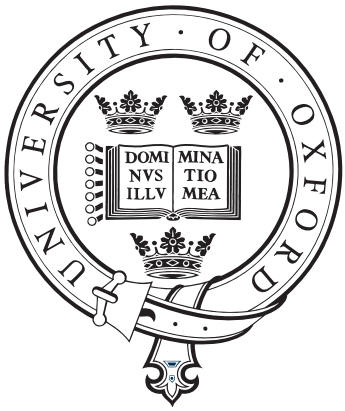
Silica-on-silicon waveguide circuits and superconducting detectors for integrated quantum information processing

Benjamin J. Metcalf
St. Peters College, Oxford



Submitted for the degree of Doctor of Philosophy
Trinity term 2014

Supervised by
Prof. Ian A. Walmsley

Clarendon Laboratory
University of Oxford
United Kingdom

Silica-on-silicon waveguide circuits and superconducting detectors for integrated photonic quantum information processing

Benjamin J. Metcalf
St. Peters College, Oxford

Submitted for the degree of Doctor of Philosophy
Trinity Term 2014

Abstract

Building complex quantum systems has the potential to reveal phenomena that cannot be studied using classical simulation. Photonics has proven to be an effective test-bed for the investigation of such quantum-enhanced technologies, however, the proliferation of bulk optical components is unlikely to be a scalable route towards building more complex devices. Instead, the miniaturisation, inherent phase stability and trivial alignment afforded by *integrated* photonic systems has been shown to be a promising alternative.

In the first half of this thesis, we describe experiments exploiting the quantum interference of three single photons on a reconfigurable integrated photonic chip. We develop a low-loss source of single photons and introduce a low-loss silica-on-silicon waveguide architecture which enables us to show the first genuine quantum interference of three single photons on an integrated platform. A loss-tolerant, element-wise characterisation scheme is developed along with a statistical test to verify that this multi-photon circuit behaves as expected. We then make use of this three-photon interference to detail the first proof-of-principle demonstration of a new intermediate model of quantum computation called boson sampling. Finally, we perform an on-chip demonstration of the quantum teleportation protocol where all key parts — entanglement preparation, Bell-state analysis and quantum state tomography — are performed on a reconfigurable photonic chip. The element-wise characterisation scheme developed earlier is shown to be critical to mitigate fabricated component errors. We develop a theoretical model to account for all sources of possible error in the circuit and find good agreement with the measured teleported state fidelities, which exceed the average teleportation fidelity possible with a classical device. We identify the elements of this error budget relevant to scaling and propose improvements to chip characterisation and fabrication in order to achieve high fidelity operation.

In the second half, we discuss the use of high efficiency superconducting transition edge sensors in enabling quantum experiments using more photons. We detail the installation and characterisation of these detectors in a new lab in Oxford. We achieve good photon number-resolution and high-efficiency operation. Work to integrate these detectors on the silica-on-silicon waveguide architecture is discussed and we detail the optical and thermal device modelling performed to optimise the on-chip detection efficiency. New, on-chip detectors, fabricated according to this design are shown to operate as expected and achieve high-efficiency and good energy resolution.

To Zoë

Acknowledgments

My first thanks must go to my supervisor, Ian Walmsley. Thank you for overseeing such an interesting and varied research project and creating an environment which was both stimulating and great fun. It is a continual source of amazement that despite your many other roles and responsibilities you have always managed to find the time for quick chat, a pint, a long discussion at the white board, or even a casual bike ride! I hope one day we will indeed get to cycle to Cambridge!

I have been extremely fortunate to have been part of a great team of people without whom I cannot imagine having enjoyed these four years as much as I did. Justin, lab buddy from the beginning, you are a true friend, a bank, provider of beer, master of brisket and defender of freedom. Nick, you were a great first mentor and taught me much, not least, the joys of the sausage sandwich during our legendary breakfast meetings. Peter, you put the 'P' in PRL. It has been a great pleasure to work alongside you although sometimes your productivity can put us all to shame - take a holiday! - oh, you already did ...

Steve, - I must admit my first impression of you was 'who is this guy with a massive beard and ridiculous ponytail??' Turns out, you were an exceptionally talented physicist. I cannot possibly list all of the things you have helped with over the years, but my special thanks must be for helping to setup the cryogenic lab. Perhaps one day I, too, will learn how to make vacuum seals without grease! Finally, I would like to thank the rest of the Ultrafast gang for being an invaluable source of knowledge, and more importantly, amusement. It has been a real pleasure working with you all over the past 4 years, thank you for making my time so enjoyable.

Outside of Oxford, I have been lucky enough to be involved in some great collaborative projects. From Southampton, my thanks go to James, Paolo and Peter for fabricating the waveguide devices. I look forward to our next sailing trip aboard the Glass Onion! I would

also like to thank the team at NIST. Especially Thomas Gerrits and Sae Woo Nam for kindly hosting me during my visit to Boulder and subsequently helping us to set up the TES detectors in Oxford. Thanks also to Adriana Lita and Brice Calkins for their work on the on-chip TES devices.

On a more personal note, I owe an enormous debt of gratitude to the Haddon family - Sarah, Dave and Laura. For all of the Sunday dinners, the barbecues, Saint's games, the holidays and being generally awesome. You have taken me in and looked after me as one of your own. I am forever grateful.

To my family, in particular my parents, thank you for everything. You have done more for me than I could possibly thank you for but I think it is safe to say that I would not be who I am or where I am without your love, support and all your hard work over the past 26 years. Harry and Sam, for being the best of siblings and for actually reading this thesis all the way through! The achievements you have both made, whilst I have been messing about with mirrors, fibres and lasers, are a continuing source of inspiration, whether it be graduating, exploring the world or getting fantastic jobs. Perhaps one day I will get one too! I am very proud of you both.

Above all, I would like to thank you, Zoë. For your support, your encouragement, your extraordinary patience and unwavering love. You are the bedrock upon which the last 10 years of my life have been built and I could not have completed this without you. But most of all, thank you for being my best friend. I dedicate this thesis to you.

Author's publications

Journal publications

1. B. J. Metcalf, J. B. Spring, P. C. Humphreys, N. Thomas-Peter, M. Barbieri, W. S. Kolthammer, X. M. Jin, N. K. Langford, D. Kundys, J. C. Gates, B. J. Smith, P. G. R. Smith, and I. A. Walmsley. "**Quantum teleportation on a photonic chip**" *Nature Photonics*, **8**, 770-774 (2014).

    - *Highlighted with News & Views by Xiao-Song Ma.*

2. P. C. Humphreys, B. J. Metcalf, J. B. Spring, M. Moore, P. S. Salter, M. J. Booth, W. S. Kolthammer, and I. A. Walmsley. "**Strain-optic active control for quantum integrated photonics**" *Optics Express*, **22**, 21719-21726 (2014).

3. P. C. Humphreys, B. J. Metcalf, J. B. Spring, M. Moore, X-M. Jin, M. Barbieri, W. S. Kolthammer, and I. A. Walmsley. "**Linear Optical Quantum Computing in a Single Spatial Mode**" *Physical Review Letters*, **111**, 150501 (2013).

    - *Highlighted with a Physics Viewpoint by Todd Pittman.*

4. B. Calkins, P. L. Mennea, A. E. Lita, B. J. Metcalf, W. S. Kolthammer, A. Lamas-Linares, J. B. Spring, P. C. Humphreys, R. P. Mirin, J. C. Gates, P. G. R. Smith, I. A. Walmsley, T. Gerrits, and S. W. Nam. "**High quantum-efficiency photon-number-resolving detector for photonic on-chip information processing**" *Optics Express*, **21**, 22657 (2013).

5. J. B. Spring, P. S. Salter, B. J. Metcalf, P. C. Humphreys, M. Moore, N. Thomas-Peter, X-M. Jin, N. K. Langford, W. S. Kolthammer, M. J. Booth, and I. A. Walmsley. "**On-chip low loss heralded source of pure single photons**" *Optics Express*, **21**, 5932 (2013).

6. J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthammer, X. M. Jin, M. Barbieri, A. Datta, N. Thomas-Peter, N. K. Langford, D. Kundys, J. C. Gates, B. J. Smith, P. G. R. Smith, and I. A. Walmsley. "**Boson Sampling on a Photonic Chip**" *Science*, **339**, 798 (2013).

    - *Highlighted with a Perspective by James Franson.*

7. B. J. Metcalf, N. Thomas-Peter, J. B. Spring, D. Kundys, M. A. Broome, P. C. Humphreys, X. M. Jin, M. Barbieri, W. S. Kolthammer, J. C. Gates, B. J. Smith, N. K. Langford, P. G. R. Smith, and I. A. Walmsley. "**Multiphoton quantum interference in a multiport integrated photonic device**" *Nature Communications*, **4**, 1356 (2013).

8. P. S. Salter, A. Jesacher, J. B. Spring, B. J. Metcalf, N. Thomas-Peter, R. D. Simmonds, N. K. Langford, I. A. Walmsley, and M. J. Booth. "**Adaptive slit beam shaping for direct laser written waveguides**" *Optics Letters*, **37**, 470 (2012).

9. T. Gerrits, N. Thomas-Peter, J. C. Gates, A. E. Lita, B. J. Metcalf, B. Calkins, N. A. Tomlin, A. .E. Fox, A. Lamas Linares, J. B. Spring, N. K. Langford, R. P. Mirin, P. G. R. Smith, I. A. Walmsley, and S. W. Nam. "**On-chip, photon-number-resolving, telecommunication-band detectors for scalable photonic information processing**" *Physical Review A*, **84**, 060301(R) (2011).

    - *Highlighted as an Editors' Choice in Science*

10. N. Thomas-Peter, N. K. Langford, A. Datta, L. Zhang, B. J. Smith, J. B. Spring, B. J. Metcalf, H. B. Coldenstrodt-Ronge, M. Hu, J. Nunn, and I. A. Walmsley. "**Integrated photonic sensing.**" *New Journal of Physics*, **13**, 055024 (2011).

Contents

Acknowledgments **iii**

Author's publications **v**

List of Figures **x**

List of Tables **xiii**

List of abbreviations **xiv**

List of symbols **xviii**

1 Introduction **1**
1.1 Quantum information 2
1.1.1 The principle of superposition 4
1.1.2 Large Hilbert spaces 6
1.2 Building an optical quantum computer 7
1.3 Thesis outline 11

2 Linear Optical Quantum Computing **13**
2.1 Quantum optics 15
2.1.1 Indistinguishable particles 16
2.1.2 Quantising the electromagnetic field 19
2.1.3 Linear optics 21
2.1.4 The Hong-Ou-Mandel effect 26
2.2 The Knill-Laflamme-Milburn (KLM) scheme 29
2.2.1 Single qubit gates 30
2.2.2 Two-qubit gates 30
2.2.3 Logic gates by teleportation 36
2.3 Beyond KLM 41
2.3.1 One-way quantum computing 41
2.4 Summary 44

3 Three Photon Quantum Interference **49**
3.1 Multi-photon states on-chip . . . 50
3.1.1 Single-photon source . . . 51
3.1.2 UV-written silica-on-silicon waveguide device . . . 58
3.1.3 Waveguide circuit design . . . 65
3.2 Loss-tolerant circuit characterisation . . . 68
3.2.1 Beam splitters . . . 70
3.2.2 Phase shifters . . . 71
3.3 Simulating interference visibility . . . 73
3.3.1 Chi-squared test for observed visibilities . . . 76
3.4 Multi-photon on-chip interference . . . 79
3.4.1 Two-photon interference . . . 79
3.4.2 Three-photon interference . . . 82
3.5 Boson Sampling . . . 85
3.5.1 Computational complexity of linear optics . . . 86
3.5.2 Results . . . 88

4 On-chip Quantum Teleportation **92**
4.1 Quantum teleportation . . . 93
4.1.1 The teleportation protocol . . . 94
4.1.2 Previous implementations . . . 99
4.2 The circuit . . . 101
4.2.1 Programmable unitary operations . . . 102
4.3 Circuit characterisation . . . 107
4.3.1 Phase shifter cross-talk . . . 109
4.3.2 Unbalanced detector losses . . . 111
4.4 Circuit simulation . . . 112
4.4.1 Simulation based on matrix permanents . . . 112
4.5 Corrective unitary rotations . . . 118
4.5.1 Teleportation with non-maximally entangled states . . . 118
4.5.2 Numerically calculated unitary rotations . . . 122
4.6 Simulation results . . . 124
4.7 Experimental results . . . 126
4.7.1 Experimental method . . . 126
4.7.2 Teleportation results . . . 129
4.8 Future devices . . . 132
4.8.1 Reasons for poor device performance . . . 133
4.8.2 Proposed modular chip design . . . 133

5 Transition-Edge Sensors **140**
5.1 Single Photon Detectors . . . 140
5.2 Available single-photon detector technology . . . 142
5.2.1 Conventional detectors . . . 144
5.2.2 Superconducting detectors . . . 146
5.3 Transition edge sensors . . . 151
5.3.1 Operating principle . . . 151

5.3.2 Electro-thermal feedback . . . 154
5.3.3 Readout . . . 156
5.3.4 Noise . . . 160
5.4 Characterisation of TES detectors . . . 166
5.4.1 Characterisation of detector noise . . . 167
5.4.2 Counting photons from coherent states . . . 171
5.4.3 Characterisation results . . . 177
5.4.4 Estimate of detection efficiency . . . 183

6 On-chip TESs 188
6.1 Evanescently coupled detectors . . . 191
6.1.1 Proof-of-principle detector . . . 192
6.1.2 Numerical modelling . . . 194
6.1.3 First results . . . 198
6.2 Increasing efficiency . . . 200
6.2.1 The 'Tadpole' design . . . 201
6.2.2 Device optimisation . . . 203
6.3 High-efficiency on-chip detector results . . . 211
6.3.1 Multiplexed detectors . . . 211
6.3.2 Results . . . 215

7 Conclusion 222
7.1 Summary . . . 223
7.1.1 Silica-on-silicon waveguide circuits . . . 223
7.1.2 TES detectors . . . 227
7.1.3 On-chip TES detectors . . . 229
7.2 Outlook . . . 230

Bibliography 233

A Extended quantum teleportation results 251

B List of symbols used in description of the TES 254

C Installation of TES detectors 256
C.1 Cryogenic system . . . 256
C.2 Wiring up the experimental insert . . . 258
C.2.1 Vacuum feedthroughs . . . 258
C.2.2 SQUIDs and breakout board . . . 259
C.2.3 Detector box and blackbody radiation . . . 260
C.3 Electronics and readout . . . 262
C.3.1 Noise reduction . . . 263

List of Figures

2.1 Single qubit unitary operations 24
2.2 Efficient construction of a general $N \times N$ linear interferometer 26
2.3 Hong-Ou-Mandel Interference 26
2.4 Construction of a two-qubit CZ-gate. 31
2.5 Generalised Hong-Ou-Mandel interference. 32
2.6 Heralded non-linear sign-shift gates. 33
2.7 The Ralph-Langford-White CZ-gate 34
2.8 Cascaded Ralph-Langford-White CZ-gates. 36
2.9 KLM scheme for efficient quantum computation using linear optics 37
2.10 Cluster-state model of computation 42
2.11 Single qubit rotations using cluster states 43

3.1 Heralded single photon source 51
3.2 Schematic of SPDC single photon source and characterisation methods 54
3.3 Gerchberg-Saxton phase recovery algorithm for efficient mode-matching of SPDC to single-mode fibre. 56
3.4 Predicted source coupling efficiency as a function of lens focal length 56
3.5 Spectral overlap of three single photons 57
3.6 Heralded two-photon HOM-dip in a fibre-beam-splitter 59
3.7 Silica-on-Silicon waveguide fabrication 60
3.8 Waveguide circuit schematic and experimental setup 65
3.9 Loss-tolerant beam splitter characterisation 70
3.10 Phase shifter characterisation 72
3.11 Results of two-photon quantum interference in a multi-port device 80
3.12 Three photon quantum interference in a multi-port interferometer 82
3.13 Boson Sampling concept 86
3.14 Three-photon Boson Sampling results. 90

4.1 The quantum teleportation protocol 94
4.2 The optimal classical teleportation strategy. 97
4.3 Photonic chip realisation of teleportation 103
4.4 A tunable beam splitter based on a Mach-Zehnder Interferometer. 105
4.5 Accessible states on the Bloch sphere 106

4.6 Experimentally prepared input states. . . . 108
4.7 Experimentally implemented measurement operators. . . . 109
4.8 Non-ideal behaviour of phase shifters. . . . 110
4.9 Partial photon distinguishability. . . . 116
4.10 Numerically computed unitary rotations 124
4.11 Histograms of simulated teleportation fidelity 125
4.12 Temporal alignment of the photons. . . . 127
4.13 Reconstructed density matrices of the teleported states. . . . 130
4.14 Measured and simulated fidelity of on-chip quantum teleportation. . . . 131
4.15 Modular chip design 134
4.16 Chip-to-chip gluing and characterisation with Bragg gratings. . . . 136
4.17 Reflection spectra from glued chip modules 139

5.1 The superconducting-to-normal transition 151
5.2 Basic operation of a Transition Edge Sensor 153
5.3 Schematic of the TES bias circuits 157
5.4 Noise sources in a TES 161
5.5 Noise contributions for our TES detector. . . . 164
5.6 Tuning the SQUIDs 167
5.7 Experimental setup for TES characterisation. . . . 172
5.8 Example TES voltage trace. . . . 174
5.9 Comparison of pulse filtering methods. . . . 174
5.10 Counting photons from coherent states. . . . 178
5.11 Measured TES energy resolution 179
5.12 Non-linear response of the SQUID array. . . . 180
5.13 Detector jitter. . . . 182
5.14 Setup to calibrate attenuators 183
5.15 Fitting a Poisson distribution to TES output. . . . 186
5.16 Measured detector efficiency 186

6.1 Evanescently coupled photon counting detector 191
6.2 The first proof-of-principle ECPCD device. . . . 193
6.3 Modelling results of the first ECPCD device 197
6.4 Results from the first ECPCD test-device 198
6.5 The ‘tadpole’ device design 200
6.6 Effect of the thermal gold-spine. . . . 202
6.7 Optical modelling of the tadpole design 203
6.8 Summary of thermal modelling 206
6.9 Effect of device length on energy resolution and optical absorption. . . . 209
6.10 Final device design simulation. . . . 211
6.11 Multiplexed detector concept. . . . 212
6.12 On-chip TES experimental setup 214
6.13 Photos of fabricated device. . . . 215
6.14 Results of high efficiency multiplexed detector. . . . 216
6.15 Classical optical loss measurements 220

A.1 Real-parts of the reconstructed density matrices. . . . 252

A.2 Imaginary-parts of the reconstructed density matrices. 253

C.1 Detail of data acquisition electronics 262
C.2 Installing detectors in the fridge. 264

List of Tables

3.1 Characterised beam splitter ratios. 72
3.2 Characterised phase shifter data. 73

4.1 Simulation of teleportation fidelities 126

5.1 Comparison of detector technologies 152
5.2 Estimated detector parameters 169
5.3 List of noise contributions. 172
5.4 Summary of measured detector parameters. 187

6.1 Comparison of waveguide-integrated single photon detectors 190
6.2 Final design parameters for the high-efficiency on-chip TES detector. 210
6.3 Predicted properties for high-efficiency on-chip TES detector 212
6.4 Derived individual detector efficiencies 217
6.5 Classically obtained losses at each detector element 221

A.1 Results of on-chip quantum teleportation. 251

List of abbreviations

UQC	universal quantum computer
TES	transition-edge sensor
KLM	Knill, Laflamme, and Milburn
LOQC	linear optical quantum computing
CV	continuous variable
DV	discrete variable
QIP	quantum information processing
CNOT	controlled-NOT
SU(2)	the group of 2×2 unitary matrices with determiniant 1
HOM	Hong-Ou-Mandel
SO(3)	the group of all 3D rotations
CZ	controlled sign-flip
NS	non-linear sign-shift
BSM	bell-state measurement
OWQC	one-way quantum computing
UV	ultraviolet
NV	nitrogen-vacancy
SPDC	spontaneous parametric down-conversion
TMSV	two-mode squeezed vacuum
BBO	Beta Barium Borate
KDP	Potassium Dihydrogen Phosphate
PMF	polarisation-maintaining fibre

