phonon-induced decoherence in solid-state single-photon sources with cavity quantum electrodynamics. *Phys. Rev. Lett.*, 118(253602), 2017.
- [84] Guoqiang Cui and M. G. Raymer. Emission spectra and quantum efficiency of single-photon sources in the cavity-QED strong-coupling regime. *Phys. Rev. A*, 73(053807), 2006.
- [85] E. A. Muljarov and R. Zimmermann. Dephasing in quantum dots: Quadratic coupling to acoustic phonons. *Phys. Rev. Lett.*, 93(237401), 2004.
- [86] G. Muñoz-Matutano, D. Barrera, C. R. Fernández-Pousa, R. Chulia-Jordan, L. Seravalli, G. Trevisi, P. Frigeri, S. Sales, and J. Martínez-Pastor. All-optical fiber Hanbury Brown & Twiss interferometer to study 1300 nm single photon emission of a metamorphic InAs quantum dot. *Sci. Rep.*, 6(27214), 2016.
- [87] R. Kalra, A. Laucht, J. P. Dehollain, D. Bar, S. Freer, S. Simmons, J. T. Muhonen, and A. Morello. Vibration-induced electrical noise in a cryogen-free dilution refrigerator: Characterization, mitigation, and impact on qubit coherence. *Rev. Sci. Instrum.*, 87(073905), 2016.

APPENDICES

Appendix A

Quantum Teleportation via Bell State Measurements

In Chapter 1, we reviewed the procedure for quantum teleportation:

1. Particles A and B, which form an entangled pair, are separated.
2. A Bell state measurement (BSM) performed on A and a third particle C immediately transfers the quantum state of C to B, modulo a unitary transformation.
3. Classical channels are used to communicate the results of the BSM so that B can be transformed into the original state of C.

In this Appendix, we will look a little closer at step 2 and attempt to understand the mechanism by which the BSM assists in the teleportation of the quantum state.

Recall our teleportation example from Chapter 1, where we had qubits A and B in entangled Bell state $|\Psi^+\rangle_{AB}$, while particle C was in some unknown state $|\psi\rangle_C = \alpha|0\rangle_C + \beta|1\rangle_C$. This three-particle system can be written as:

$$|\Psi^+\rangle_{AB} |\psi\rangle_C = \frac{1}{\sqrt{2}}(\alpha|010\rangle_{ABC} + \alpha|100\rangle_{ABC} + \beta|011\rangle_{ABC} + \beta|101\rangle_{ABC}) \quad (\text{A.1})$$

In Section 1.7.1, we showed that we can easily rewrite this equation as:

$$\begin{aligned} |\Psi^+\rangle_{AB} |\psi\rangle_C = & \frac{1}{2} [|\Psi^+\rangle_{AC} (\beta|1\rangle_B + \alpha|0\rangle_B) \\ & + |\Psi^-\rangle_{AC} (\beta|1\rangle_B - \alpha|0\rangle_B) \\ & + |\Phi^+\rangle_{AC} (\alpha|1\rangle_B + \beta|0\rangle_B) \\ & + |\Phi^-\rangle_{AC} (\alpha|1\rangle_B - \beta|0\rangle_B)] \end{aligned} \quad (\text{A.2})$$

where we have simply collected together certain terms; no actual operations have been performed on the system yet. This equation does make it appear as if we are actually entangling particles A and C, whereas we know from Section 1.7.1 that a BSM actually destroys entanglement, merely revealing the parity and relative phase of two qubits occupying a Bell state. So how can the initial quantum state of C be transferred to particle B if we are not entangling C with A and B?

To answer this question, let's return to Equation A.1. For simplicity, let's relabel the three-particle state $|\Psi^+\rangle_{AB} |\psi\rangle_C$ as $|X\rangle$. Recall from Section 1.7.1 that a BSM can be performed through the application of a CNOT gate, followed by a Hadamard gate, and then direct measurement of the qubit states.¹ (Other BSM methods exist, but this one yields some easily-interpreted results.) Let's perform these operations directly on $|X\rangle$, starting with the CNOT gate, which uses A as the control qubit and C as the target.

$$U_{\text{CNOT}}^{\text{AC}} |X\rangle = \frac{1}{\sqrt{2}}(\alpha |010\rangle_{\text{ABC}} + \alpha |101\rangle_{\text{ABC}} + \beta |011\rangle_{\text{ABC}} + \beta |100\rangle_{\text{ABC}}) \quad (\text{A.3})$$

Note that the states of A and B are still perfectly anti-correlated after this operation. In fact, we can re-write this equation as:

$$U_{\text{CNOT}}^{\text{AC}} |X\rangle = \frac{1}{\sqrt{2}}[(\alpha |0\rangle_C + \beta |1\rangle_C) |01\rangle_{\text{AB}} + (\alpha |1\rangle_C + \beta |0\rangle_C) |10\rangle_{\text{AB}}] \quad (\text{A.4})$$

While this is not a maximally-entangled state, our inability to completely factor out the state of particle C indicates that there is now some entanglement between all the particles in this system. The CNOT operation also encodes C with information about how the probability amplitudes α and β are assigned to $|01\rangle_{\text{AB}}$ and $|10\rangle_{\text{AB}}$.² The only exception is when $\alpha = \beta$, which eliminates the entanglement between C and the other two particles. This has no effect on the end result, however, as the assignment of α and β no longer matters in that case.

To continue with the BSM procedure, we now apply the Hadamard gate to particle A:

$$\begin{aligned} U_{\text{Had}}^{\text{A}} U_{\text{CNOT}}^{\text{AC}} |X\rangle &= \frac{1}{2} [\alpha(|0\rangle_{\text{A}} + |1\rangle_{\text{A}}) |10\rangle_{\text{BC}} + \alpha(|0\rangle_{\text{A}} - |1\rangle_{\text{A}}) |01\rangle_{\text{BC}} \\ &\quad + \beta(|0\rangle_{\text{A}} + |1\rangle_{\text{A}}) |11\rangle_{\text{BC}} + \beta(|0\rangle_{\text{A}} - |1\rangle_{\text{A}}) |00\rangle_{\text{BC}}] \\ &= \frac{1}{2} [\alpha(|010\rangle_{\text{ABC}} + |110\rangle_{\text{ABC}} + |001\rangle_{\text{ABC}} - |101\rangle_{\text{ABC}}) \\ &\quad + \beta(|011\rangle_{\text{ABC}} + |111\rangle_{\text{ABC}} + |000\rangle_{\text{ABC}} - |100\rangle_{\text{ABC}})] \end{aligned} \quad (\text{A.5})$$

¹See Section 1.7.1 for definitions of these two gates.

²For example, a measurement of C which yields $|0\rangle_C$ will result in a state $\alpha |01\rangle_{\text{AB}} + \beta |10\rangle_{\text{AB}}$, while the measurement result $|1\rangle_C$ forces A and B into the state $\alpha |10\rangle_{\text{AB}} + \beta |01\rangle_{\text{AB}}$.