PBS polarising beam splitter

EMCCD electron-multiplying charge coupled device

CCD charge-coupled device

SEM scanning electron microscope

FHD flame hydrolysis deposition

PEVCD plasma-enhanced chemical vapour deposition

DUW direct UV-write

FPGA field-programmable gate array

MZI Mach-Zehnder interferometer

GHZ Greenberger-Horne-Zeilinger

LOCC local operations and classical communication

POVM positive-operator valued measure

QST quantum state tomography

QPT quantum process tomography

APD avalanche photo-diode

NME non-maximally entangled

VGA v-groove array

SHG second-harmonic generation

NiCr Nichrome

SLED superluminescent light emitting diode

OSA optical spectrum analyser

FBS fibre beam splitter

PMT photo-multiplier tube

InGaAs Indium Gallium Arsenide

STJ superconducting tunnel juntion

MKID microwave kinetic inductance detector

FET field effect transistor

SNSPD superconducting nanowire single photon detector

NbN niobium nitride

WSi tungsten silicide

ETF electro-thermal feedback

SQUID superconducting quantum interference device

NIST national institute of standards and technology

FWHM full-width half-maximum

DCR dark count rate

PNR photon number resolution

AR anti-reflection

FLL flux-locked loop

PCB printed circuit board

DAQ data acquisition

PCI peripheral component interface

RF radio frequency

TFN thermal fluctuation noise

NEP noise equivalent power

RMS root mean square

FC ferrule connector

SMF single-mode fibre

FMM film mode-matching

EME eigenmode expansion

SPP surface plasmon polariton

TM transverse magnetic

TE transverse electric

GaAs Gallium Arsenide

NbTiN Niobium Titanium Nitride

SoS Silica-on-Silicon

SoI Silicon-on-Insulator

SiN Silicon-Nitride

CW continuous-wave

List of symbols

$\hat{U}$	a unitary opertor
$\hat{R}_{\hat{n}}(\theta)$	the rotation operator about the axis, $\hat{n}$, by an amount, θ
$\hat{\sigma}_x$	the Pauli-X operator
$\hat{\sigma}_y$	the Paui-Y operator
$\hat{\sigma}_z$	the Pauli-Z operator
$\mathbb{1}$	the identity operator
$\hat{X}$	the Pauli-X (bit-flip) gate
$\hat{Y}$	the Pauli-Y gate
$\hat{Z}$	the Pauli-Z (phase-flip) gate
$\hat{H}$	the Hadamard gate
$\|\psi\rangle$	a pure quantum state
ρ	a density matrix
$\|vac\rangle$	the vacuum state
$\|n\rangle$	the n-photon Fock state
$\|\phi_\nu(x)\rangle$	the single-particle state with quantum number, ν, at position, x
$\|\mathbf{0}\rangle$	the logical 0 state (in the computational basis)
$\|\mathbf{1}\rangle$	the logical 1 state (in the computational basis)
$\|+\rangle$	the rotated logical state - $\frac{1}{\sqrt{2}}(\|\mathbf{0}\rangle + \|\mathbf{1}\rangle)$
$\|-\rangle$	the anti-rotated logical state - $\frac{1}{\sqrt{2}}(\|\mathbf{0}\rangle - \|\mathbf{1}\rangle)$
$\|H\rangle$	the horizontally polarised photon state
$\|V\rangle$	the vertically polarised photon state
$\|D\rangle$	the diagonally polarised photon state

$\|A\rangle$	the anti-diagonally polarised photon state
$\|R\rangle$	the right-circularly polarised photon state
$\|L\rangle$	the left-circularly polarised photon state
$\|S\rangle$	a specific output distribution of photons
$\|T\rangle$	a specific input distribution of photons
$\|\Phi^+\rangle$	the maximally entangled two-mode triplet Bell state - $\frac{1}{\sqrt{2}}(\|\mathbf{00}\rangle + \|\mathbf{11}\rangle)$
$\|\Phi^-\rangle$	the maximally entangled two-mode triplet Bell state - $\frac{1}{\sqrt{2}}(\|\mathbf{00}\rangle - \|\mathbf{11}\rangle)$
$\|\Psi^+\rangle$	the maximally entangled two-mode triplet Bell state - $\frac{1}{\sqrt{2}}(\|\mathbf{01}\rangle + \|\mathbf{10}\rangle)$
$\|\Psi^-\rangle$	the maximally entangled two-mode singlet Bell state - $\frac{1}{\sqrt{2}}(\|\mathbf{01}\rangle - \|\mathbf{10}\rangle)$

Chapter 1

Introduction

"Information is physical"

Rolf Landauer

WHAT is information? As children we learn about information using physical objects. We use wooden blocks to represent numbers. Computations such as addition and subtraction, are carried out by physically moving blocks from one pile into another. The abacus is still used to perform calculations by traders and merchants across Asia and Africa. As we grow older, we learn more abstract methods for manipulating and storing information, carrying out additions 'in our heads' and using calculators to perform complex sums. We get used to thinking about information separately from the physical objects which represent it — the connected neurons in our brain or the localised electric charge in a calculator.

Yet, information *is* physical[1]. Information must be encoded in the state of a physical substrate and computation is achieved by manipulating this medium. A fundamental understanding of information thus requires reference to the physical laws which govern these devices. What physical computers can or cannot compute is determined by the laws of physics and not simply by pure mathematics. Therefore, it should be no surprise that we

must adapt our understanding of information and its computation in light of new physical theories.

1.1 Quantum information

Since its development at the beginning of the twentieth century, quantum mechanics has emerged as the most complete physical theory we have of our world. However, it was decades later before its implications for information theory were considered and a coherent discipline devoted to its study began to emerge. It was realised early on that modelling quantum mechanical systems exactly, or even approximately, becomes exponentially difficult as the number of particles, N, increases. This is because the abstract state space in which quantum mechanics operates — the Hilbert space — is exponentially larger than the familiar Euclidean parameter space of classical mechanics.

At the start of the 1980's Yuri Manin [2] and Richard Feynman [3] recognised one could instead circumvent this difficulty by building a *controllable* quantum system which one could use to simulate these complex quantum dynamics. Such a quantum simulator would be far more computationally powerful than any information processor relying on purely classical physics. David Deutsch significantly expanded and formalised these early ideas when he developed the model of a universal quantum computer (UQC) [4]. He showed how, in principle, one could build up quantum circuits acting on quantum bits (qubits), analogous to the bits of now-familiar classical electronic circuits, to create a quantum computer able to simulate *any* other quantum system efficiently. Although toy uses of such a computer were suggested by Deutsch soon after [5], it was another 10 years before Peter Shor discovered the first useful algorithm to run on such a universal quantum computer [6,7]. He developed an

algorithm that allowed a quantum computer to factor large co-primes efficiently and immediately transformed the field of quantum information from a niche, theoretical discipline into something of enormous practical importance. If one could build a computing device, based on the laws of quantum mechanics, it might be possible to carry out computations far more efficiently than a regular, classical computer could. The final key ingredient of *quantum error correction* was developed soon after [8–11] — making an actual physical realisation of such a machine feasible. The experimental race to build a universal quantum computer began when Cirac and Zoller showed how the elementary quantum logic operations outlined by Deutsch could be implemented using a string of trapped atomic ions [12]. Since then, there has been a huge world-wide effort to find and build suitable physical architectures on which the universal quantum computer may be realised [13–15].

This thesis is concerned with an *optical* implementation of a quantum information processor. Optical quantum information has been an important test bed for quantum protocols over the last decade and remains a promising route toward realising a universal quantum computer. In this approach, individual quantum particles of light, *photons*, can store coherent *quantum information* in any of the many degrees of freedom of the electromagnetic field and represent a promising avenue toward realising a universal quantum computer. One can imagine such a hypothetical optical quantum computer being composed of three parts: a *source* which generates appropriate quantum states of light; a *circuit* to manipulate the photons such that they perform a desired computation; and a *detector* such that one can make measurements of the output of the computation at the end. The work in this thesis is primarily concerned with the experimental implementation of the latter two parts: investigating the challenges in building the circuits and detectors necessary to realise a large-scale

photonic quantum information processor. In particular we argue that an *integrated* photonic architecture is likely required for future scalability and explore the challenges in using such an architecture to realise more complex circuits and achieve higher detection efficiencies. We leave the full overview of the work carried out in this thesis to the end of this chapter, and first look closer at the theory behind quantum computation.

1.1.1 The principle of superposition

What is it about quantum mechanics which enables a computational advantage? What makes it so different to the classical theories of Newton's mechanics and Maxwell's electromagnetism? To understand the origin of this difference we first look at how information is represented in quantum mechanics.

The *bit* is the fundamental unit of classical information and analogously we define the *qubit* to be the fundamental unit of information in quantum mechanics. Since all information must be represented by a physical state we need a way of describing them in quantum mechanics. Associated to any physical system is a complex vector space with an inner product — the Hilbert Space — in which every physical state [1] is described completely by a complex state vector of unit norm, $|\psi\rangle$. By definition, this state vector encompasses all the properties that determine the output probabilities of all the possible measurements performed on that state. The dynamical evolution of a state vector is governed by a first-order, linear differential equation — the Schrödinger equation. So far, this does not reveal anything fundamentally different from classical systems. To draw an analogy with classical mechanics, particles there exist in a Euclidean state space whose physical state can be

[1] For now we consider only 'pure' states which have no classical uncertainty. We will introduce the more general 'mixed' states later.

described by its position and momentum vector, $(\vec{r}, \vec{p})$. The evolution of this classical state vector can then be described by Hamilton's classical equations of motion.

The fundamental difference between classical and quantum states is captured by the *principle of superposition*. This states that whenever a quantum system can be found in either one of two pure states, $|\psi_1\rangle$ or $|\psi_2\rangle$, it can also be found in a pure state $|\Psi\rangle = a_1 |\psi_1\rangle + a_2 |\psi_2\rangle$, where a_1 and a_2 are complex numbers that satisfy $|a_1|^2 + |a_2|^2 = 1$. Connecting the resulting complex-valued vector describing a quantum state to a measurement outcome is done with the so-called Born rule. To give an example: when we measure the above state we will find it in the state $|\psi_1\rangle$, say, with a probability given by the Born rule — $P_1 = \langle\Psi| (|\psi_1\rangle\langle\psi_1|) |\Psi\rangle = |a_1|^2$. However, this does not simply reflect our limited ability as experimentalists to know which of the two states our physical system was originally prepared in, but represents the *fundamentally* probabilistic nature of the measurement outcome. We have defined $|\Psi\rangle$ to represent our total knowledge about a system and superposition states — very loosely described as occupying multiple states at once — are a new concept for physical reality without a clear classical analogue. In terms of encoding information, this means that if a physical system in the state $|\psi_1\rangle$ represents a logical $|\mathbf{0}\rangle_L$ and the state $|\psi_2\rangle$ represents a logical $|\mathbf{1}\rangle_L$ then a qubit may exist in the superposition:

$$|\psi\rangle = \alpha |\mathbf{0}\rangle_L + \beta |\mathbf{1}\rangle_L . \tag{1.1}$$

The states used to encode this superposition may be realised in a wide variety of physical systems but in photonic systems they are often represented using two orthogonal modes of the electromagnetic field — two distinct spatial modes for example.

1.1.2 Large Hilbert spaces

The consequences of the superposition principle for encoding information are revealed when we study ensembles of many quantum particles. If we consider N two-state classical particles (i.e. bits), the total state space is given by the Cartesian product and is of dimension $2N$. However, the superposition principle could not be formulated in such a small space — instead, we require the tensor product of N Hilbert spaces, $\mathcal{H}^{(N)} \equiv \mathcal{H}^{\otimes N}$. If each particle has an individual Hilbert space dimension of 2 (as states encoding qubits do), then N such quantum particles have a Hilbert space of total dimension 2^N.

This distinction between Cartesian and tensor products is precisely the phenomenon of quantum entanglement — the existence of pure states of a *composite* system that are *not* product states of the individual parts. Although some classical systems do exhibit linear superpositions (e.g. the space of states of a vibrating string with fixed endpoints), classical states of *separate* physical systems can never be superposed and so the Cartesian product is sufficient to describe all possible states in this system. In contrast, in quantum mechanics, there exist entangled states, which allows for any *general* superposition across many particles (i.e. including those which cannot be factored into products of individual particle states). This a state space which grows exponentially (2^N) as the number of particles, N increases.

From this example we arrive at the crux of the power of quantum mechanics. If we accept the superposition postulate, and the existence of entanglement, we are forced to accept that the state space describing an ensemble of N quantum bits is exponentially large. This is exactly the reason why large quantum systems are difficult to simulate. Even for relatively small system sizes with $N = 30$ we already need to keep track of $2^{30} \approx 10^9$ complex numbers throughout our calculation — an impossible task for even modern personal

computers. However, if we instead encode *information* in quantum particles and carefully control the system's evolution using an appropriate quantum circuit, the system can explore an exponentially large space. One can imagine, for a suitably designed algorithm, such a quantum system exploring a number of possible solutions to a problem which grows exponentially with the number of input qubits. This is the foundational basis for the potential computational speed-up of quantum computers.

1.2 Building an optical quantum computer

The previous section makes clear that in order to achieve the size of Hilbert space required to demonstrate a true quantum advantage over classical systems, we must construct large quantum states built from many particles. Of course, this alone is not enough. To gain any computational advantage, we must be able to control the evolution of this large quantum system to undergo non-trivial multi-particle quantum interference.

Quantum gates

As in classical computation we require two types of controllable operations to be able to implement any quantum circuit[15]. The first are single particle operations. The classical version is the NOT-gate which simply performs a bit-flip ($|\mathbf{0}\rangle \leftrightarrow |\mathbf{1}\rangle$). The quantum version simply corresponds to a rotation of the qubit's complex state vector in Hilbert space. The second type of operation is a two-particle gate which performs some operation on one of the particles depending on the state of the other. A classical example of such an operation is the NAND-gate which is known to be universal. One quantum example is the controlled-not (CNOT)-gate which performs a conditional bit-flip on particle two dependent on the state of

particle one. The unitary evolution of quantum states evolve demands that quantum gates must be reversible and so in this example, particle one must also be output in its original state. It turns out that if one can implement arbitrary single qubit operations and just *one* two-qubit operation then you can build any universal quantum computer acting on N qubits[15].

The great challenge of building a working quantum computer is to engineer a way of implementing these two-qubit interactions without introducing any interactions with the environment which would otherwise cause decoherence of the system into a useless classical *mixture* of states[14]. Significant experimental effort has been devoted to minimising this decoherence in a wide range of physical architectures. In general this requires ultra-cold temperatures, vacuum, cavities, optical traps or some combination of the above. Adding each additional qubit to such matter-based systems increases the potential for decoherence and has limited the number of qubits which can be stored and manipulated to around 14[16].

Encoding qubits

It is on this point that quantum computers based on optical systems have a great advantage[13]. Optical photons can be manipulated robustly through well-controlled parametric material interactions. The generation of pure single photons for use as qubits is an ongoing experimental challenge[17], but excellent approximations to such states may be readily achieved using only a pulsed laser, a suitable non-linear crystal and a single-photon detector (see section 3.1.1). With polarisation, photons also have a natural two-level system in which we can define our qubit. In fact, this belongs to the more general class of 'dual-rail' qubits where a single photon distributed over two orthogonal optical modes is used to define the

qubit[18]. Optical qubit encodings may use polarisation[19], spatial modes[20], frequency modes[21], time-bins[22] or even orbital angular momentum[23,24].

For example, in the polarisation basis we might define our logical zero state as a horizontally polarised photon - $|H\rangle$ and our logical one being a vertically polarised photon - $|V\rangle$. Thus our generalised qubit exists in the superposition state $|\psi\rangle = \alpha|H\rangle + \beta|V\rangle$ which we recognise as simply a photon with elliptical polarisation.

Single qubit gates are easy...

Photons have degrees of freedom that allow for very simple single qubit gates. For example, in the polarisation encoding, single qubit operations simply translate into rotations of half- and quarter-wave plates. More generally, single-qubits gates are easily realised for any photonic encoding which exhibit a natural $SU(2)$ symmetry[18]. Apart from polarisation, dual-rail spatial modes also exhibit this symmetry, where single-qubit gates are implemented using variable beam splitters and phase shifters (see chapter 2). However, the time and frequency degree of freedom do not posses this natural internal $SU(2)$ symmetry and single-qubit operations in these systems are more difficult to implement.

...two qubit gates are not

Of course, whilst the absence of photon-photon interactions is great for achieving a low-decoherence system, it makes implementing two-qubit gates very challenging[18]. However, it is now understood that the process of photon detection itself can lead to effective photon-photon nonlinearities, when used in systems of entangled photons[25] and we devote much of chapter 2 to detail how this works.

With detectors, we have everything we need

Finally, we must be able to make measurements of which state each photon is in at the end of our computation. Fortunately, the task of efficiently measuring single photons has attracted widespread interest across a variety of fields. Therefore, a range of high-performance detectors are now commercially available.

We have briefly seen how one can achieve the generation, initialisation, manipulation and detection of photonic qubits which means that, along with their exceptionally low decoherence, they satisfy the so-called DiVincenzo criteria for the implementation of a quantum computer[26].

Additionally, photonic qubits can be naturally interfaced with existing communication networks where their rapid spatial propagation makes them ideal carriers of quantum information over long distances[27]. For these reasons, it is perhaps no surprise that photonic experiments have played a leading role in demonstrating many of the basic principles of quantum-enhanced technologies[19,28–32]. Most of these initial demonstrations were carried out using 'bulk' optical components: mirrors, beam splitters and wave-plates, mounted to large optical tables. Whilst providing considerable flexibility in building these early small-scale experiments, it is unlikely that the proliferation of bulk optical components will be a scalable approach for building more complex quantum photonic systems. Instead, integrating these components on to chips along with waveguides provides a monolithic, alignment-free and inherently stable platform on which to build more complex circuits[20,33–35]. These integrated photonic circuits come with their own problems of additional photon loss and device deviations from design[36,37]. Nevertheless, the benefits of integration make it worthwhile to pursue solutions for these problems as we build more complex integrated circuits.

This thesis represents some effort toward this goal.

1.3 Thesis outline

Chapter two formalises the discussion in the previous section and details explicitly how one builds a universal quantum computer using photonic qubits. In particular we introduce linear optical quantum computing and discuss how both single- and two-qubit gates may be implemented using only beam splitters, phase shifters and photon detection. The technical details introduced in this chapter motivate the specific experiments carried out in the rest of the thesis - namely, realising more complex integrated circuits using more photons and developing high-efficiency number-resolving photon detection.

In chapter three we demonstrate genuine quantum interference between three single photons on a complex integrated circuit. We discuss optimisation of the heralding efficiency of a pure single photon source and introduce a low-loss silica-on-silicon waveguide platform which allows us to perform three-photon experiments. We use a novel element-wise characterisation technique and develop a statistical test to confirm that we have observed genuine three photon quantum interference. We show that this result finds immediate application in the first proof-of-principle test of the Boson Sampling problem.

Chapter four builds upon these results to show the first on-chip three-qubit protocol. We review quantum teleportation and show how it can be mapped onto our integrated circuit. The element-wise characterisation method is shown to be critical in mitigating fabricated deviations from design. A full theoretical model of the experiment is developed and it is shown that the fabricated components cause a modification of the final unitary rotations which must be accounted for in the data analysis. We successfully teleport three linearly

independent quantum states with an average fidelity beyond the classical limit using the reconfigurable integrated chip. The issues highlighted by this experiment leads us to propose a new design approach for future integrated photonic circuits.

The final two chapters discuss how high-efficiency superconducting transition-edge sensors (TESs) can be used to increase the performance of integrated photonic circuits. Chapter five introduces the TES and explains the principle of operation. We then detail the installation and characterisation of these detectors in a new lab in Oxford. In particular, the cryogenic and electronic read out systems along with associated noise sources are discussed and high-efficiency photon-number resolved measurements are made on attenuated coherent states.

Chapter six describes the integration of these TESs onto the silica-on-silicon waveguide architecture. We introduce the evanescently coupled detector design and present the results of the first proof-of-principle device. We then detail the theoretical modelling performed to optimise the device efficiency and show that the new device fabricated according to the optimised design behaves as expected. In particular, detector multiplexing is shown to be able to increase the total device efficiency and we discuss some novel applications of the evanescently-coupled design.

Chapter 2

Linear Optical Quantum Computing

"Each photon then interferes only with itself.

Interference between different photons never occurs"

The Principles of Quantum Mechanics, Paul Dirac

THE tremendous computational advantages promised by a universal quantum computer has ensured a huge range of physical architectures are being explored in the attempt to build such a machine[14]. Despite years of progress it is still not clear which one, if any, will be best suited to large-scale quantum computing tasks.

Information processors based on the quantum nature of light remain an attractive prospect[13,18]. Electromagnetic fields provide an innumerable array of optical modes with which to encode information. Since the energy of an optical photon far exceeds room temperature, each of these modes acts as a nearly ideal quantum harmonic oscillator prepared in its ground state. Moreover, there is essentially no coupling between these modes and the environment resulting in almost no decoherence. That we get this for 'free' with optical systems fundamentally sets it apart from other, matter-based, architectures such as trapped ions or atoms[38,39], superconducting qubits[40,41], electron spins within semiconductors[42],

or nitrogen vacancy centres[43]. These require every additional qubit to be introduced separately, carefully isolated from its environment and initialised into its ground state, often requiring extremely cold temperatures, very high vacuums or both. Of course, we now not only have to figure out how to populate these optical modes with single excitations but also how to manipulate them. Whilst the excellent coherence properties of optical systems provides the perfect platform for quantum communication[44], it effectively prevents us from engineering straightforward interactions between optical qubits.

In this chapter, we introduce the theoretical foundations of how linear optical systems can be used to build universal quantum computers. The results of this chapter will motivate the work carried out in the remainder of this thesis as well as providing the necessary theoretical background to understand the results of the experiments performed. We begin with the main theoretical tools needed to understand quantum states of light and the mathematical description of the linear optic elements we will use to manipulate them. The first proof that these linear optical elements are sufficient to build a universal quantum computer was provided in the seminal work of Knill, Laflamme and Milburn (KLM)[25]. We briefly summarise their original scheme since it clearly introduces most of the fundamental concepts behind all linear optical quantum computing (LOQC) schemes and thus already motivates the work in this thesis. In the decade since this work there have been many improvements to this scheme[18], and although they do not fundamentally alter the experimental requirements, they are worth reviewing.

2.1 Quantum optics

Information can be carried by either discrete (digital) or continuous (analogue) signals. In this thesis we will concentrate exclusively on the discrete-variable (DV) regime of quantum information processing — using qubits, each encoded by a single photon. We define our qubit using a single photonic excitation distributed over two orthogonal optical modes (dual-rail encoding). The quantum counterpart to analogue signals is the continuous-variable (CV) regime, which instead encodes information in the phase and amplitude quadratures of the optical field[45].

Historically, the digital revolution came about because digital (discrete) information is more amenable to error correction techniques, making computation and communication more robust[46]. Similarly, the first fault-tolerant quantum error-correcting codes were developed for the discrete variable regime[9,10] making it a popular choice for concentrating experimental efforts. However, there has recently been some proposals to apply these ideas to the CV regime by making use of qu*d*its defined in the continuous space[47–49]. It is a matter of continuing debate which is ultimately better-suited for large-scale QIP but the language, intuition and tools of the two regimes are different enough that it makes sense to limit discussion to only the discrete-regime for the rest of this thesis.

We first review the quantum mechanics of ensembles of indistinguishable particles due to its central importance for linear optical quantum computing before defining more precisely what we mean by 'linear optics' and 'photons'.

2.1.1 Indistinguishable particles

The existence of identical particles has interesting consequences for the symmetry of the overall quantum state. For truly identical particles, the probability density of any observable must be invariant under the action of the particle exchange operator, $P_{i,j}$,

$$P_{i,j} \left|\phi_1, \phi_2, \ldots, \phi_i, \phi_j, \ldots, \phi_m\right\rangle = e^{i\varphi} \left|\phi_1, \phi_2, \ldots, \phi_j, \phi_i, \ldots, \phi_m\right\rangle . \tag{2.1}$$

Under a further exchange operation we must recover the initial state, from which we immediately find that $\varphi = 0$ or π. Therefore we have states for which $|\Psi\rangle$ is *even* under permutation, $P|\Psi\rangle = |\Psi\rangle$ — *Bosons*, or states for which $|\Psi\rangle$ is *odd* under permutation, $P|\Psi\rangle = -|\Psi\rangle$ — *Fermions*[1].

The work in this thesis is based on systems of identical non-interacting bosons — photons. We want to construct an appropriate many-particle state from a collection of non-interacting single particle wavefunctions, $\phi_\nu(x)$, where ν are the quantum numbers identifying the state and x is some coordinate of the particle. Writing an N-particle state in the basis of one-particle states requires that under particle exchange $x_j \leftrightarrow x_k$, the overall wavefunction obeys the correct symmetry. For bosons, this wavefunction must be symmetric

[1]For particles which are effectively constrained to move in only 2 dimensions there turns out to be more exotic possibilities called *anyons*, the simplest versions of which permit any intermediate phase angle φ under the exchange of two particles. The correctly symmeterised wavefunctions for anyons can be constructed by forming the so-called matrix *immanent* which is a generalisation of the permanent / determinant as used for Bosons and Fermions.

and can be expressed as,

$$\Phi(x_1, ..., x_n) = \frac{1}{\sqrt{N!}\sqrt{\prod_\nu n_\nu!}} \begin{vmatrix} \phi_{\nu 1}(x_1) & \cdots & \phi_{\nu 1}(x_N) \\ \vdots & \ddots & \vdots \\ \phi_{\nu N}(x_1) & \cdots & \phi_{\nu N}(x_N) \end{vmatrix}_+ , \tag{2.2}$$

where $|\cdot|_+$ denotes the matrix permanent and n_ν counts the number of particles in each state ν. For fermions, the corresponding anti-symmetric wavefunction is given by exchanging the permanent in equation (2.2) for the Slater determinant, in which half of the terms have negative signs. This small difference turns out to have huge consequences for the computational complexity of systems of identical fermions and bosons and will be discussed further in section 3.5 in the context of the Boson Sampling problem [50].

This notation soon becomes very cumbersome to use as the number of identical particles in the system increases. Instead, we use the formalism of *second quantisation* where the intrinsic symmetry crucial to the physics of ensembles of identical particles is now carried by the operators acting on states rather than by the states themselves.

Central to this formalism is the definition of a *creation operator* whose action is defined by,

$$\hat{a}^\dagger |\psi_1, ..., \psi_N\rangle = |\psi_1, ..., \psi_N, \psi_{N+1}\rangle , \tag{2.3}$$

i.e., the creation operator, $\hat{a}^\dagger$ maps the correctly symmeterised N-particle state $|\psi_1, ..., \psi_N\rangle$ to the correctly symmeterised $(N+1)$-particle state $|\psi_1, ..., \psi_N, \psi_{N+1}\rangle$. We also define the

annihilation operator $\hat{a}$ to be the adjoint of $\hat{a}^\dagger$,

$$\langle \chi_1, ..., \chi_{N-1} | \, \hat{a} \, | \psi_1, ..., \psi_N \rangle = \langle \psi_1, ..., \psi_N | \, \hat{a}^\dagger \, | \chi_1, ..., \chi_{N-1} \rangle^* \tag{2.4}$$

$$= \langle \psi_1, ..., \psi_N | \chi_1, ..., \chi_N \rangle^* . \tag{2.5}$$

If we expand these symmeterised states in the matrix permanent fashion as in equation (2.2), it is not difficult to show [51] that for these creation and annihilation operators to generate states of the correct bosonic symmetry they must obey the set of canonical commutation relations,

$$[\hat{a}_\alpha, \hat{a}_\beta] = \left[\hat{a}^\dagger_\alpha, \hat{a}^\dagger_\beta\right] = 0 \qquad \left[\hat{a}_\alpha, \hat{a}^\dagger_\beta\right] = \delta_{\alpha,\beta}. \tag{2.6}$$

The entire symmetry of the multi-particle state is contained within these relations. The creation operators for fermionic systems are likewise endowed with a set of canonical *anti*-commutation relations. We are now in a position to choose a more convenient representation of our many-particle quantum state. We use the *occupation-number* representation in which states are labelled by the number of particles, n_k in the single-particle state k. This set of occupation numbers is exactly equivalent to the expression in equation (2.2) but results in an inherently symmetric basis vector which we denote,

$$|\Phi\rangle = |n_1, n_2, ..., n_k\rangle . \tag{2.7}$$

For bosons, n_k are non-negative integers whereas for fermions they are either 0 or 1 (Pauli exclusion principle). The space spanned by these basis vectors is called Fock space. In this representation, it is straightforward to show that the creation and annihilation operators

for bosons behave like,

$$\hat{a}_\alpha^\dagger \left|n_1, ..., n_\alpha, ...\right\rangle = \sqrt{n_\alpha + 1} \left|n_1, ..., n_\alpha + 1, ...\right\rangle \tag{2.8}$$

$$\hat{a}_\alpha \left|n_1, ..., n_\alpha, ...\right\rangle = \sqrt{n_\alpha} \left|n_1, ..., n_\alpha - 1, ...\right\rangle . \tag{2.9}$$

We labour this discussion of symmetry because of its fundamental importance for optical quantum computing. The defining 'feature' of using optical qubits is the absence of any interaction between single photons. The only method of realising non-trivial multi-qubit operations is by relying on this exchange symmetry of identical multi-particle states. This so-called exchange interaction is able to generate the necessary non-classical multi-particle correlations (entanglement) without needing a direct particle-particle interaction (see section 2.1.4 on Hong-Ou-Mandel interference, for example). However, the power of the representation of second-quantisation means that this underlying symmetry is often forgotten. We hope this brief review of second quantisation serves as a reminder of the fundamental importance the symmetry of ensembles of identical particles plays in realising a photonic quantum computer.

2.1.2 Quantising the electromagnetic field

The previous section was rather abstract - we argued that indistinguishability plays a fundamental role in quantum mechanics and this leads to important symmetry requirements for multi-particle states. The second-quantisation formalism is a convenient method of representing such states which ensures the correct symmetry is preserved. Now we wish to connect this representation to something more physical - i.e. photons of the quantised electromagnetic field. We first note that states of definite particle number, $|n_k\rangle$, are eigenstates of the

quantum harmonic oscillator with Hamiltonian $\hat{H} = \sum_k \hbar\omega_k(\hat{a}_k^\dagger \hat{a}_k + 1/2)$ and eigenvalues $E_{n_k} = \hbar\omega_k(n_k + 1/2)$, where ω_k defines the frequency of each mode k. Thus we can attribute a physical meaning to these previously abstract operators — $\hat{a}_k^\dagger$ creates- and $\hat{a}_k$ destroys- a single excitation of a specific mode, k, of a general quantum harmonic oscillator.

From Maxwell's equations we know that the vector potential $\boldsymbol{A}(\boldsymbol{r}, t)$ satisfies the wave equation,

$$\nabla^2 \boldsymbol{A}(\boldsymbol{r}, t) = \frac{1}{c^2} \frac{\partial^2 \boldsymbol{A}(\boldsymbol{r}, t)}{\partial t^2}. \tag{2.10}$$

which we can expand in terms of a set of orthogonal mode functions labelled by k. This has a separable solution in which the time-varying part is the solution to a simple harmonic oscillator, $\boldsymbol{A}(t) \propto \sum_k (a_k e^{-i\omega_k t} + a_k^* e^{i\omega_k t})$, where a_k are dimensionless complex Fourier mode coefficients. We now quantise the field by associating these complex mode coefficients with their counterpart quantum operators in the solution to the quantum harmonic oscillator. Thus, the final quantised field takes the form,

$$\boldsymbol{A}(\boldsymbol{r}, t) = \sum_k \sqrt{\frac{\hbar}{2\omega_k \epsilon_0}} \left[\hat{a}_k \boldsymbol{u}_k(\boldsymbol{r}) e^{-i\omega_k t} + \hat{a}_k^\dagger \boldsymbol{u}_k^*(\boldsymbol{r}) e^{i\omega_k t} \right], \tag{2.11}$$

where the functions $\boldsymbol{u}_k(\boldsymbol{r})$ describe the spatial distribution and polarisation of each mode k and ϵ_0 is the permittivity of free-space.

It is worth reflecting a little more on the analogy between the quantised electromagnetic field and the quantum harmonic oscillator. The most important difference is that in the harmonic oscillator model we think of a single particle having a definite energy $E_n = \hbar\omega(n + \frac{1}{2})$ so that the quantum number n becomes a shorthand for the energy of that particle. In contrast, we define photons as the quantised excitations of a particular mode, k of the

electromagnetic field carrying a definite energy $E_k = \hbar\omega_k$. Now we must identify n as literally the number of photons in that given mode such that the overall energy agrees with that of the harmonic oscillator. Thus, $\hat{a}_k^\dagger$ creates, and $\hat{a}_k$ destroys, a single photon in the optical mode, k.

2.1.3 Linear optics

The aim of this chapter is to show how a universal quantum computer may be built from linear optical elements. Whilst very recent results have shown matter-mediated two-photon interactions are feasible, they are currently extremely technically challenging to implement [52,53]. Thus, we restrict ourselves to using only linear optical elements; defined as anything whose Hamiltonian is bi-linear in its creation and annihilation operators and conserves particle number,

$$\hat{H}_{\text{linear}} = \sum_{i,j} A_{i,j} \hat{a}_i^\dagger \hat{a}_j, \tag{2.12}$$

such that each particle evolves independently of every other and thus does not permit any particle interaction. Optical elements of this form have the convenient property that mode transformations of creation operators take the particularly simple form,

$$\hat{a}_j^\dagger \to e^{-i\hat{H}_{\text{linear}}} \hat{a}_j^\dagger e^{i\hat{H}_{\text{linear}}} = \sum_{k=1}^{n} U_{j,k} \hat{b}_k^\dagger. \tag{2.13}$$

This is a powerful tool in linear optics and means that once the effective $N \times N$ unitary transformation $U_{j,k}$ of a linear optical device is known, it is possible to compute the output of this device for any possible input state (although this is often not an efficient calculation!). For the Hamiltonian to be Hermitian, we require that $\hat{A}$ be a unitary matrix. Since the

smallest unit of quantum information, the qubit, is defined over two optical modes we will first restrict our discussion to linear optical transformations on these two modes before considering the general $N \times N$ case. Any 2×2 unitary matrix can be expanded in terms of two Pauli matrices, $\hat{A} = \frac{\theta}{2} i\hat{\sigma}_y + \frac{\phi}{2}\hat{\sigma}_z$ from which emerge two distinct classes of linear mode transformations — phase shifters and beam splitters.

Phase Shifters

For $\hat{A} = \frac{\phi}{2}\hat{\sigma}_z$ we have $\hat{H}_{PS} = \frac{\phi}{2}(\hat{a}_0^\dagger \hat{a}_0 - \hat{a}_1^\dagger \hat{a}_1)$ which, in the Heisenberg picture, leads to a time-evolution of creation operators,

$$\hat{a}_0^\dagger \rightarrow e^{i\hat{H}_{PS}} \hat{a}_0^\dagger e^{-i\hat{H}_{PS}^\dagger} = e^{i\phi/2}\hat{a}_0^\dagger \tag{2.14}$$

$$\hat{a}_1^\dagger \rightarrow e^{i\hat{H}_{PS}} \hat{a}_1^\dagger e^{-i\hat{H}_{PS}^\dagger} = e^{-i\phi/2}\hat{a}_1^\dagger. \tag{2.15}$$

Alternatively, in the Schrödinger picture, we can represent this as a unitary rotation on a two-mode state,

$$\hat{U}_{PS} = e^{-i\phi\hat{\sigma}_z/2} = \hat{R}_z(\phi) = e^{i\phi/2} \begin{pmatrix} 1 & 0 \\ 0 & e^{-i\phi} \end{pmatrix}, \tag{2.16}$$

from which the connection to rotation operators can clearly be seen. More formally, there is an homomorphic mapping between the group of unitary transformations on 2 modes, $SU(2)$, and the group of rotations in 3 dimensions, $SO(3)$. Accordingly, a phase shifter mode transformation can be thought of as a unitary rotation around the z-axis by an amount ϕ. We will come back to this point below when we introduce the Bloch sphere.

Physically, this phase shifter operation can be implemented by selectively retarding the phase of one optical mode with respect to another. This can be achieved by changing the

linear dielectric properties (refractive index) of the propagation medium via the strain-, electro- or thermo-optic effect, or by physically increasing the propagation length of one of the optical modes. In this thesis, we will mostly use the thermo-optic effect to increase the refractive index of a specific waveguide and thus induce controllable relative phases between optical modes, effectively performing a rotation around the Bloch sphere z-axis.

Beam splitters

The second type of linear optical element is described by $\hat{A} = \frac{i\theta}{2}\hat{\sigma}_y$ which has a Hamiltonian of the form, $\hat{H}_{BS} = \frac{i\theta}{2}(\hat{a}_0^\dagger \hat{a}_1 + \hat{a}_1^\dagger \hat{a}_0)$ and induces transformations,

$$\hat{a}_0^\dagger \rightarrow e^{i\hat{H}_{BS}}\hat{a}_0^\dagger e^{-i\hat{H}_{BS}} = \cos(\theta/2)\hat{a}_0^\dagger + \sin(\theta/2)\hat{a}_1^\dagger \tag{2.17}$$

$$\hat{a}_1^\dagger \rightarrow e^{i\hat{H}_{BS}}\hat{a}_1^\dagger e^{-i\hat{H}_{BS}} = -\sin(\theta/2)\hat{a}_0^\dagger + \cos(\theta/2)\hat{a}_1^\dagger, \tag{2.18}$$

which, again, we see is equivalent to a unitary rotation in the Bloch sphere, this time around the y-axis,

$$\hat{U}_{BS} = e^{-i\theta/2\hat{\sigma}_y} = \hat{R}_y(\theta) = \begin{pmatrix} \cos\theta/2 & -\sin\theta/2 \\ \sin\theta/2 & \cos\theta/2 \end{pmatrix}. \tag{2.19}$$

The physical realisation of such a transformation depends on the type of optical mode we wish to manipulate. For spatially encoded qubits, this is exactly the transformation of a lossless beam splitter, or half-silvered mirror [54]. The reflectivity of this beam splitter, η is related to the Hamiltonian above through, $\eta = \cos^2\theta/2$. For polarisation encoded qubits, this beam splitter transformation can be realised using a half-wave plate and for time or frequency-encoded qubits, a time non-stationary device such as a modulator or switch is required [21,22,55,56].

The Bloch Sphere

Any arbitrary 2×2 unitary operation may be performed with a combination of these beam splitter and phase shifter elements as shown in figure 2.1, which follows directly from our earlier decomposition of $\hat{A}$ in terms of the two Pauli operators, $\hat{\sigma}_z$ and $\hat{\sigma}_y$, used to construct each element. The connection between these unitary transformations in $SU(2)$ and rotations in $SO(3)$ becomes particularly useful when considering the representation of a qubit on the *Bloch sphere*. An arbitrary qubit, $|\psi\rangle = \cos\theta/2\,|\mathbf{0}\rangle + e^{-i\phi}\sin\theta/2\,|\mathbf{1}\rangle$ may be represented as a point on a unit sphere with polar coordinates, θ and ϕ. Arbitrary single qubit unitary operations may then be thought of as a series of rotations of this point around two orthogonal axes. In particular, using the phase shifters and beam splitters as described above we may implement an arbitrary unitary operation with the following three rotations,

$$\hat{U} = \hat{R}_z(\phi_1)\hat{R}_y(\theta)\hat{R}_z(\phi_2), \tag{2.20}$$

for appropriate choices of ϕ_1, ϕ_2 and θ. If the qubit is either prepared or measured in the computational basis (that is, in the $|\mathbf{0}\rangle$ or $|\mathbf{1}\rangle$ basis) then the first or second z-rotation may be dropped respectively. We will return to the experimental implementation of this in section 4.2.1.

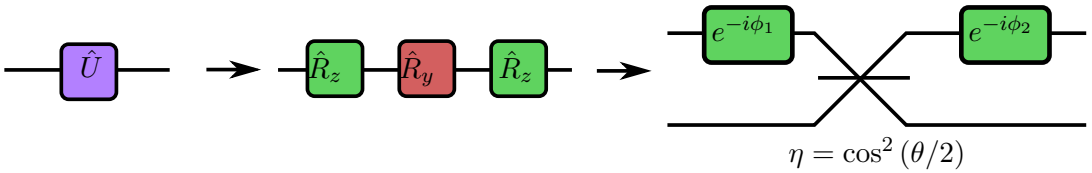


Figure 2.1: Single-qubit unitary operation. Any single qubit operation, $\hat{U}$ can be decomposed into three rotations on the Bloch sphere around the z- and y-axes. Physically, these rotations can be deterministically enacted using single mode phase shifters and beam splitters of tunable reflectivity, $\eta = \cos^2(\theta/2)$.

General $N \times N$ mode transformations

Using these two linear optical elements, two-mode phase shifters and beam splitters, it is possible to efficiently construct any unitary transformation on N modes. In this way we can build up more complex quantum circuits which act on many photons or qubits at once from these two elements for which we have a completely physical representation. The general scheme for this decomposition was given by Reck et al.[57] where they showed a maximum of $N(N-1)/2$ general two-mode elements were required to implement any N-port mode transformation as shown in figure 2.2. Specific transformations on N modes may be generated with far fewer elements than this bound and indeed we exploit this fact in the design of the circuit used in chapters 3 and 4.

It is important to note that this construction alone does not allow us to efficiently build a quantum computer. Although we may build any N-port interferometer efficiently (that is with a number of optical elements that scales as a polynomial of N), a quantum computer that encodes each term of a state vector as an optical mode would require exponentially many modes (2^N modes for an N-qubit state) and thus an exponential number of optical elements. In fact, such interferometers were used to implement a semi-classical simulation of quantum circuits before the development of an efficient linear optics protocol[58,59]. A single photon would start in one of the 2^N-modes and propagate through the N-port interferometer ending in a superposition over all 2^N output states from which the result of the computation could be read out. It is straightforward to see why such a scheme could never be an efficient implementation of a quantum computer: the dynamics of a single photon are exactly equivalent to that of a classical wave and rely on only classical, first-order interference. Genuine, efficient, optical quantum computation based on linear optics and single photons

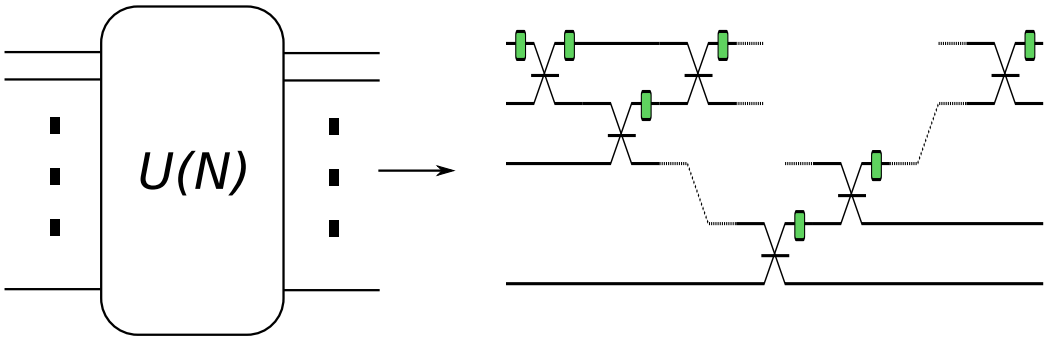


Figure 2.2: Efficient construction of a general $N \times N$ linear interferometer. Given a unitary linear transformation on N-modes, $U(N)$, Reck and Zeilinger showed that it is possible to decompose it into a network of beam splitters and phase shifters (green squares) as shown on the right. Using this scheme requires at most $N(N-1)/2$ beam splitters.