With this operation complete, A and B are no longer maximally entangled (i.e. they aren't in a Bell state), although they are still partially entangled. It is difficult to see in this equation, but it now happens that A holds phase information about the final state of B. If we rearrange the terms and separate particle B from A and C, we find:

$$\begin{aligned}
 U_{\text{Had}}^A U_{\text{CNOT}}^{\text{AC}} |X\rangle = & \frac{1}{2} [|01\rangle_{\text{AC}} (\beta |1\rangle_{\text{B}} + \alpha |0\rangle_{\text{B}}) \\
 & + |11\rangle_{\text{AC}} (\beta |1\rangle_{\text{B}} - \alpha |0\rangle_{\text{B}}) \\
 & + |00\rangle_{\text{AC}} (\alpha |1\rangle_{\text{B}} + \beta |0\rangle_{\text{B}}) \\
 & + |10\rangle_{\text{AC}} (\alpha |1\rangle_{\text{B}} - \beta |0\rangle_{\text{B}})]
 \end{aligned} \tag{A.6}$$

Individual measurements of A and C will now destroy all entanglement and force B into a superposition state. The measured value of A dictates the relative phase of B's superposition state ($|0\rangle_{\text{A}}$ for positive, $|1\rangle_{\text{A}}$ for negative) and the value of C indicates whether α is associated with $|0\rangle_{\text{B}}$ or $|1\rangle_{\text{B}}$. By comparing this equation with Equations 1.56 and A.2, we can clearly see that we have the same result, but the mechanisms by which the BSM acts to teleport a quantum state are now more obvious.

This returns us to the question we asked earlier: how is the qubit C transferred to B if A and C are never actually entangled? As we have seen, C is in fact entangled with the other two. In the first stage C is partially entangled with A and B, which are still maximally entangled with each other. Then the AB state is reduced to partial entanglement, and final measurements of A and C destroy all entanglement while forcing B to take on the initial state of C (modulo a unitary transformation).

The key to understanding the quantum teleportation procedure is in the nature of the BSM itself. Recall from Chapter 1 that the action of the BSM on a pair of entangled particles is to reveal the parity and relative phase of the two possible states. Equation A.6 shows that this procedure has a similar result when applied to one particle of an entangled pair and a third unknown qubit: a BSM of A and C reveals the parity of α/β and $|0\rangle_{\text{B}} / |1\rangle_{\text{B}}$, as well as the relative phase of B. We can therefore think of the BSM during quantum teleportation as a method of comparing the states of A and C. Since A and B are entangled in a known way, we can then use the result of our A/C comparison to transform B into the initial state of C. The BSM also generates partial entanglement between C and the Bell state of A and B, causing the probability amplitudes α and β to be applied to the entire three-particle entangled state. This is how we are able to transfer the unknown state $|\psi\rangle_{\text{C}} = \alpha |0\rangle_{\text{C}} + \beta |1\rangle_{\text{C}}$ from C to B.

Appendix B

Platypus Experiment Automation Software

In Chapter 2, we briefly discussed Platypus, the Python-based equipment control and experiment automation software used in the SPIN Lab. In this Appendix, we will look at it in a little more detail, since it forms such an important part of our work in the lab.

Platypus was originally created by Dr. Alex Bogan while he was a graduate student at NRC. Dr. Louis Gaudreau and myself have contributed a fair amount to the code, particularly with the device and experiment files which are unique to our lab.

B.1 Introduction to Platypus

Platypus is a custom-built, modular software package which acts as a framework for remotely controlling lab equipment and running experiments with near-autonomy. Its main purpose is to enable experiments, especially those involving multiple and repeated procedures, to be run without need for manual operation of the equipment. Additional features include automatic plotting and recording of data while experiments are running, recording of initial and final states of all equipment, and logging of all experiments. This software is especially useful in scenarios requiring remote control of instruments located in hostile environments, such as the interior of the dilution refrigerator in our cryogenics lab.

Platypus has many advantages over more common experiment automation programs, such as LabVIEW and other GUI-based software. Since the Platypus user interface is a Python command line interpreter, it allows for the creation of highly complex experiments and complete user control over code. It also enables easy adaptation of existing experiments for new purposes and for use with new pieces of equipment. Being custom-built, it is license-free and can be installed in all of our labs without expense.

While Platypus' core code is generally altered only by Dr. Bogan, the average user will typically interact with the following three components:

- 1) **Device Drivers:** In order to automate an experiment, control software must be created for each piece of equipment.
- 2) **Experiment Automation Files:** A script specifies the actions of each instrument during the experiment and the order in which they should occur.
- 3) **User Interface:** The Platypus user interface is composed of a window for monitoring instrument readings and a Python shell where users can interact with instruments and experiments.

Once an experiment file has been written, the experiment may be run as many times as desired, and it is usually designed to be general enough that it will work with several different pieces of equipment. Platypus thereby provides versatility and efficiency, allowing the experimenter to focus on other tasks while complex and time-consuming experiments are being run by the computer. In the following sections, we will briefly review the three main components of Platypus.

B.1.1 Devices

The first step in automating an experiment with Platypus is the creation of a device driver file for each piece of equipment. With this code, the user is able to completely control the device from the Platypus user interface without adjusting any settings manually. Platypus divides equipment into two categories: ‘control’ devices and ‘measurement’ devices. Measurement devices are those which are used to record data by making measurements. Control devices are those which control the conditions of an experiment. For example, when measuring the current flowing through a wire, the control device would be the voltage source, and the measurement device would be the ammeter reading the amplitude of the current.

Device files usually consist of a single Python class, with each method of the class controlling a different function of the instrument in question. There are several standard operations which nearly every device file contains, while some instruments have unique abilities and require much more complex code. Measurement devices typically have a simple **read** method which, when called, causes them to acquire data. Control devices will have code unique to their function. For example, voltage sources will have code to change the output voltage and mechanical rotators will have code to rotate them to a specific locations. If a control device has a particular function, the Platypus command which causes it to execute that function is usually called **set** (e.g. set the voltage to a specific value, set the rotator to a certain location, etc.). All commands must be sent in the particular format required by each device, such as truncated words or hexadecimal strings. The user manual or programming manual for a given device will usually provide the necessary information for controlling it remotely.

Some devices require the ability to perform more complicated tasks than their basic functions. As we saw in Chapter 3, for example, our lab uses three piezoelectric nanopositioning stages, which normally only accept commands to move a certain number of steps

backwards or forwards. They possess a separate system for determining their actual position, however, so more complicated device files were required in order to incorporate other equipment and more advanced positioning abilities.

The following is a list with a few examples of instruments for which I have written or updated Platypus device drivers:

- Thorlabs PM100D powermeter
- Kiethley 2230-30-1 power supply
- Attocube rotator/linear positioners and ANC300/ANC350 controllers
- Agilis AG-UC8 piezo motor controller
- Thorlabs KDC101 rotator motor controller
- Kiethley 2010 multimeter
- Kiethley 2400 source meter
- Tektronix TDS 6154C oscilloscope
- Newport NSC200 rotator controller
- PathWay 7334 electronic switcher
- Tektronix TDS 5052 oscilloscope
- Tektronix AWG70002A arbitrary waveform generator

All device files comply with the requirements of Platypus as coded by Dr. Bogan and are therefore based upon his outline for writing them.