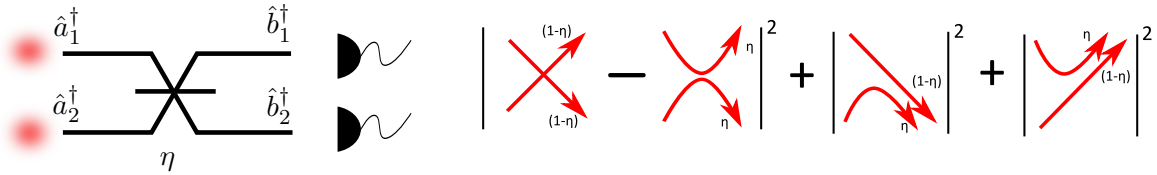


Figure 2.3: Hong-Ou-Mandel Interference. Two single photons are coupled into each input port of a beam splitter of reflectivity η. The output ports are monitored with number-resolving detectors. The four possible outcomes of such a setup are shown on the right. The two outcomes which lead to photons exiting at different ports can be made indistinguishable provided the beam splitter is balanced and the photons are identical. The bosonic symmetry of indistinguishable photons leads to destructive interference of these terms and thus the photons never emerge from different output ports. This is HOM interference.

requires higher-order non-classical interference, such as the Hong-Ou-Mandel effect.

2.1.4 The Hong-Ou-Mandel effect

We will now bring together the discussions of linear optical elements and the indistinguishable particles formalism to study what happens when two identical photons undergo a beam splitter transformation. This elementary case will allow us to appreciate the effect of correlation without interaction which underpins the fundamental working principles of linear optical quantum computing.

We consider an experiment as described in figure 2.3. Each of the two input ports to a beam splitter are populated with a single photon state such that the initial state can be

written,

$$|\psi_{\text{ini}}\rangle = |1,1\rangle_{1,2} = \hat{a}_1^\dagger \hat{a}_2^\dagger |\text{vac}\rangle \,. \tag{2.21}$$

The action of the beam splitter transforms the photons into a coherent superposition of possible output states. Using the transformations derived in equation (2.17) we can express the final state as,

$$\begin{aligned} |\psi_{\text{fin}}\rangle &= \left(\hat{U}_{BS}\hat{a}_1^\dagger \hat{U}_{BS}^\dagger\right)\left(\hat{U}_{BS}\hat{a}_2^\dagger \hat{U}_{BS}^\dagger\right)|\text{vac}\rangle \qquad (2.22)\\ &= \left(\cos(\theta/2)\hat{b}_1^\dagger + \sin(\theta/2)\hat{b}_2^\dagger\right)\left(-\sin(\theta/2)\hat{b}_1^\dagger + \cos(\theta/2)\hat{b}_2^\dagger\right)|\text{vac}\rangle \\ &= \left(-\cos(\theta/2)\sin(\theta/2)(\hat{b}_1^\dagger)^2 + \cos(\theta/2)\sin(\theta/2)(\hat{b}_2^\dagger)^2 + \cos^2(\theta/2)\hat{b}_1^\dagger\hat{b}_2^\dagger - \sin^2(\theta/2)\hat{b}_2^\dagger\hat{b}_1^\dagger\right)|\text{vac}\rangle \end{aligned}$$

where we denote the *output*-mode creation operators using $\hat{b}^\dagger$ for clarity. The four terms correspond to the four possible ways the photons can pass through the beam splitter — either both are reflected, both are transmitted, or one is reflected and the other transmitted which can happen with either photon. In the special case of a 50:50 beam splitter we have $\eta = 1/2 \rightarrow \theta = \pi/2$ and we find,

$$\begin{aligned} |\psi_{\text{fin}}\rangle &= \left(-\frac{1}{2}(\hat{b}_1^\dagger)^2 + \frac{1}{2}(\hat{b}_2^\dagger)^2 + \frac{1}{2}[\hat{b}_1^\dagger, \hat{b}_2^\dagger]\right)|\text{vac}\rangle \qquad (2.23)\\ &= -\frac{1}{\sqrt{2}}\left(|2,0\rangle - |0,2\rangle\right), \qquad (2.24) \end{aligned}$$

where we have crucially made use of the bosonic commutation relation, $[\hat{b}_1^\dagger, \hat{b}_2^\dagger] = 0$. This is quite a remarkable outcome! The photons emerge either *both* from one port or *both* from the other - the photons are perfectly correlated despite never having interacted. This effect is known as Hong-Ou-Mandel (HOM) interference after the three physicists who first observed

this[60]. It is one of the most striking consequences of the existence of indistinguishability in quantum mechanics. It is the interference between the two manifestly indistinguishable paths for photons exiting in different ports which leads to this result. That such two-photon correlations can be generated without requiring a non-linear photon interaction underlies all of the known schemes for building a quantum computer from linear optics.

We often characterise this effect by counting how many photons emerge from the different ports of the beam splitter simultaneously. From this we can determine the probability of the photon emerging from different ports which we call $P_{1,2}^{(\text{indist})}$. When the photons are completely indistinguishable, this ideally falls to exactly zero. From equation (2.23), we see that if the photons can be distinguished (for example by introducing a temporal delay) then no interference occurs and instead $P_{1,2}^{(\text{dist})} = 1/2$. By varying the relative temporal delay of the photons we can trace out a 'dip' in the number of coincident detections. Determining the resulting visibility of this 'HOM-dip', $V = (P_{1,2}^{(\text{indist})} - P_{1,2}^{(\text{dist})})/P_{1,2}^{(\text{indist})}$ is a common method of characterising this effect.

In contrast to the first-order interference which governs the evolution of single photons through N-port interferometers and forms the basis of single-qubit operations, this second-order interference effect cannot be understood using the classical wave equations of Maxwell. It can be shown that when we are restricted to using the classical wave equations, it is not possible to observe any second-order interference effect with unit visibility if there is no first-order interference (i.e. if there is no definite fixed phase relation between the two photons). This leads to a classical visibility limit of $V_{\text{class}} = 1/2$ for a balanced beam splitter[61], an observed visibility greater than this forces us to accept that quantum interference (i.e. interference of indistinguishable outcomes) has occurred[62,63].

The difficult of mutually matching all of the physical properties of the two photons constitutes one of the greatest challenges in the laboratory and it is still a cause for celebration when a HOM-dip with high visibility is observed. In this thesis, we will often make use of this visibility to estimate how indistinguishable our photons are and to verify that quantum interference is taking place.

2.2 The Knill-Laflamme-Milburn (KLM) scheme

Single photons were one of the very first systems used to study quantum information protocols due to the relative ease with which one can generate, manipulate and measure single quanta of light [28,29]. However, the lack of a direct photon-photon interaction was thought to require a matter-mediated non-linearity to create a two-qubit gate [15].

In 2001 Knill, Laflamme and Milburn (KLM) caused a revolution in the optical quantum computing community by constructing a scheme for efficient universal quantum computation using only linear optics, photo-detection and feed-forward [25]. They showed how the 'correlation-without-interaction' of the HOM effect could be combined with the inherent non-linearity of photo-detection and conditional feed-forward operations to build near-deterministic two-photon quantum logic gates.

Whilst their proof that linear optics quantum computing was in principle scalable, the scheme as originally presented required enormous overheads, needing many hundreds of ancillary single photon states to create a single two-qubit gate. However, the renewed interest in optical quantum computing as a serious contender soon saw many improvements to this protocol which were far more resource efficient. In fact, it is now incredibly unlikely that any full-scale universal optical quantum computer would ever be built according to

the original KLM scheme[18]. However, the underlying principles for all of these schemes (measurement-induced non-linearity and gate-teleportation) are essentially equivalent and the KLM scheme remains one of the clearest ways of introducing these concepts. The purpose of this section is to more clearly motivate the work undertaken in later chapters, in particular the requirement for many indistinguishable single photons, large, multi-port optical interferometers, and high efficiency number-resolved photo detection. Experiments are still at a stage where the specific protocol to be implemented has little bearing on the devices being built. However, for completeness, we will end this section with a brief discussion of the major advancements in linear optics quantum computing protocol design.

2.2.1 Single qubit gates

Section 2.1.3 detailed how deterministic unitary transformations on single qubits can be enacted with variable beam splitters and phase shifters. The specifics on how such components are built will be left until section 4.2.1 but it suffices for now to know that these deterministic arbitrary single qubit operations are possible.

2.2.2 Two-qubit gates

We know that along with arbitrary single qubit gates we also need a suitable two-qubit gate in order to construct a quantum computer[64]. In fact, almost any two-qubit gate will suffice[65] and for the remainder of this section we will concentrate on the controlled-phase or controlled-Z gate which we will write as CZ. The action of this gate on two qubits, traditionally labelled the control and target qubits, is defined as,

$$\alpha\left|\mathbf{00}\right\rangle_{c,t}+\beta\left|\mathbf{01}\right\rangle_{c,t}+\gamma\left|\mathbf{10}\right\rangle_{c,t}+\delta\left|\mathbf{11}\right\rangle_{c,t}\xrightarrow{\text{CZ}}\alpha\left|\mathbf{00}\right\rangle_{c,t}+\beta\left|\mathbf{01}\right\rangle_{c,t}+\gamma\left|\mathbf{10}\right\rangle_{c,t}-\delta\left|\mathbf{11}\right\rangle_{c,t}, \quad (2.25)$$

where the bold font kets are understood to represent the two logical qubit states to distinguish them from photon number states. In the dual rail scheme, these can be written in the Fock basis as $|\mathbf{0}\rangle_c = |1,0\rangle$ and $|\mathbf{1}\rangle_c = |0,1\rangle$. Therefore, a logical two-qubit state would be expanded into four physical modes, for example $|\mathbf{01}\rangle_{c,t} \rightarrow |1,0\rangle_c \otimes |0,1\rangle_t$.

We can imagine such a gate being constructed in the dual-rail scheme as shown in figure 2.4. Using this scheme we have reduced the two qubit gate to being able to implement

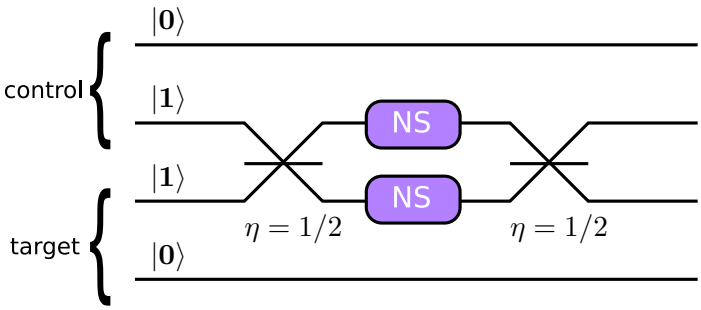


Figure 2.4: Construction of a two-qubit CZ-gate from two non-linear sign shift (NS) gates. See the main text for details of operation.

the non-linear sign shift gate (NS-gate), which acts on superpositions of Fock states as,

$$\alpha\,|0\rangle + \beta\,|1\rangle + \gamma\,|2\rangle \xrightarrow{\text{NS}} \alpha\,|0\rangle + \beta\,|1\rangle - \gamma\,|2\rangle\,. \tag{2.26}$$

With two of these gates embedded in an interferometer as shown in figure 2.4 we can see that when either the control and target qubits are in the logical state $|\mathbf{0}\rangle$ the NS-gate has no effect and the overall operation is the identity. However, when both the control and target qubits are in the logical $|\mathbf{1}\rangle$ state, there are two photons incident on the first beam splitter which then undergo HOM interference, $|1,1\rangle \rightarrow \frac{1}{\sqrt{2}}(|2,0\rangle - |0,2\rangle)$. This state acquires a phase flip under the action of the NS-gates which, after a 'reverse'-HOM interference at the second beam splitter completes the desired action of the CZ-gate — $|\mathbf{11}\rangle_{c,t} \rightarrow -|\mathbf{11}\rangle_{c,t}$.

We know that physically we cannot construct such a deterministic NS-gate because

photons do not interact with each other. We have seen above that the only way photons can influence each other is via their bosonic symmetry in HOM-type effects. KLM were the first to show that one can combine this effect with ancilla resource photons and photo-detection to create probabilistic, but heralded, two-photon gates. To start, we consider the simplest generalisation of the HOM-effect as shown in figure 2.5. This consists of a beam splitter of

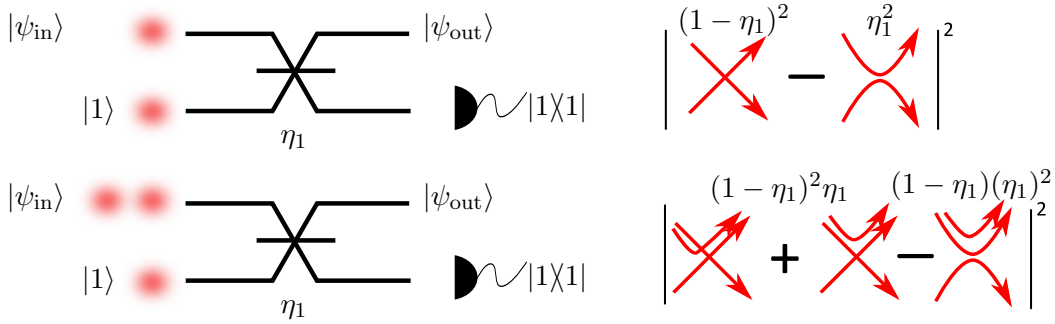


Figure 2.5: Generalised Hong-Ou-Mandel interference. In this simple extension to the HOM-effect we now consider up to two photons being in the upper input mode, $|\psi_{\text{in}}\rangle$. Conditioning on detecting one, and only one, photon in the lower mode, the output state, $|\psi_{\text{out}}\rangle$, depends on the HOM-like interferences of the probability amplitudes as shown on the right. Choosing $\eta_1 < 1/2$ results in a phase flip occurring for the two photon input state when compared to the single photon state. However, a second beam splitter is required to balance the post-selected probability amplitudes to implement the full NS-gate as seen in figure 2.6

reflectivity η_1 with our input state of zero, one or two photons coupled into the top mode, whilst an ancilla single photon state is coupled into the lower mode . Since the NS-gate conserves photon number, our herald event must include detection of the same number of ancilla photons as we start with. Therefore, we ask what the output state of the top mode is conditional on detecting one (and only one) photon in the lower mode as shown in figure 2.5. Using the beam splitter transformations we can show that this post-selected gate implements,

$$\begin{aligned} |\psi_{\text{in}}\rangle &= \alpha\,|0\rangle + \beta\,|1\rangle + \gamma\,|2\rangle \qquad (2.27)\\ \longrightarrow |\psi_{\text{out}}\rangle &= \sqrt{\eta_1}\alpha\,|0\rangle + (1-2\eta_1)\beta\,|1\rangle - \sqrt{\eta_1}(2-3\eta_1)\gamma\,|2\rangle\,. \end{aligned}$$

Although we see that for $\eta_1 < 1/2$ this correctly adds a phase flip to the "2" component, unfortunately there is no solution for η_1 such that the "0", "1" and "2" components have equal amplitudes, leading to a biased gate.

This can be solved by using another beam splitter of reflectivity η_2, and heralding on vacuum at its output, as shown in figure 2.6(a), which leaves us with,

$$|\psi_{\text{out}}\rangle = \sqrt{\eta_1}\alpha\,|0\rangle + \sqrt{\eta_2}(1-2\eta_1)\beta\,|1\rangle - \eta_2\sqrt{\eta_1}(2-3\eta_1)\gamma\,|2\rangle\,. \tag{2.28}$$

This allows us to re-balance the gate using, $\eta_1 = (3-\sqrt{2})/7$ and $\eta_2 = 5-3\sqrt{2}$. The state $|\psi_{\text{out}}\rangle$ as defined above is not properly normalised which reflects the fact that the gate is probabilistic and only succeeds with probability, $\eta_1 \sim 0.23$. However, by detecting a single photon at the output of the first beam splitter, and vacuum in the second, we know for certain that the gate has succeeded. Unfortunately this success probability cannot be much improved as the theoretical upper bound is 0.25 [66,67]. An alternative NS-gate which saturates this bound was given by KLM and is shown in figure 2.6(b). Since we require two of these gates to both herald a success event for the CZ-gate to work the overall success probability is at best $P = 1/16$.

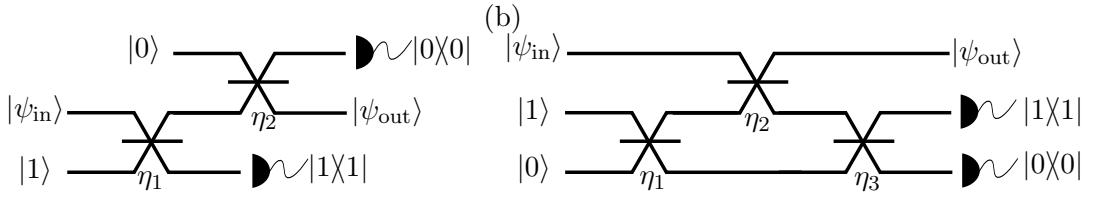


Figure 2.6: Heralded non-linear sign-shift gates. (a) The addition of a second beam splitter and vacuum heralding is sufficient to turn the Wang-Kobayashi interferometer shown in figure 2.5 into a NS-gate with success probability $P \sim 0.23$. (b) The original NS-gate as proposed by KLM saturates the maximum possible success probability of $P = 0.25$ but at the expense of requiring a third beam splitter and interferometric stability.

An alternative heralded two-qubit gate was proposed by Pittman, Jacobs and Franson

soon after in 2001 [68]. Their scheme requires an entangled ancilla state and succeeds with probability 1/4. Of course, generating this entangled ancilla cannot be performed deterministically with linear optics. The best known protocol for heralding an entangled Bell pair from four separable photons succeeds with probability 3/16 [69] thus reducing the overall success probability to just below 1/16.

These heralded gates on their own still do not allow us to build an efficient quantum computer since the probability of the computation succeeding becomes exponentially small as the number of two-qubit gates increases. In the next section, we show how KLM used quantum teleportation to increase this gate success probability close to 1 with only a polynomial increase in resources. Before we move on, however, we will discuss one more implementation of the CZ-gate which will be useful in understanding the teleportation experiment in chapter 4.

Ralph-Langford-White CZ-gate

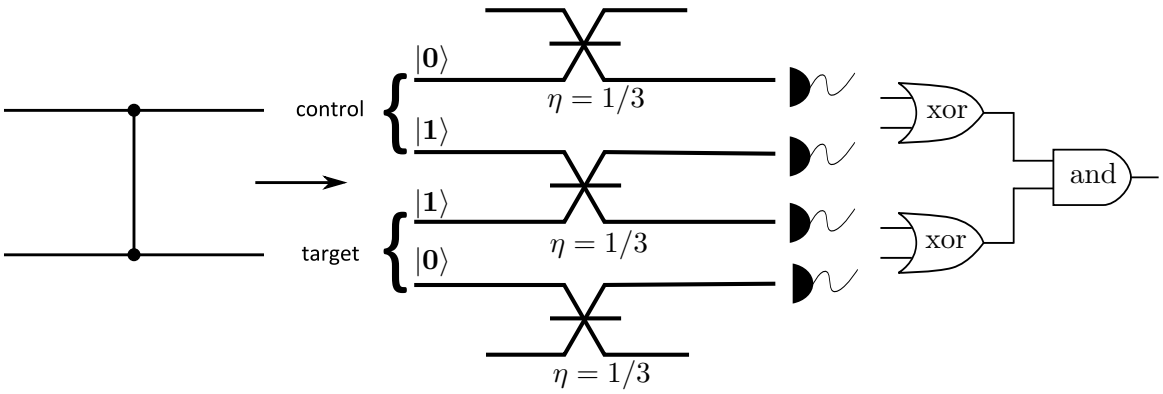


Figure 2.7: A post-selected CZ-gate operating in the coincidence basis. All three beam splitters are identical and have reflectivity $\eta = 1/3$. Successful operation is post-selected on detecting one photon in each qubit — i.e. one photon in either mode 2 or 3, and one photon in either mode 4 or 5. This succeeds with probability $P = 1/9$. Non-classical interference at the central beam splitter is responsible for the conditional phase flip whilst the outer beam splitters are required to correctly balance the probability amplitudes.

The addition of two single photon ancilla, coupled with number-resolved and vacuum detection make the experimental implementation of the above CZ-gate challenging. Instead of heralding, one can sometimes work in the 'coincidence-basis', using the photon qubits as their own ancilla such that the success of the gate is determined by measuring the very photons on which that gate acts. Obviously, such a protocol is not scalable but it provides a very convenient method of implementing a few parallel two-qubit operations with fewer resources than demanded by their heralded counterparts discussed above. The optical circuit used to implement this scheme is shown in figure 2.7. It consists of three identical beam splitters of reflectivity $\eta = 1/3$, thus defining 6 optical modes. The dual-rail logic used to encode the two qubits is defined by the central four modes whilst the other two outside modes have vacuum input and act as ancillary loss modes to balance the probability amplitudes.

By post-selecting on always detecting one photon in each of the two qubits it is easy to see using only single photon propagation that the three logical states, $|\mathbf{00}\rangle_{c,t}$, $|\mathbf{01}\rangle_{c,t}$ and $|\mathbf{10}\rangle_{c,t}$ emerge on the output modes unchanged — albeit with only a probability of 1/9. The phase flip of the $|\mathbf{11}\rangle_{c,t}$ state again occurs when two photons are incident on the central beam splitter. In this case, the photons undergo partial HOM-interference leading to the state, $|1,1\rangle \rightarrow \frac{1}{3}(\sqrt{2}\,|2,0\rangle - \sqrt{2}\,|0,2\rangle - |1,1\rangle)$. By post-selecting on detecting one photon in each qubit we only retain the final term and thus project onto the phase-flipped state as required. Therefore, with this design, we can implement the CZ-gate using only two photons, with a success probability of 1/9 but we must necessarily destroy the qubits in the process.

An important extension of this scheme was derived by Ralph[70] who showed that two of these gates can be operated in parallel by sharing each others' ancillary loss modes and therefore increasing the total success probability of both gates working from $P = (1/9)^n \rightarrow$

$P = (1/3)^{n+1}$. A sketch of this scheme is shown in figure 2.8 and we will make use of this design later in the construction of the on-chip teleportation circuit in chapter 4.

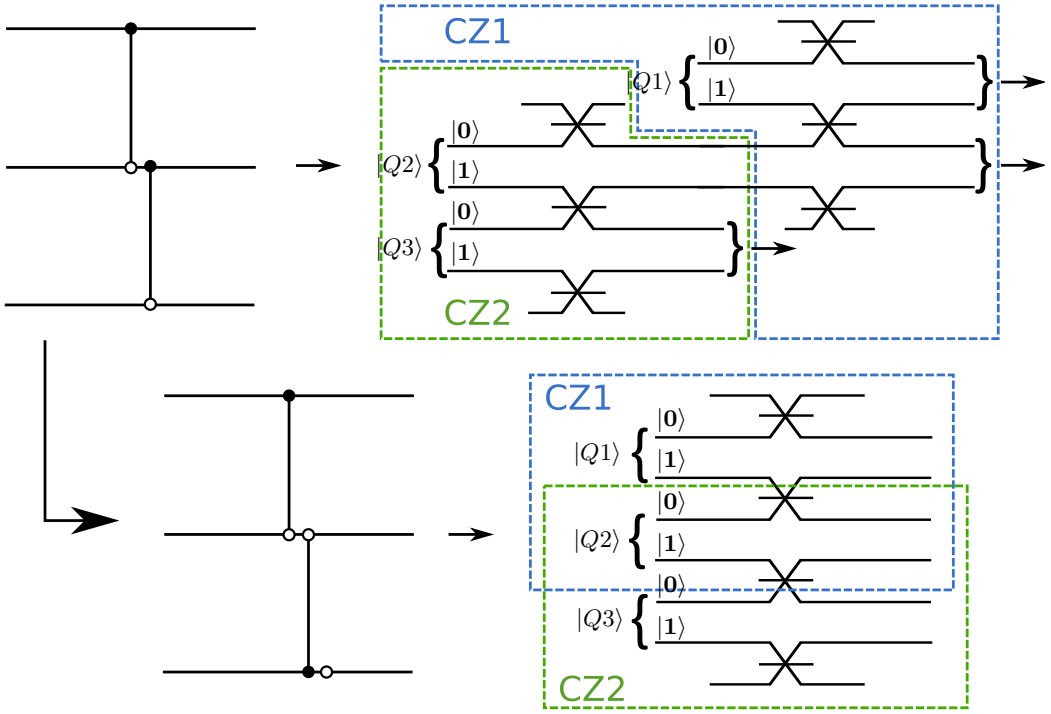


Figure 2.8: Optical circuit for two cascaded Ralph-Langford-White CZ-gates. Instead of requiring two copies of the circuit depicted in figure 2.7, the ancillary loss modes can be re-used, which increases the success probability from $P = (1/9)^2 = 1/81$ to $P = 1/27$. Note, that in this case we have re-defined the dual-rail logical encoding from before which leads to a slightly modified circuit operation, implementing a sign-flip on the logical zero state - denoted in the computational circuit on the left with an open circle.

2.2.3 Logic gates by teleportation

In 1999, Gottesman and Chuang outlined a method for moving these probabilistic gates from the quantum circuit into the state-preparation stage. Known as the 'teleportation trick' it relies on the quantum teleportation protocol. This is discussed at length in chapter 4, but we briefly summarise it now. Quantum teleportation is a process of transferring quantum information between two parties using shared quantum entanglement and classical feed forward without transferring any physical state[71]. The shared entanglement most often

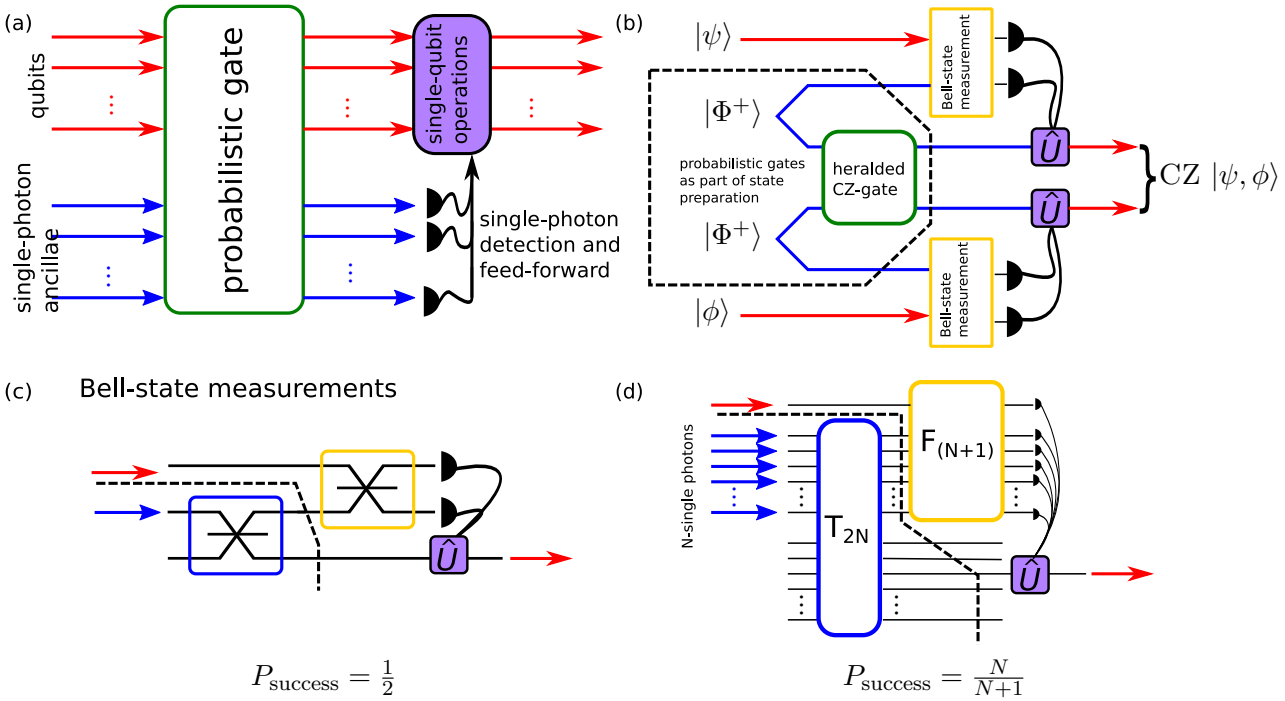


Figure 2.9: KLM scheme for efficient quantum computation using linear optics. (a) KLM showed that by combining a register of qubits with single-photon ancilla states, detection and feed-forward, efficient LOQC is possible. (b) The 'teleportation-trick'. Using two ancilla entangled qubits,$|\Phi^+\rangle$,one can probabilistically apply a two-qubit gate to one mode of each Bell-pair. Once the gate succeeds, a joint Bell-state measurement (BSM) is made between each qubit and the other half of each Bell pair. This teleports the action of the two-qubit gate on to the two qubits. In this way, all the probabilistic elements of linear optical two-qubit gates are reduced to a state-preparation problem, thereby allowing repeat-until-success strategies to be used without destroying the valuable qubits (red) from earlier computation. (c) In the simplest linear-optical teleportation scheme, modal entanglement is generated by sending a single photon ancilla through a 50:50 beam splitter. A partial BSM is made between one half of this pair and the qubit using another 50:50 beam splitter. This gate only succeeds with probability $P_{\text{success}} = 1/2$. (d) To increase the success of the BSM, we must prepare a more general entangled ancilla state of N-photons distributed over $2N$-modes. This can be achieved probabilistically using a quantum circuit built from heralded two-qubit gates and single qubit rotations as outlined in the text. The Bell-state measurement is then implemented using the $(N+1)$-point discrete Fourier transform, $F_{(N+1)}$ between the first N-modes of the entangled ancilla and one-mode of the dual-rail qubit. This can be implemented deterministically using an $(N+1)$-port linear optical interferometer. Detecting k photons at the output of this BSM teleports the qubit to the $(N+k)^{\text{th}}$-mode with probability $P_{\text{success}} = N/(N+1)$. In all cases, single qubit corrective unitary operations, $\hat{U}$, must be performed on the output qubits dependent on the exact outcome of the BSM.

takes the form of a maximally entangled Bell pair ($\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$). The protocol proceeds by making a joint measurement (a Bell measurement) between one half of this entangled pair and the state we wish to transfer. This projects the remaining half of the Bell pair onto the original state modulo one of the Pauli rotations. The specific rotation is revealed by the joint measurement outcome and can be communicated classically to the other party, thus completing the protocol.

The 'teleportation-trick' of Gottesman and Chuang now proceeds like this: beginning with two entangled Bell-pairs, $|\Phi^+\rangle$, one of the probabilistic, but heralded, two-qubit CZ-gates can be implemented between one mode of each Bell pair. The preparation of the Bell-pairs and the execution of the CZ-gate can be performed using a repeat-until-success strategy without worrying about destroying any valuable quantum information from earlier computations. Once this preparation-stage has succeeded, a Bell-state measurement (BSM) can be made between the other modes of the Bell-pairs and one mode of each qubit (see figure 2.9(b)). Because the CZ-gate commutes with the single-qubit Pauli rotations applied at the end of the teleportation protocol, the resulting qubits which emerge have effectively had a CZ-gate deterministically implemented between them.

This scheme is closely related to the Pittman et al. gate introduced above. In their scheme and in the 'teleportation-trick', entangled Bell pairs plus conditional measurements are used to make a probabilistic, but heralded two-qubit gate.

However, a perfect, deterministic Bell state measurement using only linear optics has been shown to be impossible[72,73], as is Bell-state creation. The simplest BSM using linear optics is shown in figure 2.9(c). In this scheme we teleport only one mode of the qubit $|\psi\rangle = c_0 |1,0\rangle + c_1 |0,1\rangle$, which in general is a superposition of zero and one photons, $|\psi\rangle_0 =$

$\alpha|0\rangle + \beta|1\rangle$. The entangled ancilla in this case can be prepared using a single photon incident on a beam splitter, $|\Phi^+\rangle = |1,0\rangle + |0,1\rangle$. A partial BSM can then be implemented using a balanced beam splitter between the qubit mode and one mode of this entangled state. Mode teleportation is achieved when we detect one photon in either arm of this beam splitter which occurs only half the time. Since two qubit modes need to be teleported in this way the overall success probability of our CZ-gate would only be 1/4 (c.f. the Pittman et al. gate above). This would appear to make this teleportation scheme useless for solving our original problem of having only probabilistic two-qubit gates.

KLM showed that if we allow for an arbitrarily small imperfection in this teleportation protocol, then we can increase the success probability arbitrarily close to 1 using only linear optical elements. We first prepare a large number, N, ancillary photons in the entangled state, $|t_N\rangle$ defined by,

$$|t_N\rangle = \sum_{j=1}^{N} |\mathbf{0}\rangle^j |\mathbf{1}\rangle^{N-j} \tag{2.29}$$

where $|\boldsymbol{k}\rangle$ are logical qubit states and $|\boldsymbol{k}\rangle^j$ denotes $|\boldsymbol{k}\rangle_1 |\boldsymbol{k}\rangle_2 , ..., |\boldsymbol{k}\rangle_j$. Because this state is part of the preparation stage, we can prepare it using the probabilistic two-qubit gates discussed above, only introducing it to the computation when we have heralded its successful creation. The teleportation then proceeds by sending one of the qubit modes and the first N-modes of the state $|t_N\rangle$ into a balanced $(N+1)$-port interferometer which erases all which-path information of the input photons[2]. All $N+1$ output modes are measured in the number-basis and the total number of photons detected, k, is recorded. The input mode is then teleported to mode $N+k$ of the entangled ancilla state modulo some unitary rotation, identified by the specific configuration of output photons detected. In this manner, the

[2]Such a balanced interferometer actually implements a discrete Fourier transform on these modes and as such is often referred to as a Fourier multi-port.

success probability of the BSM can be increased to $N/(N+1)$. In this scheme, we must now apply N successful CZ-gates between modes $N+m$ and $2N+m$ for all $m = 1, ..., N$ of the entangled ancilla. This results in a CZ-gate being implemented between the two qubits with a probability of $N^2/(N+1)^2$. With a sufficient number of ancilla photons this probability can be made smaller than fault-tolerant thresholds and it is thus possible to implement a scalable quantum computer using single photons, linear optics, photon detection and feed-forward.

Of course, this scheme can still be undone by photon loss and the (now small) probability of teleportation failures. Like any other quantum computation scheme, it must be possible to encode against these errors in a fault-tolerant fashion. We will not go into any details of the proof here since they are not relevant for the work carried out in this thesis. We have covered the major points of interest already, namely the requirements for high-visibility quantum interference between photons, construction of large optical interferometers and high-efficiency number-discriminating detectors. KLM showed a fault-tolerant error-correction scheme can be implemented in the near-deterministic teleportation architecture detailed above. Of particular interest, they showed that using error correction it is possible to recover from failed teleportation attempts and thus significantly reduce the size of the entangled ancilla state, $|t_N\rangle$ to achieve a given success probability. They were able to prove that such error-correction could be implemented efficiently with additional linear optics and single photon ancilla states and thus complete the first proof showing LOQC was, in principle, possible.

The physical resources for this KLM proposal, although formally scalable, are still incredibly daunting. For example, to achieve a CZ success probability of 95% requires of order

10,000 linear optical elements for each CZ-gate one wishes to implement! The reason LOQC continues to attract significant attention is the development of more efficient quantum gates. We will briefly review the progress made so far in reducing the resources required for LOQC and find that, although the resource requirements are much reduced, the basic experimental building blocks remain the same.

2.3 Beyond KLM

The defining breakthrough of the original KLM scheme was realising that the probabilistic two-qubit gates could be performed on 'cheap' ancillary photons off-line before teleporting them onto the 'valuable' qubits carrying the quantum information, reducing it to a state-preparation problem. The major advances in LOQC since the KLM scheme have focused on taking this paradigm further and reducing the entire quantum computation to a state-preparation problem — leaving only single qubit measurements and feed-forward to define the actual computation carried out. This was first formalised by Raussendorf and Briegel in what they called one-way quantum computing (OWQC) [74].

2.3.1 One-way quantum computing

It has been known since the seminal work of Deutsch in 1985 that any quantum computer can be implemented using arbitrary single qubit rotations and two-qubit gates. In the 'circuit-model' approach, quantum information is encoded in qubits which undergo a sequence of single and two-qubit operations (see figure 2.10(a)). Using linear optics the single qubit operations are implemented deterministically with beam splitters and phase shifters whilst each two-qubit operation involves a complex and resource-intensive procedure like the one

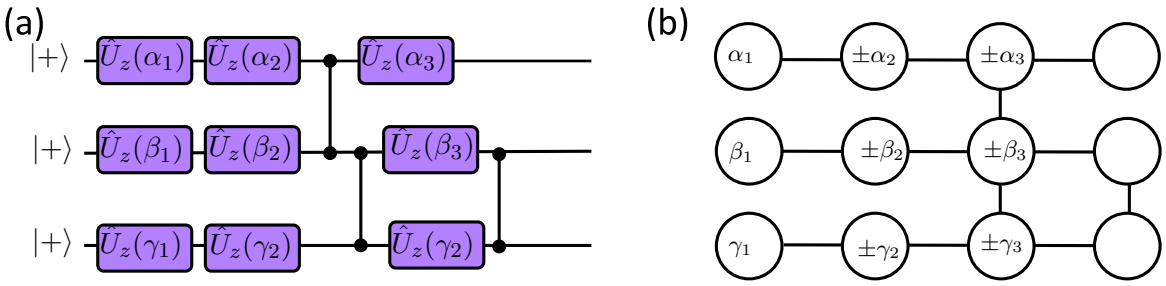


Figure 2.10: Cluster-state model of computation. (a) Any quantum circuit can be built up using single qubit operations (purple boxes) and two-qubit CZ-gates. (b) In the one-way computing model, a highly entangled cluster state acts as a universal resource. Each circular node in the cluster represents a single qubit prepared in the superposition state, $|+\rangle$, and connecting vertexes represent CZ-gates. Each horizontal row in the cluster represents one logical qubit and the computation proceeds from left-to-right by making measurements on each of the nodes in turn. Figure 2.11 explains how one can feed-froward these measurement results to enact arbitrary single qubit operations, thereby enacting the equivalent circuit in (a).

outlined by KLM.

A new model of quantum computation was introduced by Rausendorf and Briegel in 2001 called one-way quantum computing (OWQC)[74]. In this model, the entire computational resource state is created prior to any computations taking place. This resource is termed a cluster state and consists of single qubits prepared in the $|+\rangle$ state connected by CZ-gates as shown in figure 2.10(b). It is simplest to see how this implements a quantum computer by first considering how single-qubit operations are performed. Figure 2.11 explains how making only single qubit measurements on a linear cluster teleports the information contained in the measurements to the remaining, un-measured qubits, effectively implementing arbitrary rotations.

The mapping between the circuit-model and cluster-states can now be seen in figure 2.10. Each row of the cluster state represents one qubit in the circuit model. Measurements are made left-to-right with specific settings depending on previous measurement outcomes to implement the desired single qubit unitaries. Like the circuit model, these qubits are joined by vertical connections implementing the two-qubit CZ-gates. Thus it can be seen

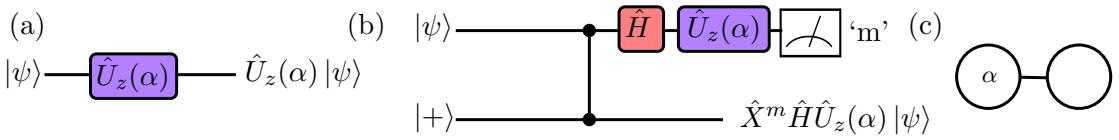


Figure 2.11: Single qubit rotations using cluster states. (a) A single qubit rotation around the z-axis can be performed on a remote qubit using the equivalent teleportation circuit in (b). The measurement basis is modified by the appropriate z-rotation to project the lower qubit into the desired state, modulo a Pauli rotation depending on the measurement outcome. Only z-rotations can be performed in this way since they commute with the CZ-gate. Arbitrary rotations around any axis can be performed by concatenating these z rotations three times as shown in. The measurement basis of each subsequent qubit must be modified depending on the outcome of previous measurement in order to correct for Pauli bit flips. This determines a temporal ordering for the measurements and is why this model is referred to as 'one-way' quantum computing. (c) The teleportation circuit can be conveniently represented using cluster states. Each circular node represents a qubit prepared in the $|+\rangle$ state with horizontal lines denoting a CZ-gate. This can be considered a universal resource state and single qubit rotations can now be applied to the right-most qubit by making appropriate measurements on the left-most qubit. Quantum information travels through the cluster state via teleportation and feed-forward corrective operations in a similar fashion to the original circuit-based KLM scheme.

that two-dimensional cluster states can be a resource for quantum computation. In fact, any regular two-dimensional clusters are *universal* resources for quantum computation and any three-dimensional regular cluster is universal for realising error-correction and fault-tolerant quantum computing. Therefore, once the cluster is constructed only single-qubit measurements and feed-forward is required to implement the entire fault-tolerant quantum computation - things which are 'easy' to implement using linear optics. Hence most of the challenge in building an optical quantum computer is reduced to being able to build these highly-entangled cluster states.