B.1.2 Experiments

To automate an experiment, the user must write a script which: a) dictates the order in which devices are to perform the required actions and b) determines how the data will be saved and displayed on screen. As with the device files, each Platypus experiment file has a fairly standard format. Experiments are Python classes with three specific methods: **start**, **iterate**, and **finish**.

The **start** function is activated when the experiment is initiated from the Platypus user interface. It contains everything necessary to begin the experiment, including initial states of instruments and code to create a graph window so that incoming data may be viewed while the experiment runs.

The **iterate** portion of the experiment begins after **start** has finished. This segment contains the main body of the experiment, and Platypus repeats it (hence the name ‘iterate’) until a specific condition has been fulfilled, the nature of which is specified by the user. A simple example would be an experiment in which a multimeter makes a voltage measurement once per second for 10 minutes. In such a case, **iterate** would contain just the multimeter **read** command, followed by a 1-second delay. Platypus would then repeat these two lines 600 times, i.e. repeatedly for 10 minutes. The **iterate** method also contains the code to update the plot window as measurements are made, as well as instructions for saving newly-acquired data to a text file with the extension *.dat* once per iteration. The latter procedure ensures that no information is lost if the experiment is stopped prematurely and allows preliminary analysis of the results before the experiment has even finished running.

The last part of the experiment file, called **finish**, is activated after the final repetition of **iterate**. This section contains all end-of-experiment instructions to the equipment, such as having the devices put themselves into safety mode. Images of data plots are typically saved at this point as well.

I have written or edited dozens of different Platypus experiments during my time at the SPIN Lab, and all of the research presented in this thesis was conducted using Platypus. Examples of Platypus experiment files to which I have contributed include:

- Many experiments in which Attocube positioners move a fibre while optical measurements are made. An example would be imaging nanostructure samples with sub-micron resolution using photocurrents, reflectance, and photoluminescence emissions
- Calibration programs for the Attocube positioners
- Experiments which acquire spectral data via the Princeton Instruments Lightfield software
- Experiments to rotate waveplates and observe the intensity of polarized light which has passed through them
- Experiments for monitoring instrument readings over time
- Experiments for monitoring instrument readings as the control settings are varied

All experiment classes inherit from the **platypus.base.StandardExperiment** class created by Dr. Bogan and are based on his outline for writing them.

B.1.3 User Interface

The Platypus user interface consists of two components. The first is the Python shell window, which is shown in Figure [B.1](#). The user may run experiments or send commands to

instruments from the shell, which automatically loads device drivers and opens communication channels to all connected equipment on startup. A second, smaller window, called the State Monitor, displays the current readings of all measurement devices, allowing the user to monitor the status of the instruments during an experiment.¹ Figure B.2 shows the State Monitor window with readings for four instruments.




Figure B.1: The Platypus user interface shell.

Running an experiment is a simple matter of importing the file in the Platypus shell, creating an instance of the experiment class with the appropriate arguments, then starting the experiment with a call to its **start** method. Before running the **iterate** method, the values in the State Monitor are automatically updated and recorded in a text file with the extension *.start*. After running through the appropriate number of iterations, the **finish** method is executed. Platypus then updates the State Monitor once more and saves the final states of all the devices in a text file with the extension *.end*. The *.start* and *.end* files may act as a reference for future experiments, and all commands typed into the Platypus shell are also saved to a log file. This automatic logging of all commands and experiment conditions has proven to be a useful resource, especially when we attempt to replicate or recall the details of a given experiment.

One of Platypus' most useful features is the **STOP** button in the State Monitor window (see Fig. B.2). This button allows one to pause any experiment which is in progress. Once an experiment has been stopped, the user may alter its parameters before typing the **resume** command into the shell. Examples of parameters which one might wish to alter could include the time delay between measurements or the step size of an incrementally-increasing voltage. Once resumed, the experiment continues on from that point with the new parameters. Platypus therefore allows for in-situ modifications to be made to experiments, rather than cancelling them or waiting for them to be finished.

In addition to automating individual experiments, Platypus can also run scripts, wherein multiple experiments may be called and executed consecutively. This option is enormously

¹The State Monitor does not have real-time updating of instrument readings, as this would slow down communication with other devices. Instead, it merely displays the result of the most recent measurement of each device, which may have occurred as part of an experiment or may have been triggered manually by the user from the shell.

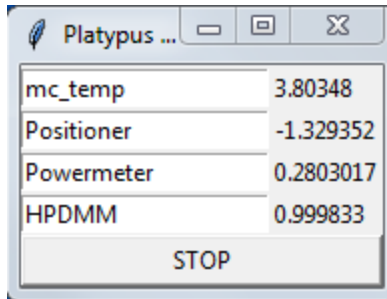


Figure B.2: The Platypus State Monitor window displaying the status of several devices. The **STOP** button allows one to pause an experiment.

useful, as one can run many experiments overnight or over a weekend without the need for user input. An obvious example would be the case where a single experiment is to be repeated several times, with one variable being changed each time. With a simple loop script, the user does not need to be present in the lab to start each new experiment, which can save a lot of time.

B.2 Significance of Platypus

While many similar device/experiment automation programs exist, Platypus is a highly-adaptable, modular² program which provides the user with complete control over the actions of their equipment. By having the user work directly in Python, it allows for far more complex experiments than are practical when using graphics-based software such as LabVIEW, where programs are written using pictures and lines. Other code-based experiment automation programs do exist, but they are generally hindered by their attempt at a ‘one size fits all’ solution. The open-ended programming options available with Platypus, however, allow one to easily overcome the challenges of automating a wide variety of highly complex and diverse experiments.

With Platypus, once experiment and device files have been written, running experiments becomes a simple matter of typing a few lines into the user interface. Time-consuming experiments can be set to run overnight or over a weekend, and multiple experiments can be set to run consecutively without need for further input. Equipment located in difficult-to-reach locations or in dangerous environments (e.g. in vacuum or ultra-low temperatures) can be easily controlled using the device driver software. Since the user is the one who creates the device driver software, instruments can be made to perform new and unique tasks which are needed for specific experiments, rather than being limited to their standard functions.

²It is modular in the sense that it is composed of many different pieces, i.e. the device and experiment files, which can be added or removed as needed.

With its unique features and versatility, Platypus offers many advantages over more common experiment automation tools. As we have expanded the range of Platypus' abilities, this software has significantly improved the efficiency of our work, allowing us to focus on other tasks while the computer runs the experiments for us.

Appendix C

Interactive Data Plotting and Analysis Software

As discussed in Chapter 2, I have created a lot of software for the SPIN Lab during the course of my degree, with one such program being a data analysis program. With automatic plotting of experimental data on start-up, this software is equipped with customizable tools for comparing, analyzing and performing mathematical operations on data. As it has become an invaluable and time-saving component of our daily lab work, I will use this appendix to describe some of its most useful features.

Many thanks to Dr. Louis Gaudreau of NRC for providing the initial idea for the program, and for suggesting important features which were later incorporated into it.

C.1 Introduction to the Interactive Plotter

Prior to the creation of the Interactive Plotting and Analysis Program, every data set produced by an experiment, be it a multi-hour overnight experiment or a brief two-minute test, would require the data to be manually retrieved, plotted and analysed. In our labs, this procedure is typically done using the Python programming language, which means that we were required to write new code every time an experiment concluded. Due to the time required for plotting and analysing each new data set, experiments could rarely be conducted in quick succession, as the parameters for each experiment typically depend on the results of the previous one. To improve the efficiency of our experimental work, I was tasked with creating a new plotting system.

C.1.1 Layout and Basic Features

The Interactive Data Plotting and Analysis Program (Interactive Plotter for short) is a Python-based graphical user interface intended to be used for quick viewing of experimental

data and preliminary data analysis. Building on the basic plotting abilities of the Python library Matplotlib, this program eliminates the need to write new code every time data is to be plotted, thanks to automatic graphing of the most-recently-acquired data. It provides users with fast access to previous data sets, the ability to perform mathematical operations on data, and multiple methods of data comparison and analysis. Users are able to interact with the plots through the use of unique functions which have been added to the basic Matplotlib toolbar, and processed data may be saved for future use. This program is not intended to entirely replace detailed plotting and analysis for publication purposes, but may be used to quickly observe, compare and analyze data sets on a day-to-day basis.

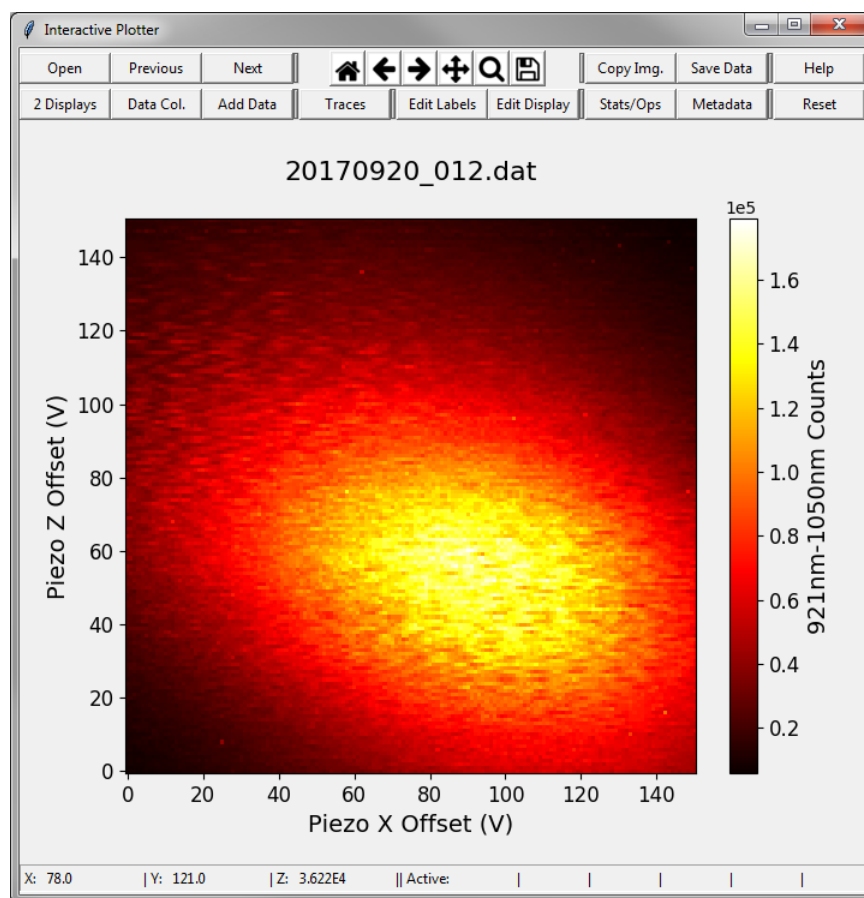


Figure C.1: The Interactive Plotting and Analysis Program


Upon start-up, the Plotter automatically retrieves and plots the most-recently-acquired data set, with no need for any user input. The program appears as a simple graphical user interface, with the plot filling most of the screen. Two rows of buttons at the top give access to analysis tools, and a status bar at the bottom indicates the x-, y-, and z-coordinates of the computer cursor as it hovers over the plot. Figure C.1 shows the basic Plotter window, which is displaying data from a high-resolution scan of nanowire

photoluminescence emissions. The program displays line plots for 2D (XY) data sets, and heat maps for 3D (XYZ) data. By default, the plot title is the data file name, and the axis labels are the labels of the data columns in the file.

The *Open* button in the top row allows the user to select another data set from a standard *Open File* dialog, while *Previous* and *Next* buttons will automatically plot past data sets in chronological order. In the centre of the top row of buttons, the standard Matplotlib navigation toolbar provides options for manually zooming and panning, as well as saving a picture of the plot. The *Copy Image* button allows the user to copy an image of the plot to the clipboard.

C.1.2 Altering the Plot's Appearance

While this program is not necessarily intended for use in creating publication-level graphs, its current features can generate plots which are of sufficient quality for presentations and talks. For aesthetic purposes, the user has complete control over the appearance of the plot. Options are provided for changing the title and axis labels, and for adding text to the plot. The axes can be switched between linear and logarithmic, and the axis tick label size can be adjusted. For 2D files, data can be plotted as line plots or scatter plots, with options for changing the line/marker style, colour, and size. Figure C.2 shows examples of display options for 3D plots, which include various colour schemes, several interpolation methods, and the ability to zoom in/out along the Z-axis. The user can also choose to display 3D data as either a heat map or a contour plot. Additionally, the x- and y-axes can be swapped by reassigning the columns in the data file.



Different colour schemes

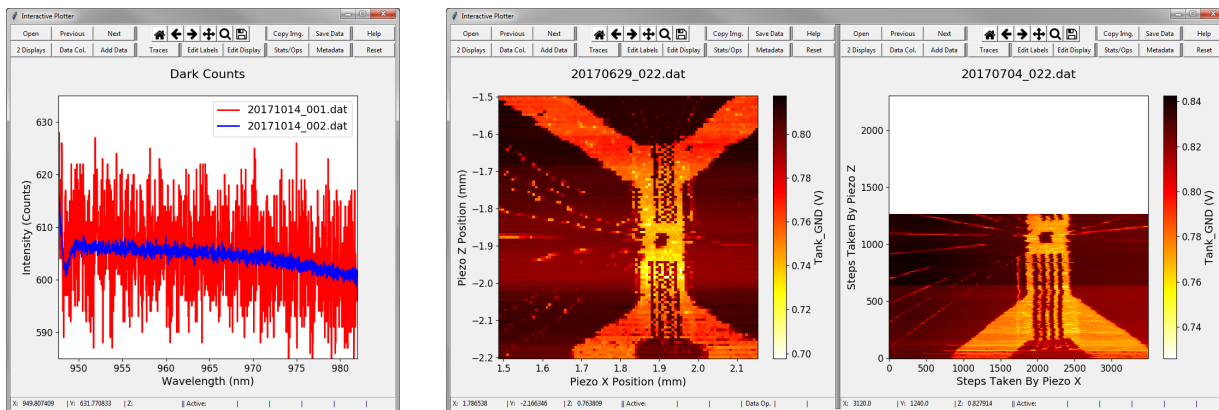
Contour plots

Z-axis zooming

Figure C.2: Options for displaying a heat map.