This has a number of advantages over the circuit-based model. Since any 2D regular cluster is universal for quantum computing, efforts can concentrate on being able to build this one resource state. As all of the two-qubit gates are implemented during state preparation, the tolerance to errors is much greater. In particular, the cluster-state model has been found to be extremely resilient to photon loss, requiring only that the product of detector and of

photon source efficiency be greater than $2/3$ [75].

The cluster-state model is also more resource efficient. For comparison, the original KLM scheme would require two sets of highly entangled 100-photon $|t_{100}\rangle$ ancilla states to perform a single CZ-gate with 98% probability [76]. Using the scheme proposed by Browne and Rudolph [69] to generate the equivalent required cluster state, requires on average only 52 2-photon Bell states [76], a significantly less demanding task.

2.4 Summary

In this chapter we have seen how the lack of a direct photon-photon interaction can be overcome in creating two-photon logic gates using only linear optical elements without requiring any matter-mediated non-linearities. At the heart of these linear-optical quantum computing schemes is the phenomenon of correlation-without-interaction arising from the enforced symmetry of ensembles of identical particles, exemplified in the Hong-Ou-Mandel effect.

Combining this effect with ancilla photons and post-selection enables us to induce effective non-linear interactions between photons. The great breakthrough of Knill, Laflamme and Milburn was to show that such probabilistic non-linear gates could be performed off-line and efficiently teleported into a quantum circuit. Although later improved upon by the introduction of one-way quantum computing, the essential ingredients of an optical quantum computer remain the same [77] — identical photons propagating through linear interferometers coupled with measurement and feed-forward to teleport quantum information through the circuit or cluster.

These schemes provide us with a prescription for building a quantum computer from

linear optical elements. In particular, we require a large number of single photon sources all producing indistinguishable pure photons such that we can achieve the high-visibility single-photon ('classical') and two-photon ('quantum') interference necessary to generate correlated photon states through HOM-type interference. Implementing arbitrary single qubit rotations requires tunable beam splitters and phase shifters and the two-qubit gates of KLM require large optical interferometers consisting of many phase-stable interfering paths all of which need to be well mode-matched.

Finally, we have seen how photo-detection is a key requirement in all of these protocols. In a sense, it is the highly non-linear act of detection (detectors either click or not) which allows us to induce the effective optical non-linearities of the post-selected states. Crucially, all of the photons used in the computer - including all of the ancillary ones - need to be detected at some point. In reality, all detectors work with some finite efficiency such that the probability of all detectors working at once decreases exponentially, rapidly degrading the fidelity of the gate. Error-correcting codes and the cluster-state model alleviate some of the requirements but the demands on efficiency are still daunting. Apart from a recent proposal[78] for generating initial cluster states with non-number-resolving detectors [3] all known LOQC schemes also rely on conditioning output modes to contain a certain number of photons. Thus detectors which are able to resolve photon-number with high efficiency are likely to be required.

[3] Although note that this scheme relies on an available on-demand source of entangled Bell pairs free of higher-order contamination - i.e. a source which is not spontaneous parametric down-conversion.

The state of the art in LOQC

Many initial experiments in quantum information were first carried out using photons. Photonic entanglement was first generated via atomic cascades [79,80], and later via parametric down-conversion [81–83] and used to show the first violation of Bell's inequality [79–85]. Soon after these fundamental tests of quantum mechanics, the first experimental demonstration of quantum teleportation [71] was made using photons [30,86].

Even before the scheme of KLM, photonic systems were used as a test-bed for realising quantum logic gates. Cerf et al. proposed a scheme for simulating quantum logic using only a single photon and linear optics [58]. This scheme required 2^N modes to simulate N qubits making it manifestly unscalable. A similarly non-scalable approach was used to optically simulate Grover's search algorithm [59].

Soon after the KLM scheme appeared, demonstrations of post-selected [19,87–90] and heralded two [91,92] and three [93] qubit gates were reported. A variety of information encodings have been explored, starting with polarisation [19] and spatial modes [13] up to using time [22,94–98] and frequency encodings [21,99] and, more recently, orbital angular momentum [23,24,100–102]. Many of the basic necessary protocols have also been demonstrated, including the heralded CNOT-gate [103], gate-teleportation [104], active feed-forward [105] and quantum error-correction [106].

There have also been demonstrations of cluster-state generation [107–113], with up to 8 photon entangled clusters being shown in post-selection [114,115]. Experiments using cluster states have also implemented various algorithms such as solving systems of linear equations [116,117] blind-quantum computing [118] and verification protocols [119]. The loss-tolerance and error-correction abilities of photonic cluster states have also been demon-

strated in experiments[114,120].

It is clear that quantum photonic systems have been remarkably successful in providing many of the first proof-of-principle experiments in quantum information. However, if it is to continue to compete with developments being made by ion trap, or superconducting-based qubits then there has to be a plausible route for scaling these photonic system up to larger numbers of qubits — more photons propagating in more modes. The above experiments were mostly performed using free-space propagation and bulk optical elements. Whilst a significant number of experiments have been performed in this way, it is not a practical approach for scalable implementations of more powerful quantum computations. Bulk optics take up significant room and require careful alignment, mode-matching and sub-wavelength path stability. This is a problem which gets much harder to overcome as the number of photons and optical modes increases further. Although fiber-optic based systems can overcome some of these issues, the typically longer optical path lengths increase the effects of dispersion, polarisation dephasing and path-length instability. Instead, an approach based on *integrated* optics mostly removes alignment issues and offers an inherently stable architecture to build complex interferometric networks. In addition, being able to define a single optical spatial mode using waveguides enables excellent mode-matching and thus can provide exceptionally high quantum interference visibilities as required for LOQC.

The first quantum integrated photonic logic gates were demonstrated in 2008 using etched silica-on-silica waveguides to realise a post-selected two-qubit gate[20]. The use of thermo-optic phase controllers was first shown in controlling the effective reflectivity of an integrated beam splitter[121] and has since been used on a reconfigurable circuit to arbitrarily manipulate two-qubit entangled states[122]. Since then a wide variety of dif-

ferent waveguide architectures have been explored, from silica-based platforms including etched[20,33,34,122–126], UV-direct written[36,121,127] and femtosecond-laser written fabrication techniques[128–140]. High index contrast materials such as silicon[141–143] or GaAs[144] have also been used and are promising for achieving high component density but currently suffer from large propagation and coupling losses. Non-linear materials such as lithium niobate have also been used and show promise in integrating both photon sources, circuits and high-speed modulators in one system[145]. There have been many implementations of two-qubit algorithms using integrated optical circuits from a compiled version of Shor's factoring algorithm[126], quantum random walks[129,132,136,143], quantum metrology[146], delayed choice experiments[147] and variational eigenvalue solvers[148].

The work described in this thesis was part of a series of recent experiments which moved beyond these two-photon integrated experiments and made the jump to *three* input photons. This has enabled the exploration of more interesting multi-particle quantum interference[36,131,149] and algorithms of emerging importance such as boson sampling[127,133,150–153].

Chapter 3

Three Photon Quantum Interference

Integrated photonics has emerged as a leading platform for realising linear optical networks of sufficient complexity to reach a regime beyond what can be classically simulated. Since the first demonstration of this technique in 2008, pioneering work has shown the viability of this approach for the manipulation of the quantum properties of photons[20,35,121,122,126,128,137,138,143,154,155]. However, these experiments have been mainly limited to using only two photonic inputs. In the next two chapters of this thesis we will demonstrate advances in the complexity of integrated photonic circuits used for LOQC by simultaneously increasing the number of modes and interference nodes in the circuit and the number of photons distributed amongst them. In this chapter we demonstrate the first three-photon quantum operation of an integrated device. This device contains three coupled interferometers, eight spatial modes and many classical and non classical interference nodes. This already represents an advance over the previous state-of-the-art and is the first on-chip non classical interference with more than two photonic inputs. This chapter ends with an application of this three-photon interference in proving the central premise of a new

limited-form of quantum computation known as Boson Sampling. Chapter 4 then builds on the progress made in this chapter to realise the first on-chip three-qubit logic gate and presents a design for a future device which will enable even more complex experiments.

We start by discussing the crucial experimental advances which were necessary for this demonstration: optimisation of a low-loss single-photon source, use of a low-loss waveguide architecture and a careful circuit characterisation scheme. We then show how one can verify quantum behaviour using classically characterised device elements and hierarchies of photon correlation functions and accurately predict the device's quantum behaviour and show operation inconsistent with both classical and bi-separable quantum models. We end with a brief outline of Boson Sampling and present results showing the first demonstration of this protocol using the chip presented in this chapter.

3.1 Multi-photon states on-chip

There are two key outstanding challenges associated with scaling integrated photonic circuits to larger sizes. First, photon loss exponentially limits the complexity achievable in a quantum circuit, both in terms of the number of circuit elements and the number of photons that can be used effectively. Second, reliable fabrication of precise integrated optical components is difficult. Moreover, the monolithic nature of integrated architectures means it is also difficult to even verify that the chip meets the design specification. In this section, we will concentrate on the former, and discuss the strategies used to mitigate photon loss in both the design of the photon source and the choice of waveguide architecture. Later, in section 3.2, we will return to the issue of verification of expected quantum behaviour within a circuit. Finally, in the next chapter, we consider the effect of device deviations from design

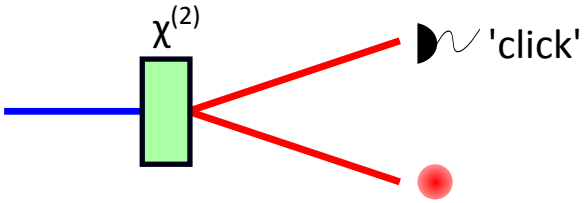


Figure 3.1: Heralded single photon source. A pump photon (blue) is spontaneously converted into two daughter photons (red) with low probability in a $\chi^{(2)}$ non-linear crystal. Although the conversion is probabilistic, placing a detector on one arm *heralds* the presence of a photon in the other, allowing discrimination of multi-pair events when more than one source is used. Ensuring the photon is heralded into a single mode in a pure quantum state with high efficiency constitutes the bulk of source engineering work and is discussed in more detail in the text.

and develop strategies to cope with these imperfections.

3.1.1 Single-photon source

Despite the rapid progress in experimental quantum optics over the past decade, generating a single excitation of single optical mode (a Fock state), remains a key challenge[17]. Ideally, one requires a source which emits, on-demand, a single photon in a pure quantum state. Additionally, one would like this photon in a mode that can be efficiently coupled to a photonic circuit. Assuming one can always post-select a pure quantum state using appropriate filters, the key metric for any photon source is the *heralding efficiency* — the conditional probability that you obtain a photon in the desired mode given a heralding signal,

$$\eta_{\mathrm{h}} = \frac{R_{\mathrm{photons}}}{R_{\mathrm{herald}}}, \tag{3.1}$$

where R_{photons} is the rate of single photons emitted into the desired mode coincident with the herald signal and R_{herald} is the rate of events which signal that a photon *should* exist in that mode. In effect, the heralding efficiency defines the loss experienced by a photon in the heralded mode and limits the overall N-photon data-rate to, $R_{\text{data-rate}} \propto \eta_{\mathrm{h}}^{N}$.

Isolated quantum matter systems such as quantum dots[156–158], single trapped atoms[159–163], and nitrogen vacancy (NV) centres in diamond[158,164–167] have shown some promise at producing truly single photons 'on-demand'. However, there remain significant challenges in engineering such sources to emit into a single optical mode which in practise means heavy filtering is needed to generate a pure quantum state. The loss induced by this filtering renders such sources increasingly probabilistic, decreasing their utility in many experiments.

Instead, the workhorse of quantum optics experiments remains the spontaneous generation of number-correlated pairs of photons through the parametric interaction of a strong pump laser with a non-linear medium (see figure 3.1). In this work we use spontaneous parametric down-conversion (SPDC) in a bulk crystal to produce two sets of heralded single photons (figure 3.2). In this scheme, although the photons are only generated probabilistically, we can use one of the photons to *herald* the presence of the other. This probabilistic emission jeopardises SPDC as a scalable system for many-photon generation. Nevertheless, the ability to generate pure, indistinguishable approximate single photons using a small collection of room-temperature optics makes it an incredibly convenient platform for small-scale demonstrations using a few photons[17].

Heralding pure single photons

The quantum state generated in a type-II SPDC process is a two-mode squeezed vacuum (TMSV),

$$|\psi_{\text{SPDC}}\rangle = \sqrt{1-\lambda^2}\sum_{n=0}^{\infty}\frac{\lambda^n}{n!}(\hat{A}^\dagger\hat{B}^\dagger)^n\,|\text{vac}\rangle \tag{3.2}$$

where λ parameterises the squeezing and $\hat{A}^\dagger$ and $\hat{B}^\dagger$ are the creation operators for the signal and idler photons into broadband wave-packets. By keeping the pump power low we can reduce λ to ensure the $n = 1$ term dominates as desired (the vacuum component is effectively eliminated via post-selection). Down conversion is governed by energy and momentum conservation between the pump and the daughter photons so that, in general, the properties of one photon are highly correlated with those of its twin. This often manifests itself as non-factorable spectral correlations between the signal and idler modes [168]. In this case, we can no longer write our state using only a single pair of broadband modes and instead we must write our state as a tensor product of orthogonal broadband modes,

$$|\psi_{\mathrm{SPDC}}\rangle = \bigotimes_{m=0}^{\infty} \sqrt{1-\lambda_m^2} \sum_{n=0}^{\infty} \frac{\lambda_m^n}{n!} (\hat{A}_m^\dagger \hat{B}_m^\dagger)^n \, |\mathrm{vac}\rangle \, . \tag{3.3}$$

Since our detectors are unable to resolve these broadband modes, detection of a herald photon will project the other into a partially mixed state, ruining any quantum interference effects we wish to observe. This is most often rectified by using tight filtering to select only one mode. This comes at the expense of source brightness and thus reduces the available experimental data rates. Instead, by carefully matching the pump bandwidth with the phase matching function, one can engineer the source to emit into only one broadband mode - negating the need for tight filtering [169].

Three-photon source

This spectrally factorable source design was demonstrated in our group several years ago, where an un-filtered Hong-Ou-Mandel dip visibility of over 95% was achieved between two independent heralded sources [170–172]. Ideally, we would add one more heralded source of

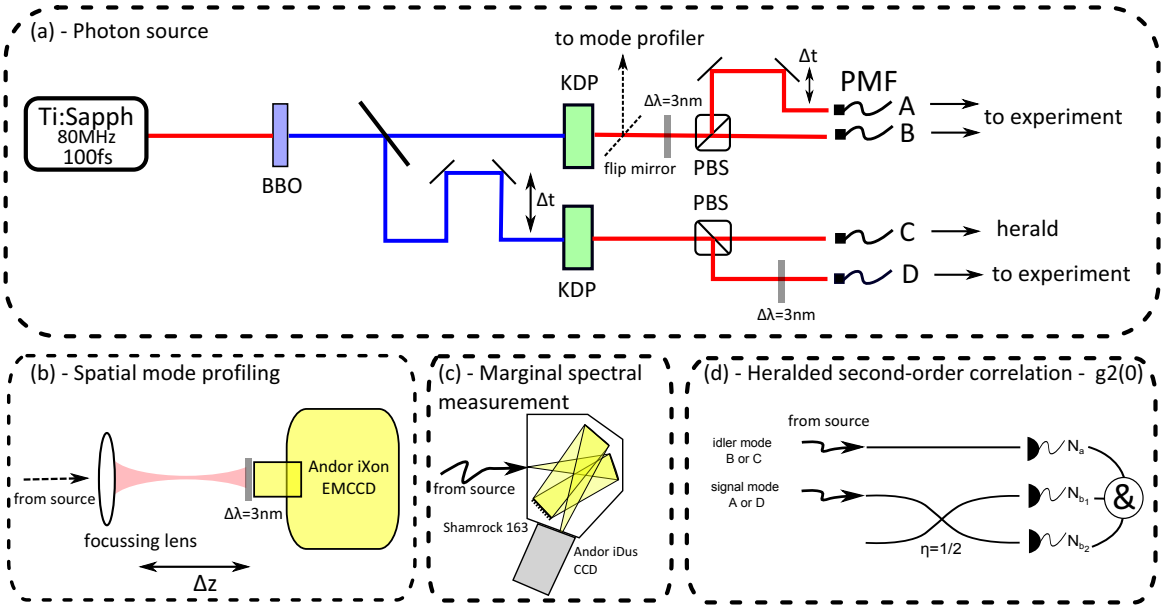


Figure 3.2: Single photon source. (a) Schematic layout of the SPDC source used for the generation of single photons in this thesis. An 80 MHz Ti:sapphire oscillator (Mai-Tai, Spectra Physics) produces 100 fs pulses at 830 nm (2.6 W cw equivalent average power) which are up-converted to 700 mW of 415 nm light in a 700 μm $\beta - BaB_2O_4$ (BBO) crystal cut for type-I second-harmonic generation. This blue light is split on a 50:50 beam splitter and used to pump two 8 mm-long AR-coated potassium dihydrogen phosphate (KDP) crystals phase-matched for degenerate type-II collinear parametric down-conversion. The signal and idler modes are split using polarising beam splitters (PBSs) and coupled into polarisation maintaining fiber (PMF). Two motorised translation stages are used to adjust the relative path delays between the two sources. Narrow-band interference filters (Semrock) ensure the photons are spectrally indistinguishable. We characterise the spatial mode (b), marginal spectra (c) and number statistics (d) of the source to optimise the heralding efficiency, spectral indistinguishability and higher-order down conversion terms respectively.

pure single photons to generate our three Fock states required for the experiments in this thesis. However, the probabilistic nature of these sources makes such six-fold coincident events extremely rare. Instead, because the source emits degenerate photon pairs, we chose to use one source in a heralded configuration and injected both modes of the other source into our circuit. We then always post-selected on detecting four-fold events as shown in the top panel of figure 3.2. Detecting four photons in this way enabled discrimination between down-conversion events with one photon in each spatial mode ($|111\rangle$) from equally likely unwanted noise events with two photons in each of only two input modes ($|022\rangle$ or $|200\rangle$).

The details of the source are given in figure 3.2.

The source uses two KDP crystals cut for degenerate, type-II, collinear parametric down-conversion. The crystal is cut to achieve asymmetric group velocity matching, resulting in the generated signal and idler photons having different spectral widths, the narrow-band ($\sim 3\,\text{nm}$) and broad-band ($\sim 12\,\text{nm}$) daughter photons[171]. We therefore must spectrally filter both signal and idler to restore spectral indistinguishability. Despite this filtering, the carefully-engineered spectral factorability of our source still has benefits. Specifically, we can choose a filter matched to the narrow-band photon width which allows heralding of that photons presence with high efficiency.

Optimisation of the source

The design of this source ensures emission into a well-defined spectral and polarisation mode. However, momentum conservation in free-space causes the signal and idler photons to emit into a cone-like spatial geometry as seen in figure 3.3. Mode-matching this spatial output into a standard single-mode fiber is critical for obtaining high heralding efficiencies. To achieve this, the down-conversion spatial mode was profiled using a single-photon sensitive electron-multiplying charge coupled device (EMCCD). By imaging the two down-conversion cones in different planes (see figure 3.2(b)) we applied the Gerchberg-Saxton phase-recovery algorithm[173] to recover the full complex electric field of the down-conversion at the point of emission.

We were then able to propagate this mode throughout the collection optics and identify the lenses which would best mode-match this to a single-mode fiber. By optimising the lenses according to this analysis, and replacing all silvered mirrors with dielectric stacks (Thorlabs

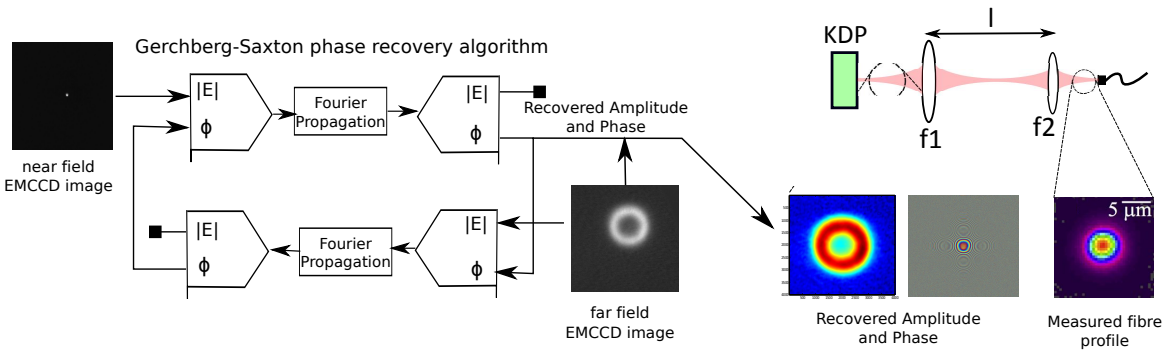


Figure 3.3: Gerchberg-Saxton phase recovery algorithm for efficient mode-matching of SPDC to single-mode fibre. (a) Using the apparatus shown in figure 3.2 we profile the down-conversion mode in the near-field and far-field. Since we know the propagation distance between these two profiles the iterative algorithm is able to reconstruct the full complex electric field of the down converted mode. (b) With this knowledge, we use two lenses (f_1 and f_2) to match this mode to that of a single-mode fiber (mode profile shown bottom right) with high efficiency.