C.1.3 Comparing and Analysing Data

The Plotting Program is often used for preliminary analysis of experimental results due to its variety of options for comparing and analysing data files. Figure C.3 shows how multiple data sets can be superimposed on a single plot and how two separate graphs can be displayed beside each other for comparison. The program can also save multiple data sets to a single file, which, when opened in the Plotter, will automatically display all of them on one graph. These features of the program have are generally used to quickly compare experimental data sets in order to determine the effects of changing experiment parameters.



Multiple data sets per plot

Two displays

Figure C.3: Options for comparing data sets.

One of the major features of the program is accessed through the *Statistics and Operations* button in the lower toolbar. In addition to providing various useful statistics for the graph(s) displayed, the *Statistics and Operations* pop-up window allows the user to perform mathematical operations on the data. Examples of such operations may include normalization, differentiation, and removal of background noise. The program contains several built-in mathematical tools and allows the user to create their own custom operations. Custom operations which are used often may be saved to an external file, where they are retrieved by the Plotter, while simple operations may be written in the program's Function Editor for one-time use. For every operation performed, the plot is instantly updated. Mathematical operations can even be performed on experimental data while the experiment is still running, allowing us to terminate early those experiments which do not appear to be yielding useful results, rather than waiting for them to finish before analysing the data.

The ability to apply multiple mathematical operations to the data with the press of a few buttons has virtually eliminated the need for us to write new Python code each time we need to perform calculations on data. Note that the program does not actually change

the data file itself; the altered data exists only in the program's memory. The software does provide the option to save altered data to a new file, however, which can then be accessed by the Plotter in the future. The *Reset* button resets the plot to its initial state by reading the original data file again.

C.1.4 Trace Viewing and Spectrum Viewing

A useful tool in analysing a 3D plot is the ability to extract a single line of data, such as a horizontal or vertical trace. The Plotting Program therefore provides the *Trace Viewing* tool for viewing of the individual line traces which comprise a 3D data set, which enables one to observe those features of the data which are not easily seen in the colours of a heat map. Figure C.4 shows a high-resolution map of nanowire emissions, where horizontal and vertical cross sections of the data are displayed next to the graph. The cursor can be used to select the traces to be displayed, or the keyboard arrow keys can be used to skim through the rows and columns of the 3D plot, with the traces being updated automatically. The program also supports custom diagonal traces, where the user clicks and drags the cursor across the plot. Traces can be expanded into a regular-sized plot which is displayed beside the 3D data set. (Expanded traces are also updated as the user clicks on the 3D plot or navigates with the arrow keys.) Multiple traces can be added to the right-hand plot simultaneously for easy comparison.

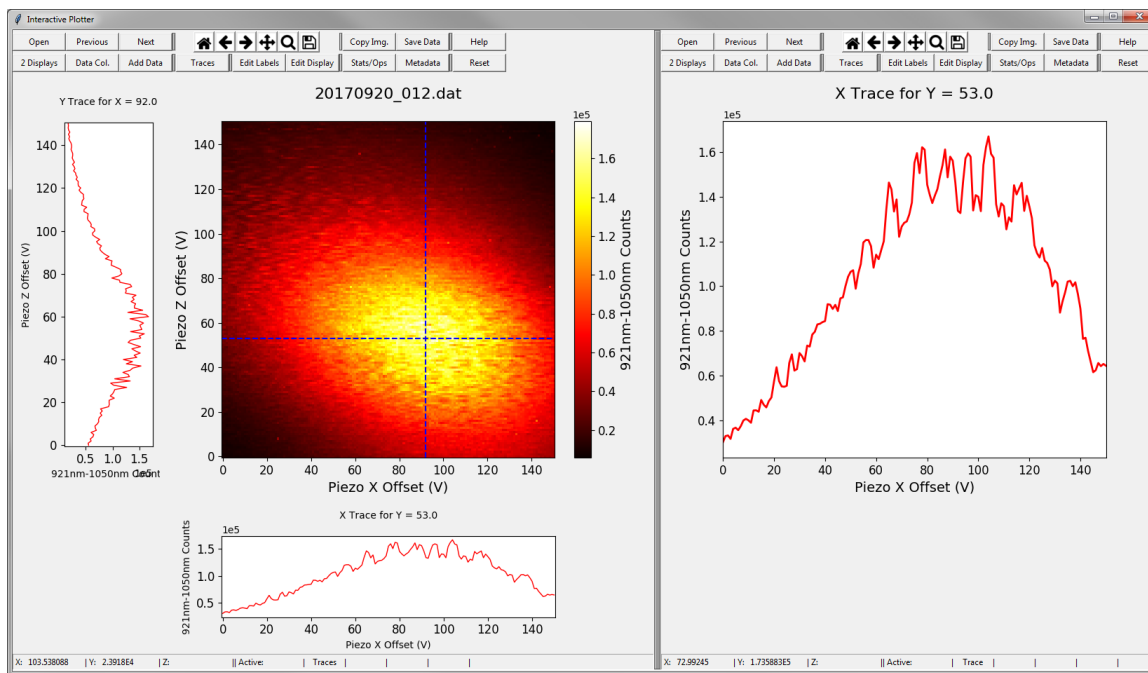


Figure C.4: *Trace Viewing* allows one to observe the individual horizontal and vertical traces which form a 3D plot. The dashed blue lines indicate the specific traces seen below and to the left of the heat map. On the right, the horizontal trace has been enlarged.

In special cases, each data point of a 3D plot may have an entire additional data set associated with it. In the example of Figure C.5, each data point of the 3D plot indicates the integrated intensity of light measured within a specific wavelength range (921nm-1050nm). The value of each data point is therefore obtained from a unique emission spectrum. The Interactive Plotter is able to display these spectrum files beside the 3D graph, as in the right-hand plot of Figure C.5, which shows the spectrum associated with the point indicated by the dashed blue lines on the 3D plot. As with *Trace Viewing*, this feature (which we call *Spectrum Viewing*) allows one to use the arrow keys to move around the 3D plot, with automatic updating of the spectrum file. Spectra from other data points can be superimposed on the right-hand plot for comparison purposes.

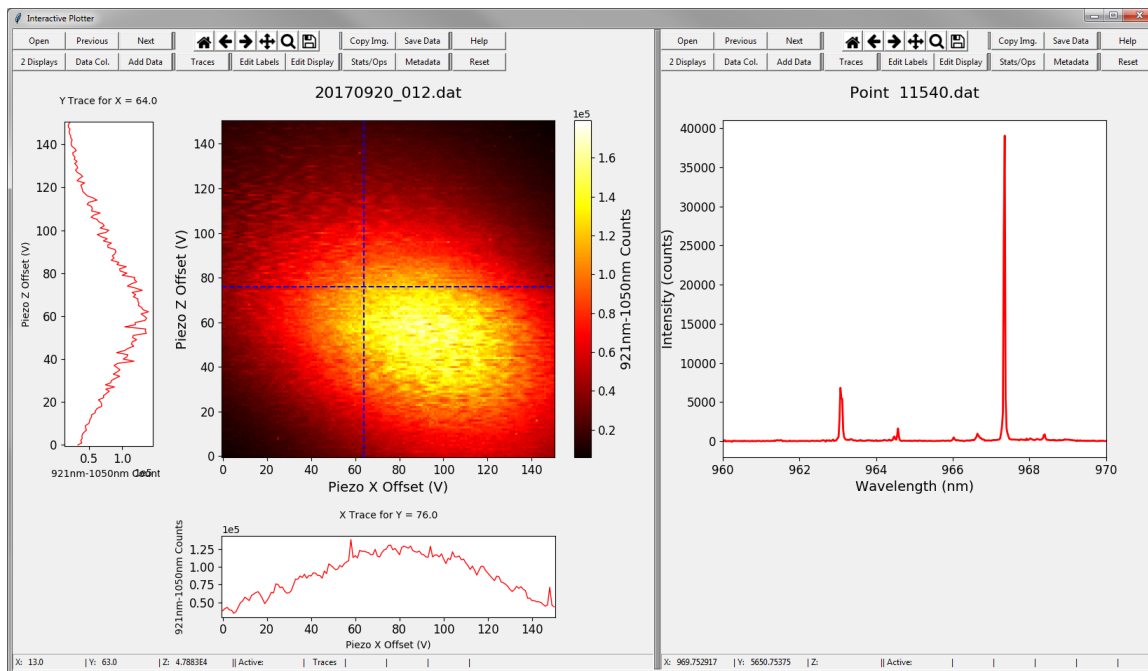


Figure C.5: *Spectrum Viewing* lets one observe a second plot associated with a specific point on a 3D plot. In this example, the right-hand plot is the measured emission spectrum of a nanostructure. By integrating the number of counts, we get the value of one data point on the 3D plot, specifically Point 11540.

The high-resolution emission profile in Fig. C.5 is comprised of the integrated counts from almost 23 thousand different spectrum files. Manually locating and plotting just a few of these files for preliminary analysis would take much longer without the use of the Plotting program.

C.2 Significance of the Interactive Plotter

While this software is by no means unique, having features of many other data plotting programs (like Origin, for example), the fact that it has been specifically tailored to our needs has made it an important part of our daily experimental work. Existing plotting software typically requires a fair amount of user input, since it is employed for a wide variety of applications, whereas many of the Plotting Program's abilities have been highly automated. While other programs require the user to specify file names and plot formats, this work is all done automatically by the Plotter, and any mathematical operations or display alterations can easily be accomplished with the press of a few buttons. We have also created a library of custom-made, often-used mathematical functions which the program can retrieve and apply to our data.

The Plotting Program is designed to operate in conjunction with Platypus, the experiment automation software used in our labs (see Appendix B). While Platypus automatically plots data as an experiment runs, it does not have data comparison or analysis features. The Plotting Program works with Platypus' file naming convention in order to load experimental data and uses the metadata output of an experiment to display additional information, such as duration of the experiment. Data from sources other than Platypus can be easily plotted as well.

This software is being increasingly used in the labs run by our group at NRC, as it provides an extremely efficient manner of reviewing and analysing data during experiments. The principle benefits of this software are its automated data plotting, its versatility in displaying different types of data (and even incomplete data sets), and its analysis features, which have been tailored to the type of experiments we do in our labs. The use of the Plotting Program has already saved us a significant amount of time during experiments, and it will no doubt continue to do so in the future as new features are added and existing ones are refined.

Appendix D

The SPIN Lab’s Piezoelectric Nanopositioning System

In Chapter 3, we introduced the SPIN Lab’s sample imaging system, as well as the devices, procedures, and software which comprise it. The purpose of this appendix is to supplement that chapter and provide further details on the nanopositioners. While not crucial to the main ideas presented in the text, the information presented here may assist the reader in understanding some of our decisions during the development of the imaging system.

D.1 Temperature-Dependent Positioner Behaviour

Piezoelectricity relates electric field strength and charge density to the strain and stress experienced by a material. As the temperature decreases, so too do the piezoelectric coefficients which relate these properties [63]. The result is a diminished voltage response to the application of stress, and a diminished deformation response to an applied voltage.

Several temperature-dependent behaviours of the piezoelectric positioners become relevant when we image our samples in the fridge. The most immediately-noticeable effect of a temperature change is the corresponding difference in positioner step size. For a given step voltage, the size of a step is much larger at room temperature than at mK temperatures. This temperature dependence is evident when one considers that the maximum range of the positioners in DC mode is $5\ \mu\text{m}$ at room temperature, but just $800\ \text{nm}$ at $4\ \text{K}$ [64, 65]. This behaviour has the convenient effect of allowing us to perform more detailed scans at low temperatures without having to use smaller step voltages.¹

The positioners moving the fibre tend to get stuck more frequently while performing their slip-stick motion at room temperature (although the smaller step sizes at low temperatures would seem to suggest the reverse should be true). An increase in the step voltage

¹If the step voltage is too low, the positioner has difficulty moving and may become stuck.

(35–40 V) and/or the step frequency (1–2 kHz) will typically suffice to get the positioner moving smoothly again. With the fridge closed at low temperatures, it can be difficult to determine whether a positioner has become stuck. As the step size during a high-resolution scan is often below the noise level of the resistive position encoder, the only way to tell whether the positioner is stuck is to instruct it to move a large distance ($\gtrsim 800$ nm). While the resistive position encoder will be able to detect this movement if it is successful, this risks moving the positioner far from the desired position. Rather, if it is suspected that the positioner is stuck, the first course of action should be to increase the step voltage. Taking several steps, rather than a single step, may also help get the positioner moving, as the repeated force seems to help it overcome its immobility.

D.2 Positioner Calibration Curves

Figure D.1 shows the calibration curves we obtained for the three types of piezoelectric positioners used in our system (reproduced from Chapter 3). In producing these curves, we assumed a constant step size as the positioner moved from one end of its range to the other. After the data was acquired, a scale of -2.5 mm to 2.5 mm was applied to the x-axis of the calibration curve (or 0° - 360° for the rotator), since 5 mm is known to be the full range of travel. In spite of an applied voltage of 1 V, the calibration curves for the linear positioners do not extend from $S = 0$ V to 1 V because the resistive elements stretch slightly beyond the 5 mm travel range.² The value of the applied voltage V_0 is itself largely irrelevant, as the measured voltage S can be plotted as a fraction of V_0 . Consequently, these curves are still applicable for any value of V_0 . Higher values of V_0 provide better positioning resolution due to a relative reduction in voltage noise, but $V_0 = 1$ V was chosen as a compromise between accuracy of positioning and the ohmic heating of the positioners at low temperatures.