EO3) we were able to increase the raw heralding efficiency from 15% to 32%. Dividing by the approximate detector efficiency at these wavelengths we arrive at a collection efficiency of 60-70% - in good agreement with the theoretical simulations shown in figure 3.4.

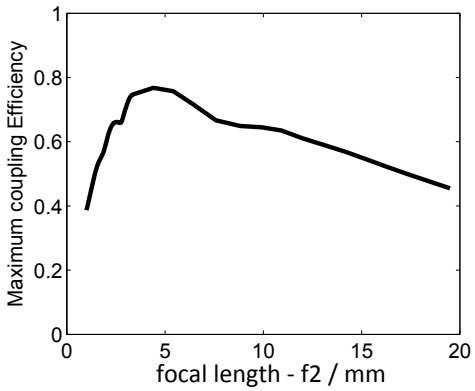


Figure 3.4: Predicted source coupling efficiency as a function of lens focal length. The result of the optimisation detailed in figure 3.3 suggests we should be able to achieve a coupling efficiency from source to single mode fibre of close to 80% in the abscence of any other losses.

The spectral indistinguishability was checked using a Czerny-Turner spectrometer (Andor Shamrock 163) fitted with a single-photon sensitive CCD (Andor iDus, figure 3.2(c)). Angle tuning the narrow-band interference filters allowed us to match the spectra of all three photons as shown in figure 3.5. Using a Hanbury-Brown Twiss interferometer (a 50:50 beam

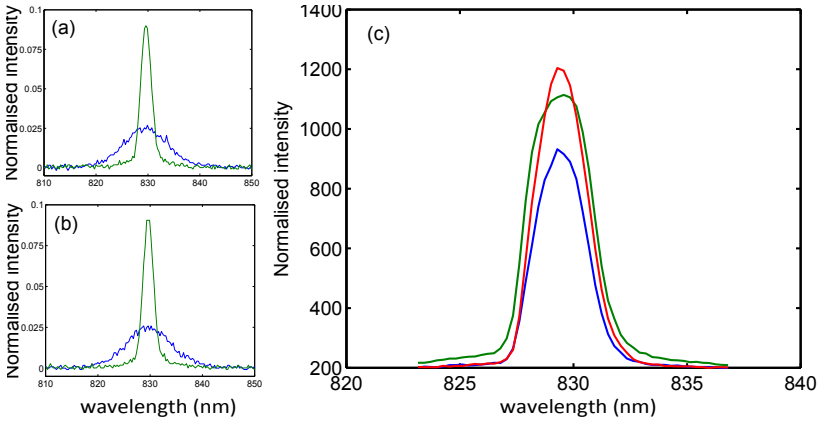


Figure 3.5: Spectral overlap of the three single photons used in this work. (a,b) The narrow (green) and broad (blue) marginal spectra for each of the two KDP down conversion sources. (c) Using two narrow-band spectral filters as shown in figure 3.2 we match the marginal spectra three of the photons (two narrow-band and one broadband). The remaining herald photon is left unfiltered to maximise count rates.

splitter with vacuum input on one mode - figure 3.2(d) we measured the second order correlation function of our source, $g_H^{(2)}(0) = N_{ab_1b_2}N_a/(N_{ab_1}N_{ab_2})$ [171]. Combining this with knowledge of our herald-arm efficiency, $\eta_H = N_{ab}/(N_{b1} + N_{b2})$, allows us to determine the squeezing parameter, λ^2 [139],

$$\lambda^2 = g_H^{(2)}(0)\frac{\eta_H}{2(1-(1-\eta_H)^2)}. \tag{3.4}$$

We determine that $\lambda^2 \approx 0.03$. This is small-enough to make sure the higher-order terms do not cause too much disruption[1] whilst ensuring we see enough down-conversion events to make the data-rate acceptable.

Finally, the indistinguishability of our source was checked by performing a Hong-Ou-Mandel dip through a 50:50 fiber beam splitter. We obtained a raw dip visibility of $92 \pm 1\%$ between the two heralded narrow-band photons (from fibers A and D - see figure 3.6) and a raw dip visibility between the heralded narrow-band photon from fiber A or D and the

[1] Although we shall see their effect is by no means negligible!

heralded broad-band photon from fiber B of $88 \pm 2\%$. However, these raw visibilities are affected by background noise from APD dark counts, higher-order photon pair emission and residual pump light leaking into the optical fibres. In order to accurately assess the impact of residual photon indistinguishability, we must subtract this background from our data. During data collection we periodically blocked one of the inputs to the fibre beam splitter using a computer controlled beam block. Without any background terms we should never record a two-photon coincidence term with one input blocked. The number of coincidence counts we record with one arm blocked provides an estimate of the background terms. Subtracting this background from our data gives a background-corrected dip visibility of $97 \pm 1\%$ between the two heralded narrow-band photons (see figure 3.6) and a background subtracted dip visibility between the heralded narrow-band photon from fiber A or D and the heralded broad-band photon from fiber B of $94 \pm 2\%$. These corrected visibilities are later used to make small first-order corrections to the expected visibilities of multi-photon interference through our waveguide device in section 3.4.

3.1.2 UV-written silica-on-silicon waveguide device

Having optimised the loss of the photon source our attention turns to the integrated circuit. In integrated photonics, significant losses arise from interfacing the circuit with both sources and detectors and become more pronounced with increased photon numbers[37,174]. Even with perfect interfacing, losses are fundamentally limited by the intrinsic optical properties of the medium which scale with the circuit size. Both of these issues can be addressed by choosing a suitable waveguide medium which exhibits low propagation loss and supports mode structures similar to those of the single-mode fibre used in both the photon sources

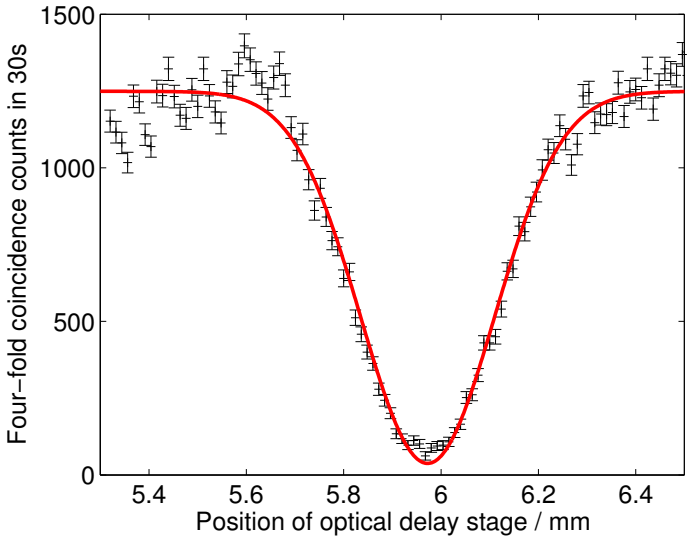


Figure 3.6: Heralded two-photon HOM-dip. we test the indistinguishability and purity of our two single photon sources by performing a heralded HOM-dip between the two narrow-band photons through a fibre-beam-splitter. The fitted visibility is $97 \pm 1\%$.

and detectors. Silica is one of the lowest loss dielectrics available and the telecommunications boom of the 1980's is responsible for making germanium-doped silica the de-facto standard for optical fiber construction[175,176]. A silica-based integrated architecture thus has many attractive properties for constructing low-loss waveguides with good mode-matching to optical fibre and consequently to almost all available photon sources and detectors.

Through a close collaboration with the Optoelectronics Research Centre at the University of Southampton we have been able to obtain silica-on-silicon waveguide circuits for use in our quantum experiments. The direct-UV writing approach they use allows for rapid prototyping of device designs at low cost. Because individual waveguide circuits can be fabricated relatively cheaply we have been able to ask for highly custom designs to be produced for a number of different experiments. This silica-on-silicon waveguide architecture detailed here is used in this thesis for all three of the experiments performed in the next two chapters and also for the on-chip detector experiments discussed later in chapter 6.

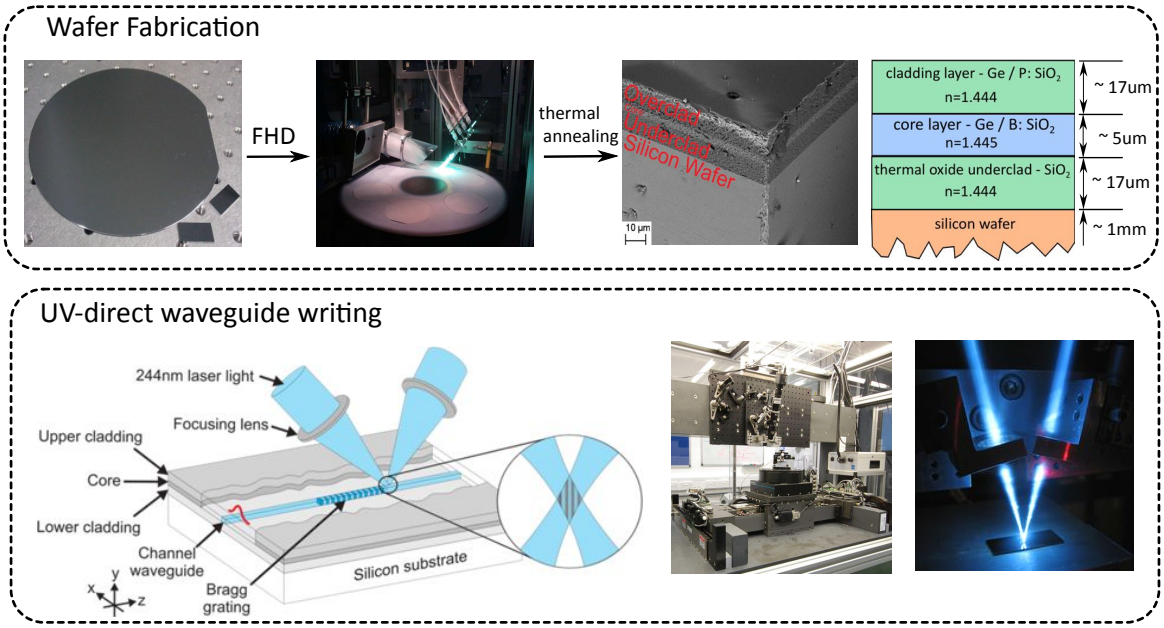


Figure 3.7: Waveguide Fabrication. Top - Wafer Fabrication. Commercial 6" silicon wafers are bought with a 17 μm thermal oxide layer already deposited. Flame-Hydrolysis Deposition (FHD) is then used to deposit the doped silica-based core and cladding layers. The green flame in the photo is due to the Boron dopants. After each FHD layer is applied the soot is consolidated into a glassy layer through a thermal annealing cycle. Germanium and Boron dopant's are added in carefully controlled amounts in the core layer to increase photosensitivity whilst ensuring the bulk refractive index closely matches that of the thermal oxide underclad. An SEM image showing the final 3-layer waveguide structure is shown in the grey-scale photo. The dimensions of the three layers are carefully controlled to provide single-mode guiding operation at 780 nm. Bottom - UV-direct waveguide writing. The wafer is diced into chips and polished before being Hydrogen-loaded over 1 week in a high-pressure Hydrogen cell to increase the photo sensitivity of the core layer. UV radiation at 244 nm is focused into the core layer to cause a permanent increase in the refractive index by $\sim 0.5\%$, thus defining a waveguide. The wafer is translated using 2D computer-controlled air-bearing stages (photo) to trace out the desired waveguide cirucit. Crossing two UV beams generates a controllable interference pattern which enables the creation of in-situ Bragg gratings along the waveguides. Photos are courtesy of the Optoelectronics Research Centre, Univerisity of Southampton.

Waveguide circuit fabrication

It is no surprise that with such excellent optical properties, waveguide fabrication in silica is a very mature technology. Silica samples are fabricated using techniques derived from the micro-electronics industry, resulting in relatively inexpensive production costs. Planar silica samples are most commonly fabricated on silicon host wafers - leading to the name silica-on-silicon (SiO_2:Si). Not only are high-quality silicon wafers readily available from

the integrated electronics industry but they also offer a very high degree of planarity, good adhesion of the deposited silica, excellent thermal dissipation and a usefully high melting point [177].

A three-layer silica-on-silicon sample design is utilised in this work as illustrated in figure 3.7. Two buffer-layers, the underclad and the overclad serve to isolate the core from the high refractive index of the silicon and the low refractive index of the air respectively. The underclad is created through thermal oxidation of the surface of the underlying silicon wafer. The wafers in this work were purchased from the Centre for Integrated Photonics (CIP) with a 17 μm thermal oxide already deposited.

The remaining two silica layers can be deposited with a variety of well-established planar deposition techniques. These are most commonly plasma enhanced chemical vapour deposition (PECVD) or flame hydrolysis deposition (FHD), which is used here [178]. In flame hydrolysis deposition, vapour precursors are introduced in a flame where they undergo oxidation to form soot particles which are then collected on a planar layer to form a low-density porous layer. The layer thickness is defined by the amount of time the wafer spends under the flame. This porous layer undergoes thermal annealing which consolidates it into a dense glassy layer. This process is repeated again to deposit the second layer. Different chemical dopants in the vapour control the refractive index, photosensitivity, thermal expansion and melting point of each layer. In the central core layer of our waveguides, germanium and boron are added to increase the photosensitivity of the glass whilst matching the refractive index of the underclad. The top cladding layer is doped with germanium and phosphorus to match the refractive index of the core whilst making it non-photosensitive and reducing the melting point to below that of the layers beneath it. This sequence is summarised in

figure 3.7 with a schematic on the top right showing the properties of the various layers involved. The wafer is then diced into appropriate sized chips and the end facets are lapped and polished with a colloidal silica solution to provide an optically flat surface suitable for fiber coupling. This process results in a small birefringence ($\sim 1 \times 10^{-4}$) in the core layer due to residual stresses [139]. This is sufficient to enable polarisation maintaining operation although this inevitably precludes its use in supporting polarisation encoded qubits [137].

UV-writing

Waveguide definition can now take place to create regions of localised increased refractive index, carefully engineered to support the guiding of only a single optical spatial mode. Although traditional fabrication methods involving photo-lithography and etching are well-suited to mass production devices, they lead to long and expensive production runs for one-off designs. Single mistakes in the mask design or any of the numerous processing steps have the potential to ruin the entire batch of devices making this an unsuitable technique for the small-scale rapid prototyping of custom designs we require for our work.

Direct UV-writing (DUW) is an alternative technique to define waveguides on photosensitive materials [179]. A UV laser of wavelength 244 nm is focused into the photosensitive core layer causing an increase in the refractive index within the high intensity focal spot of the beam. The germanium and boron dopants alone do not provide sufficient photosensitivity. Instead, the sample must be in-diffused with hydrogen in a high pressure (~ 100 bar) hydrogen-rich cell for 5 days prior to writing. The hydrogen diffuses out over a period of ~ 1 hour, limiting the available waveguide writing time. The sample is then translated transversely to the surface normal with computer-controlled 2D motion to trace out the

desired waveguide circuit. Waveguides of typical 4.5μm$\times 4.5\mu$m dimension are formed with an average refractive index contrast of $\sim 0.5\%$. This contrast is smaller than what can be achieved through traditional etching methods and leads to larger bend radii of the order of 1 cm. However, the DUW system does not require expensive cleanroom facilities for fabrication and whole waveguide circuits can be designed and written within an hour. Because the index contrast and core size is similar to that of commercial silica fiber the mode-matching between the two can be in excess of 98%[180].

Cross-couplers

Integrated beam splitter devices can be made by bringing together two guides sufficiently close so that their evanescent fields overlap. The beam splitter reflectivity is then determined by the distance between the waveguides and length of the interaction region. This 'directional coupler' has a beam splitting ratio which is sensitive to wavelength, temperature, polarisation and fabrication imperfections[181]. Instead, DUW allows the creation of compact x-couplers[182] whose beam splitting ratio is instead determined by the geometrical crossing angle of the waveguides. This makes the beam splitting ratio relatively insensitive to wavelength, polarisation and temperature although fabricating a pre-determined splitting ratio still requires accurate knowledge of the index contrast of the guide. In practise, a desired splitting ratio is achieved through trial and error. A crossing angle of roughly $\sim 2.45^\circ$ provides 50:50 coupling with a coupler length of ~ 4 mm and excess loss due to the waveguide bends of between $0.1 - 0.2$ dB[177].

Thermo-optic phase shifters

Tunable phase shifts are generated via the temperature-dependent refractive index of silica. These thermal phase shifters are fabricated by depositing nichrome heaters (0.35 μm $\times$ 50 μm $\times$ 2.5 mm, 0.85 kOhm electrical resistance) deposited directly above the optical waveguide modes. Once calibrated, a computer interfaced voltage controller can then adjust the phase to any desired value. The excellent thermal conduction of the silicon wafer ensures the heat is dissipated straight down allowing cross-talk between different parts of the chip to be reduced (see section 4.3.1). It was found that passing too much current through these heaters caused them to permanently stop working - effectively restricting the phase range achievable with these devices (see section 4.3).

Integrated Bragg gratings

Finally, DUW allows the definition of Bragg reflectors along the waveguides as they are being created[183]. The UV writing beam is split and recombined in an interferometer resulting in an interference pattern at the focal spot (see bottom left of figure 3.7). Translating the sample through this pattern averages out the interference effect leading to the creation of regular waveguides. However, by modulating one of the beams using an electro-optic modulator one can translate the interference pattern along with the sample, thus defining a Bragg grating along the waveguide[184]. We will make use of these Bragg gratings for the on-chip TES work described in chapter 6 and later in section 4.8 when we discuss future devices.

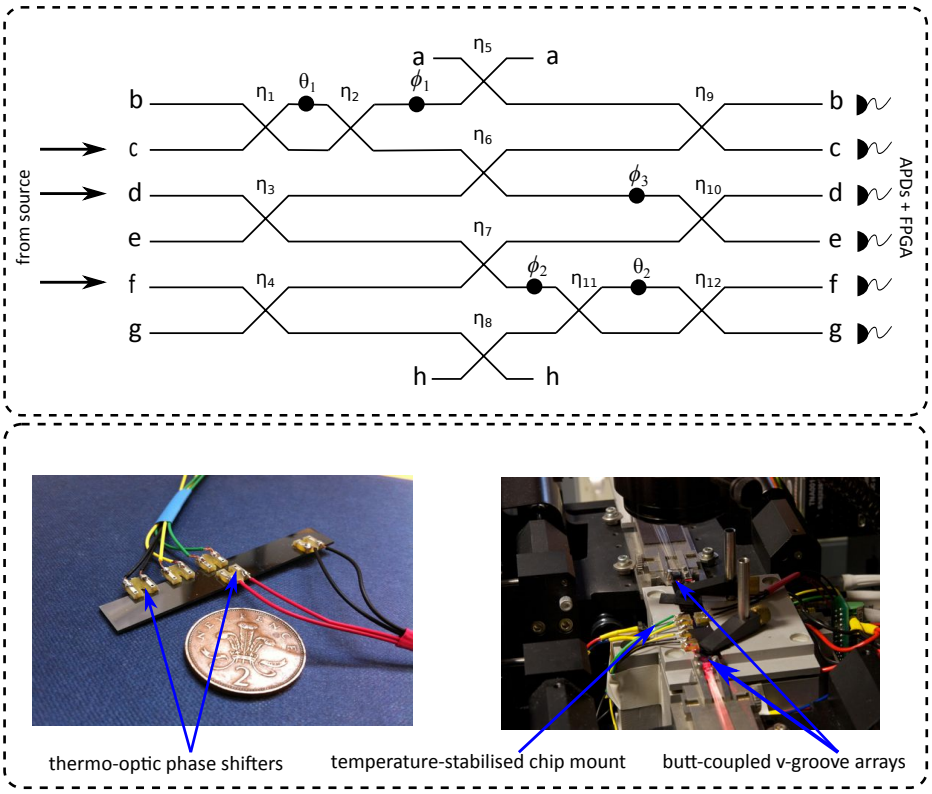


Figure 3.8: Integrated waveguide device and experimental setup. Top — the waveguide circuit consists of eight modes, labelled a to h, twelve beam splitters labelled η_1 to η_{12} and five thermo-optic phase shifters labelled θ_1-θ_2 and ϕ_1 to ϕ_3. The circuit was designed with beam splitter reflectivity's of $\eta = 1/2$ except the central ones labelled η_5-η_8 which were designed with a reflectivity of $\eta = 1/3$. Three single photons are launched into a subset of modes b to g and the output of each mode was monitored by a single photon detector. The ancillary modes a and h were not available for coupling. Bottom — a photo of the waveguide device with a two-penny piece for scale. The thermo-optic phase shifters are bonded to small PCB sections allowing us to solder control wires onto the chip. The waveguide circuit is mounted on a temperature-controlled block and butt-coupled with 10-port polarisation-maintaining v-groove arrays (VGAs) and index-matched oil using two piezo-controlled 6-axis alignment stages to optimise coupling.