Figure D.1: Calibration curves for the a) ANPx101/RES horizontal positioner, b) ANPz101/RES vertical positioner, and c) ANR240/RES rotator.

The calibration plots are not precisely linear, in part due to imperfections in the resistive strip, but mainly because the positioner step size is not precisely constant, even when an

²There is also some resistance in the wiring between the voltage source and the positioners.

identical sawtooth waveform is applied. The step size repeatability is just 5%, with higher variability likely to occur when the device starts moving [63]. We suspect the variation in step size may also depend on the location of the positioning stage, as we consistently observe across multiple calibration attempts that some sections of the curves have steeper slopes, while other sections are shallower. The non-linearity of our calibration curves has some distorting effects on our large-scale imaging scans, especially in the case of the rotator, but the microscopic imaging presented in the following sections was not significantly affected by this issue. Linear calibration curves, produced by drawing straight lines between the minimum and maximum measured voltages, produced marginally better results for large images, as they ignore the step size. (Imperfections in the resistor are not taken into account with the straight line method.)

D.3 Hysteresis Effects

One of the challenges we experienced during the development of our imaging procedure was hysteresis in the movement of the positioners. In this section, we will briefly look at two of these hysteresis effects.

D.3.1 Extreme Hysteresis in the Vertical Positioner's Step Size

The vertical positioner has a particularly large step size hysteresis. For a given voltage, it may take dozens of steps upwards to move the positioner the same distance as a single downward step. Whereas we commonly use step voltages of about 24-28 V for the horizontal and focusing positioners, we require a step voltage of 34 V for the vertical positioner in order to achieve any substantial upwards movement. The cause of this hysteresis, which is not observed to such an extreme degree when the device is operated in isolation, is likely the weight of the fibre clamp and the two other positioners, which are held up by the vertical positioner. The locations of the positioners and the fibre clamp are shown Figure 3.3. The vertical positioner must pull the full weight of those three components upwards, while their weight assists with its downwards motion. Therefore, the vertical positioner is always set to move downwards when imaging a sample, while the horizontal positioner moves back and forth, acquiring data. The vertical positioner is then moved up again before the next scan begins. A (perhaps obvious) resolution to this issue would be to move the vertical positioner to the bottom of the stack, so that it need only hold the weight of the fibre clamp. We have yet to attempt this.

D.3.2 Positioner Voltage Probe Hysteresis

We have recently learned that the small probe S used for readout of the positioner location is slightly flexible and drags along the surface of the resistor. As the stage moves, the

probe bends and lags behind. In addition to the step size hysteresis, there is therefore a second hysteresis effect in the measured position due to the probe lagging behind the stage. Compensation for this effect can be achieved by always travelling in the same direction when making position measurements and overshooting the target by $100\ \mu\text{m}$ when returning in the opposite direction. This overshoot allows the probe time to switch directions when the stage moves back to the desired start position. This procedure is still subject to the $1\text{-}2\ \mu\text{m}$ position repeatability which is fundamental to the devices, but it should at least improve the repeatability of our measurements to that point. As we were not aware of this positioner property at the time, we did not use this procedure for the scanning experiments described in this thesis.

Appendix E

Details of the HBT Plot

To supplement the information given in Chapters 1 and 4, here I discuss two features of the HBT graphs which warrant further explanation: the ‘negative’ time delay, and the symmetry of the peaks in the pulsed HBT plot.

E.1 ‘Negative’ Time Delay in an HBT Plot

It might be noted that in the auto-correlation plots, half of the data appears at a ‘negative’ time delay.¹ This is achieved with the use of a delay line attached to one of the single-photon detectors. When a photon is detected, an RF line sends a signal from the detector to the computer, and the computer measures the time delay between detections at D1 and D2. Suppose that an extra length of RF line corresponding to a 25 ns delay is now added to the D2 detector.² If a photon arrives at each detector simultaneously, the D2 detection will be measured by the computer 25 ns after the D1 detection. This means that the ‘25 ns delay’ on the horizontal axis of a $g^{(2)}(\tau)$ plot is effectively the new ‘zero time delay’, so if there is single-photon emission, one should now expect zero counts to occur at a delay time of 25 ns instead. Similar results could be achieved by making the distance between the BS and D2 larger than the distance between D1 and the BS.

In our auto-correlation plots, we shift the values on the horizontal axis so that 25 ns appears as a delay of 0 ns, to make for easier viewing of the graph. This shift merely accounts for the delay line and moves ‘zero delay’ to the real position where simultaneous detections would occur. To clarify this description, suppose that a photon is detected at D2 less than 25 ns before a detection at D1. The D1 signal will reach the computer first, followed less than 25 ns later by the D2 signal. Without the delay line, this event would

¹The HBT experiment measures the delay between a detection at D1, followed by a detection at D2. A ‘negative’ time delay appears to imply a reversal of this order: D2 followed by D1.

²We actually used a few different lengths of delay lines over the course of our experiments, including 25 ns.

not be recorded (since the D2 detection is occurring before a D1 detection, but with the delay, this becomes one of the ‘negative’ time delay values on the new, shifted axis. Any detections which occur at D2 after a detection at D1 will occur at their normal, positive values on the new horizontal axis. For continuous emission, the plot should be symmetric about the new ‘zero delay’ position,³ but without the delay line, only positive delays would be visible.

E.2 Peak Symmetry in a Pulsed HBT Plot

As the exciton lifetime measurement produced by a laser pulse is very clearly asymmetric (Fig. 4.17a), we might also expect the peaks observed in the pulsed HBT data (Fig. 4.16a) to be asymmetric as well, since they represent the decay of excitons produced by subsequent pulses. Figure E.1 attempts to explain the symmetry we instead observe in the plots.




Figure E.1: Explanation for symmetry of $g^{(2)}(\tau)$ peaks in spite of asymmetry of exciton lifetime function. (See text for explanation.)

In the pulsed HBT experiment, the counts around Point A are detected when the D1 photon is produced on the tail-end of the exciton lifetime, near Points 2 and 3, while the D2 photon is produced much earlier with respect to its pulse, at Points 1 or 2. Since the D1 photon arises late and the D2 photon is produced early, coincidence counts appear earlier. Point B of course corresponds to both photons being produced at the very beginning of the pulse, Point 1, when exciton recombination is most probable. Finally, the data at Point C comes from instances where the D1 photon is produced at Points 1, 2, or 3, while the D2 photon is produced near the edge of the lifetime, at Points 2, or 3. There is therefore a longer delay between the recombinations of their respective excitons. Table E.1 summarises the emission times for the photons detected at D1 and D2 for each point in the HBT peak.

³This symmetry is mathematically true of the $g^{(2)}(\tau)$ function.

HBT Point	Emission Point for D1 Photon	Emission Point for D2 Photon
A	2,3	1,2
B	1	1
C	1,2,3	2,3

Table E.1: Emission times for the photons detected at D1 and D2 for each point in an HBT peak. Emission points are given according to the labels in Fig. 4.17c).

Appendix F

Uniquely Identifying Polarisation States

As discussed in Chapter 5, our optical system uses a rotating QWP, a fixed linear polariser (horizontally-oriented), and a photodiode to determine the polarisation state of a beam of light. In this Appendix, we will work through the simple proof that the intensity signature produced by measuring the light intensity as a function of QWP angle is unique for every polarisation state. We will be using the Jones calculus described in Section 1.5.2, as well as basic quantum theory. I would also like to thank Dr. Alex Bogan of the NRC for having thought of this proof and demonstrated its validity.

We begin by noting that the unitary transformation of a density matrix $\hat{\rho}$ (or indeed, any operator) is given by:

$$\hat{\rho}' = \hat{U}\hat{\rho}\hat{U}^\dagger \quad (\text{F.1})$$

where \hat{U} is the unitary operator and \hat{U}^\dagger is its Hermitian adjoint. In the case of our optical setup, the unitary transformation is the action of the QWP on the input state $\hat{\rho}$.

Next, we consider the fact that the expectation value of an observable \hat{O} on the state $\hat{\rho}$ is:

$$\langle \hat{O} \rangle = \text{Tr}(\hat{\rho}\hat{O}) \quad (\text{F.2})$$

The probability of measuring a given state $|o_i\rangle$ when projecting onto the basis of observable \hat{O} is therefore:

$$\text{Pr}_i = \text{Tr}(\hat{\rho}|o_i\rangle\langle o_i|) \quad (\text{F.3})$$

This means that if we are projecting a state $\hat{\rho}$ onto the horizontal polarisation state, then $|o_i\rangle = (1, 0)^T$, and our observable is:

$$\hat{O} = |o_i\rangle \langle o_i| = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad (\text{F.4})$$

This is just the Jones matrix for a horizontally-oriented linear polariser, P_H , as we found in Section 1.5.2. It specifically isolates the horizontal component of the polarisation.

Putting these two concepts together, then the (normalized) probability of measuring the horizontal state after sending light through a rotated QWP, followed by a horizontal polariser is:

$$\begin{aligned} \text{Pr}_H &= \text{Tr}(\hat{\rho}' P_H) \\ &= \text{Tr}(\hat{U} \hat{\rho} \hat{U}^\dagger P_H) \end{aligned} \quad (\text{F.5})$$

Since the QWP matrix is dependent upon its angle, θ , we should in fact have $\hat{U} = \hat{U}(\theta)$. Additionally, our actual measurements will be fractions of the total intensity I_0 in front of the polariser, so we should scale our result. The final measured intensity is therefore:

$$I = I_0 \text{Tr}(\hat{U}(\theta) \hat{\rho} \hat{U}^\dagger(\theta) P_H) \quad (\text{F.6})$$

Now we can ask whether there are specific values of θ which are unique for every polarisation. We can infer from Figure 5.7 (and from simulations of all the signatures produced around the Poincaré sphere) that there is no one value of θ which will uniquely identify a polarisation, so we must instead resort to combinations of angles. In fact, it turns out that we need a minimum of three QWP angles to identify the polarisation. (Therefore, we technically only need three points instead of the entire curves shown in Figure 5.7.)

To prove this, we deconstruct the observable of the QWP-polariser system and show that it can be decomposed into a different equation for each value of θ , with the three equations being linearly independent. By the cyclic property of the trace, we can write this observable as: $\hat{U}^\dagger(\theta) P_H \hat{U}(\theta)$. Let's label it $\hat{\alpha}$.

For some polarisations (e.g. any linear states), the signature measured by this system repeats after every 90° rotation of the QWP, so let's choose smaller angles, say 0° , 22.5° , and 45° . (As initially shown by Dr. Bogan, this choice of angles makes for a nice example.) For $\theta = 0^\circ$, our observable becomes:

$$\begin{aligned} \hat{\alpha}(0) &= \hat{U}^\dagger(0) P_H \hat{U}(0) \\ &= \begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \end{aligned} \quad (\text{F.7})$$

We can obviously decompose this matrix using the identity, \mathbb{I} , and Pauli matrices, σ_i . In particular, $\hat{\alpha}(0) = 1/2(\mathbb{I} + \sigma_z)$. To help us determine the decomposition of $\hat{\alpha}$ for each angle, we can use the well-known equation for decomposing a 2×2 matrix M into the Pauli matrices:

$$M = \frac{1}{2} \left[\text{Tr}(M)\mathbb{I} + \sum_i \text{Tr}(\sigma_i M)\sigma_i \right] \quad (\text{F.8})$$

Using this equation, we find the following:

$$\begin{aligned} \hat{\alpha}(0) &= \frac{1}{2}(\mathbb{I} + \sigma_z) \\ \hat{\alpha}(22.5) &= \frac{1}{2} \left(\mathbb{I} + \frac{1}{2}\sigma_x + \frac{1}{\sqrt{2}}\sigma_y + \frac{1}{2}\sigma_z \right) \\ \hat{\alpha}(45) &= \frac{1}{2}(\mathbb{I} + \sigma_y) \end{aligned} \quad (\text{F.9})$$

It is quite evident that these three equations are at least partially linearly independent, which means that they form a basis in the measurement space. To determine the degree of independence, we can use the triple product of these three vectors, where 1 indicates complete independence and therefore perfect distinction between polarisation states. Discarding the common factor of 0.5 and the identity from the results in Eq. F.9, we find that the triple product for $\hat{\alpha}(0)$, $\hat{\alpha}(22.5)$, and $\hat{\alpha}(45)$ is 0.5. This indicates that these three points are perhaps not the best points to use for identifying polarisation states.

In fact, QWP angles of 0° , 30° , and 60° yield a much better result:

$$\begin{aligned} \hat{\alpha}(0) &= \frac{1}{2}(\mathbb{I} + \sigma_z) \\ \hat{\alpha}(30) &= \frac{1}{2} \left(\mathbb{I} + \sin(60) \cos(60)\sigma_x + \sin(60)\sigma_y + \frac{1}{4}\sigma_z \right) \\ \hat{\alpha}(60) &= \frac{1}{2} \left(\mathbb{I} - \sin(60) \cos(60)\sigma_x + \sin(60)\sigma_y + \frac{1}{4}\sigma_z \right) \end{aligned} \quad (\text{F.10})$$

These vectors yield a triple product of 0.75, indicating that they are more independent than those of Equation F.9 (although it is harder to see by looking at the equations). They do not form a completely linearly-independent basis in the measurement space, however, so there may be some states which are difficult to distinguish. By taking further points, however, or a full trace for a 180-degree rotation of the QWP, we can achieve far greater accuracy. We have therefore proven that by measuring the intensity at the three QWP

angles of 0° , 30° , and 60° , we can uniquely determine the polarisation state of light with a rotating QWP in front of a linear polariser.

As a final note, many have suggested that just a rotating linear polariser could be used to identify the states, with no QWP. While this would work for most polarisations, we would not be able to distinguish between the left and right circular states, which is important for our photon-to-spin work (see Section [1.7](#)).

Thank you again to Dr. Alex Bogan for devising this proof and demonstrating its validity.