3.1.3 Waveguide circuit design

The silica-on-silicon waveguide circuit design used for the experiments in the next two chapters is shown in the top half of figure 3.8. It consists of 12 beam splitters spanning 8 spatial modes and 5 thermo-optic phase shifters controlling 5 nested interferometers. The circuit was designed a number of years ago by Dr. Brian Smith, a former post-doc within our group and was fabricated by Dr. Dmytro Kundys at the University of Southampton.

The waveguide core was of dimension $5\,\mu\text{m} \times 5.5\,\mu\text{m}$, with index contrast 4.5×10^{-3} and was designed to support a single guided mode at 830 nm. The fabricated waveguide supported both the TM00 and TE00 polarised modes with slightly different effective indices owing to a residual stress-induced birefringence ($\Delta n \sim 1 \times 10^{-4}$). The waveguide mode was imaged using a mode-profiler (Spiricon SP620U), which confirmed single-mode operation and gave a circular, Gaussian mode with mode-field diameter $\sim 5\,\mu$m. Performing a mode-field overlap between the measured waveguide mode and a fibre-mode gave a theoretical maximum butt-coupling efficiency of $> 98\,\%$. A decision was made at design time not to fabricate the ancillary waveguide modes a and h to the edge of the chip. Unfortunately, this later complicated circuit characterisation. Additionally, during final polishing we accidentally removed too much material from the input face resulting in mode b no longer being exactly $250\,\mu$m spaced from modes a and c - preventing us coupling into this mode simultaneously with the others. Although the circuit was fabricated many years ago, the lack of an optimised multi-photon source, as well as a lack of resources, meant no experiments were able to be performed with this chip until the work described in this thesis.

A general 8×8 linear optical network requires 28 beam splitters and 7 phase shifters to implement an arbitrary unitary transformation on these modes [57]. Our circuit design is clearly not capable of such an arbitrary transformation but is in fact a 'compiled' design which allows the implementation of a number of three-qubit protocols. Using a restricted number of beam splitters in this way allowed us to minimise the loss of the device whilst still being able to perform interesting experiments. A theoretical description of this circuit design was first provided by Ralph [70], who showed its versatility in a range of different quantum logic applications, from teleportation, to cluster state generation and one-way

computing to basic error-correction. In these configurations, the central 4 beam splitters implement concatenated two-qubit C-PHASE-gates as previously shown in figure 2.8, whilst the surrounding beam splitters serve to implement single-qubit gates. A more complete explanation of the circuit design is saved until chapter 4.

The fragility and limited tuning range of the phase shifters made many of these possible quantum logic experiments difficult to perform. In this chapter, we instead explore for the first time three-photon interference effects in a complex integrated circuit. This removes the need for active phase control but nevertheless acts as an important primitive for future experiments. In fact, at the end of the chapter we show that such a passive circuit is still very interesting from a computational complexity point of view. In the next chapter, we are able to use the limited range of the phase shifters to perform a quantum teleportation experiment.

In all of these experiments, three of the photons from the SPDC source were coupled into polarisation maintaining (PM) fibre and launched into the waveguide circuit using a butt-coupled PM v-groove array. The waveguide modes were 250 μm spaced to match commercially available v-groove arrays, resulting in overall chip dimensions of 1 cm $\times$ 7 cm. Index matching oil is applied between the fiber array and the chip to reduce reflection losses and a 6-axis piezo-controlled alignment stage (Thorlabs NanoMax MAX603D) provides all the degrees of freedom necessary to achieve optimal simultaneous coupling into all six input modes. Two straight-channel waveguides were written either side of the circuit to help in alignment. The fiber-to-fiber transmission through the straight channels was $\sim 25\%$ whilst the total fiber-to-fiber transmission through the waveguide circuit averaged only $\sim 10\%$. This was due to the Hydrogen diffusing out of the waveguide core during writing leading to

lower index contrast (and thus higher loss) for waveguide channels written towards the end of the process. We will return to potential solutions to this problem in section 4.8. The piezo-driven axes were connected to a computer-controlled automatic alignment system (Thorlabs NanoTrak). During alignment, the piezos were set to drive the input v-grooves in a circular motion (of radius $\sim 3\,\mu$m) whilst the transmitted waveguide power was monitored on a photodiode to be used as an error-signal. This error-signal was fed-back to the alignment stages to move the fibre v-groove in order to increase transmission. Once optimal optical alignment had been achieved using this method, the piezo axes were operated in closed-loop mode, which maintained their position relative to the manual micrometer positions, in order to maintain optimal coupling throughout the experiment. An identical set-up is used on the output to a achieve maximal coupling from the chip into the single photon counting modules. A home-built FPGA-based logic unit records all desired coincidence counts simultaneously [2]. Temperature control by means of a thermo-electric Peltier element maintained stable phase control over many weeks during continuous thermo-optic heater operation.

3.2 Loss-tolerant circuit characterisation

After loss, a second major challenge in realising multi-photon interference on-chip is a robust method for verifying the device is behaving as we expect. However, existing quantum process tomography (QPT) techniques for verifying the complete quantum operation of the entire experiment [87,185] become impracticable once it becomes sufficiently complex. In particular, the number of input probe states and measurements required to perform QPT

[2] The FPGA program was written by Justin Spring and loaded onto a Xilinx Spartan 6 development board. A custom mezzanine board was also built to provide 16 input channels to the FPGA. Details of the design can be found in Justin Spring's DPhil thesis [139]

grows exponentially in the number of photons in the system ($16^N - 4^N$ parameters for a trace-preserving operation on N-photons)[186]. To alleviate this problem, alternative methods based on the compressed sensing technique have been developed which, although still scale exponentially in the system size, have significantly reduced resource requirements[187].

Part of the problem with this usual approach is that QPT provides more information than is often needed. Instead, it is very often sufficient to verify quantum behaviour for only a limited set of input states and output measurements. For example, in chapter 4, we simply wish to verify that our device exceeds the best possible classical teleportation fidelity, which we can do using only 3 input states despite the full three-qubit QPT requiring 4032 separate measurements. Similarly, we often have some idea about what has been realised in the laboratory and so we may only want to identify a particular state or operation within a restricted set of possibilities. These types of ideas are encapsulated in the recently developed tools of randomised benchmarking[188] and Monte-carlo process certification[189] and provide a way to verify quantum behaviour for a large range of states and processes without resorting to full QPT. We rely on these ideas in this chapter - we only want to identify the theoretical description, from a restricted set of possibilities, that best matches our experimental data.

In the first instance we could assume that the fabricated circuit has met its design specification and use this to model the device behaviour for the range of interesting possibilites - for example, for an input of classical or quantum particles. However, it is often necessary to be able to independently confirm that the design specification has been met. It is not generally possible to access individual components *in situ* using the external input and output ports, nor is it always possible to configure ancillary access ports for injecting probe states or performing detection locally. Alternatively, it is possible to infer the overall linear

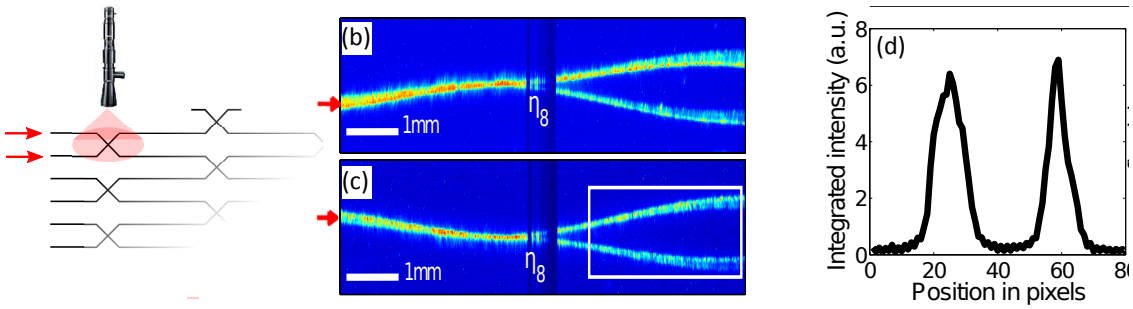


Figure 3.9: Loss-tolerant beam splitter characterisation. (a) The waveguide beam splitter element to be characterised is imaged from above using a 12x zoom telescope. (b,c) Charge-coupled device (CCD) images of light scattered in the transverse direction when light is coupled into the bottom (b) and top (c) input of the beam splitter. Red arrows denote the direction of propagation of the light. The white box denotes the integration region at the output for the ratiometric analysis. (d) The integrated intensity is used as part of a ratiometric analysis to determine the reflectivity of the beam splitter.

transformation implemented by a device from a series of classical interference experiments using only the nominal input and output ports[190,191]. Although these methods may be useful for cases where only the operation of the device as a whole is of concern, they do not give access to individual components. As will become clear in the next chapter, it is sometimes necessary to be able to separate out the behaviour of different parts of the device, meaning an element-wise characterisation is needed. Below, we present a method for the classical characterisation of individual fabricated circuit elements *in-situ*.

3.2.1 Beam splitters

We measured the beam splitter reflectivities, η_1 to η_{12}, by sending continuous-wave laser light through each splitter in turn and calculating the reflectivity using a ratiometric analysis which is independent of coupling and transmission losses[37,192]. This technique uses four measurements for each beam splitter, sequentially coupling light into each input port and recording the power at each output port. Due to the complex circuit topology, it is not generally possible to independently access the input and output modes, so instead, light

scattered out of the chip surface was used to measure the output powers. The ratiometric calculation is insensitive to different scattering efficiencies in the same way as to different interface coupling efficiencies. Figures 3.9(a) and (b) show a typical example, with 100 mW of laser light coupled into the chip (note that splitters 5 and 8 only had one available input, since modes a and h were not accessible for coupling, see figure. 3.8). The input polarisation was adjusted to maximise the amount of the transverse scatter, which was imaged using a CCD with a highly linear response. Integrating over a specified region then provides the required intensity measurements (Fig 3.9(c)). In this manner we were able to characterise ten of the twelve beam splitters with an accuracy of 2% as shown in table 3.1 along with their designed reflectivity. The final two beam splitters (η_2 and η_{11}) could not be characterised in the same manner because we did not have independent access to the input and output modes. Instead, the value of these beam splitters were inferred from an additional measurement of the effective reflectivity of the Mach-Zehnder unit when the phase shifter was scanned through its tuning range as discussed below.

3.2.2 Phase shifters

For each MZI, we characterised the phase by considering the interferometer as an effective beam splitter with a reflectivity determined by the phase shift and the reflectivities of the four relevant beam splitters as shown in figure 3.10. This was again characterised using the ratiometric technique. Adjusting the voltage, V, across the thermo-optic element varies the effective reflectivity, as shown in Fig. 3.10. Fitting this data then provides both an estimate of the zero-voltage phase of the interferometer (shown in table 3.2), and an independent consistency test of the four beam-splitter reflectivities which define each interferometer.

Beam splitter	Characterised ratio
η_1	0.54
η_2	0.47
η_3	0.61
η_4	0.65
η_5	0.23
η_6	0.28
η_7	0.29
η_8	0.28
η_9	0.56
η_{10}	0.57
η_{11}	0.35
η_{12}	0.55

Table 3.1: Characterised beam splitter ratios. All 12 beam splitters were characterised using scattered light as described in the text. Rows shaded gray are designed with a reflectivity of 1/3 whilst all others should be 1/2. The values of η_2 and η_{11} were inferred using additional phase shifter data since it was not possible to access the ports on these beam splitters independently. An error of 0.02 is estimated for each value based on comparing this method to directly measuring the beam splitting ratio using transmitted light [192].

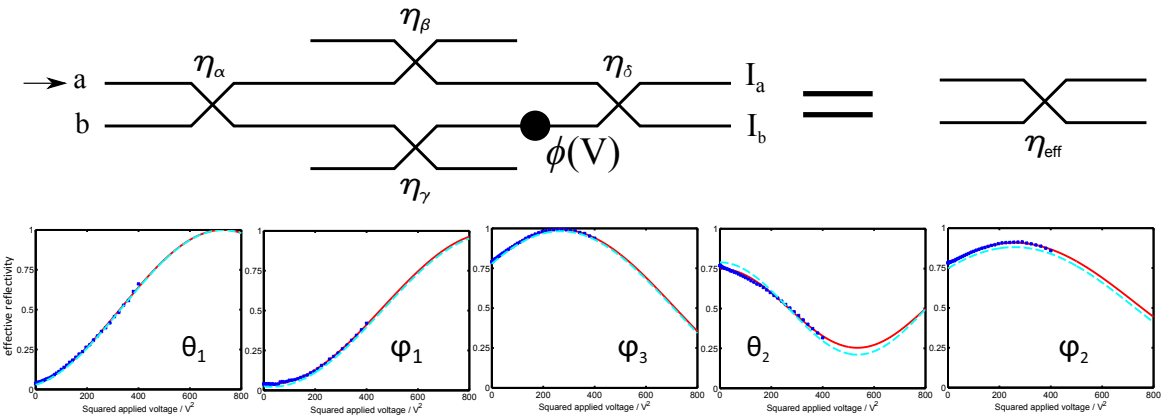


Figure 3.10: Phase shifter characterisation Top - Each of the five phase shifters are characterised by considering the Mach-Zehnder Interferometer in which it is embedded. In general, the MZI's are formed by four beam splitters η_α through η_δ which must be known in order to recover the phase from intensity measurements at the output. We use a ratiometric method, as described in the text, in order to eliminate the effect of input and output losses on the characterisation. Bottom - interference fringes produced by scanning each of the five phase shifters. The zero-voltage phase offset and voltage-phase relationship are obtained from a theoretical fit (red line). A second theoretical fit substitutes the independently characterised beam splitting ratios into the model (blue line). The good agreement between the two fits confirms the self-consistency of our characterisation.

	$\phi(V) = a + bV^2$	
Shifter	a	b
θ_1	-1.12	0.0048
ϕ_1	3.48	0.0038
θ_2	-5.65	0.0040
ϕ_2	2.55	0.0039
ϕ_3	4.01	0.0035

Table 3.2: Characterised phase shifter data. We record intensity data from each MZI as a function of applied voltage to the phase shifter to recover the zero voltage phase offset, a and the phase-voltage coefficient, b.

These checks agreed within the measurement uncertainty.

Once all of the beam splitters and phase shifters have been characterised, the unitary linear transformation of the entire circuit (U_{circ}) can then be computed. This unitary transformation is calculated by multiplying the individual beam splitter and phase shifter matrices together respecting the appropriate time-ordering.

3.3 Simulating interference visibility

The results of the circuit characterisation can now be used to simulate the expected behaviour of our circuit. For the first experiments we want to run, we are interested in the visibility of higher-order multi-photon correlation functions. After discussing how such visibilities may be calculated, we move on to discuss how one can test whether the experimental data is consistent with a fully quantum model of interference.

The simulated visibilities in our interference experiments were calculated by first simulating the complete quantum output state, $|\psi_{\text{out}}\rangle$, using the characterised circuit unitary, $\mathbf{U}_{\text{circ}}$. Photons are coupled into the input modes, a_1 through a_M, and undergo a unitary transformation governed by $\mathbf{U}_{\text{circ}}$ before ending up in the output modes, b_1 through b_M

where the light is detected.

The state of a general separable input to the circuit is specified by the matrix of Fock state coefficients **C**. For a network with M spatial modes and a maximum of N photons in any of those modes, the input state is,

$$|\psi_{\text{in}}\rangle = \prod_{i=1}^{M} \left(\sum_{j=1}^{N+1} C_{i,j} \frac{\hat{a}_i^{\dagger(j-1)}}{\sqrt{(j-1)!}} \right) |\text{vac}\rangle \,, \tag{3.5}$$

where the rows of the matrix **C** must each satisfy the normalisation condition, $\sum_{j=1}^{N+1} |C_{i,j}|^2 = 1$. The effect of our linear optical network is to transform the vector of input creation operators, $\hat{\mathbf{a}}_i^\dagger$ into linear combinations of output state creation operators, $\hat{\mathbf{b}}_i^\dagger$ via the unitary mapping, $\hat{\mathbf{a}}^\dagger = \mathbf{U}_{\text{circ}} \hat{\mathbf{b}}^\dagger$, where $\mathbf{U}_{\text{circ}}$ is readily calculable from the parameterised circuit. The final output state of the network is then simply given by substituting this expression into equation 3.5 [192].

Calculating interference visibilities

We can now use this output state to find the intensity cross-correlation functions at zero and infinite temporal delay and hence calculate the interference visibility. For photons launched into the set of modes $\{a_i\}$ of a linear optical circuit the intensity cross-correlation between the set of output modes $\{b_i\}$ at zero temporal delay is

$$\Gamma^{(0)}_{\{b_i\}} = \left\langle \psi_{\text{out}}^{\{a_i\}} \right| \prod_{b_i} \hat{\text{I}}_{b_i} \left| \psi_{\text{out}}^{\{a_i\}} \right\rangle , \tag{3.6}$$

where the intensity operator on mode i is $\hat{\text{I}}_i = \hat{b}_i^\dagger \hat{b}_i$. Suppose the photon in mode a_d undergoes a temporal delay, then that photon does not contribute to any quantum interference

effects and can be propagated separately. However, because we use slow-detectors, we cannot resolve this time-of-arrival difference and so our output click statistics will be determined by a classical mixture,

$$\rho_{\text{out}} = \left|\psi_{\text{out}}^{\{a_i, i\neq d\}}\right\rangle\!\left\langle\psi_{\text{out}}^{\{a_i, i\neq d\}}\right| + \left|\psi_{\text{out}}^{(a_d)}\right\rangle\!\left\langle\psi_{\text{out}}^{(a_d)}\right|. \tag{3.7}$$

The intensity cross-correlation function at infinite delay is then,

$$\Gamma_{\{b_i\}}^{(\infty)} = \text{Tr}\{\rho_{\text{out}} \prod_{b_i} \hat{\text{I}}_{b_i}\}, \tag{3.8}$$

and the visibility of the interference pattern between the n^{th}-order output coincidence counts is then given by,

$$V_{\text{quant}} = \frac{\Gamma_{\{b_i\}}^{(\infty)} - \Gamma_{\{b_i\}}^{(0)}}{\Gamma_{\{b_i\}}^{(\infty)}}. \tag{3.9}$$

The corresponding classical visibility of this interference pattern is given by injecting three equal amplitude coherent states with mutually randomised phases into modes $\{a_i\}$. This ensures we mimic independent sources of light which will have no first-order correlation as required to compare against Hong-Ou-Mandel-type quantum interference. The resulting output vector of complex amplitudes is, $\mathbf{e}_{\text{out}} = \mathbf{U}_{\text{circ}}\mathbf{e}_{\text{in}}$ where $\mathbf{e}_{\text{out}}$ is the vector of time-independent electric fields in each of the output modes ($e_j = e^{i\theta_j}$) and similarly for $\mathbf{e}_{\text{in}}$. The phase-averaged intensity cross-correlation function is then

$$\Gamma_{\{b_i\}}^{\prime(0)} = \frac{1}{(2\pi)^n} \int_0^{2\pi} \ldots \int_0^{2\pi} \prod_{b_i} |(\mathbf{e}_{\text{out}})_{b_i}|^2 \, d\theta_1 \ldots d\theta_n. \tag{3.10}$$

The classical cross-correlation function at infinite delay is calculated in a similar manner,

taking incoherent sums between the delayed and non-delayed photons,

$$\Gamma'^{(\infty)}_{\{b_i\}} = \frac{1}{(2\pi)^n}\int_0^{2\pi}\cdots\int_0^{2\pi}\prod_i\left(|(\mathbf{e}_{\text{out}}^{\{a_i,i\neq d\}})_{\text{b}_\text{i}}|^2 + |(\mathbf{e}_{\text{out}}^{(a_d)})_{\text{b}_\text{i}}|^2\right)d\theta_1\ldots d\theta_n. \tag{3.11}$$

From which we calculate the classical interference visibility,

$$V_{\text{class}} = \frac{\Gamma'^{(\infty)}_{\{b_i\}} - \Gamma'^{(0)}_{\{b_i\}}}{\Gamma'^{(\infty)}_{\{b_i\}}}. \tag{3.12}$$

Monte-Carlo Simulations

In order to estimate the errors on our simulated visibilities we used a Monte-Carlo approach to run the above calculations 1000 times each, using a different $\mathbf{U}_{\text{circ}}$ for each run. The $\mathbf{U}_{\text{circ}}$ for each run was calculated by drawing each of the 15 measured parameters characterising the circuit from a normal distribution centred on the parameter mean and with a standard deviation given by 0.02 for each of the beam splitters and those shown in table 3.2 for the phase shifters. The final expected visibilities were then calculated from the mean of this distribution and the standard deviation gave us an estimate on the error of our simulations.

3.3.1 Chi-squared test for observed visibilities

With a method of calculating quantum interference visibilities now established, how can we certify we have achieved genuine multi-photon quantum interference given a set of experimentally measured visibilities? Naively, we might apply a chi-square test to the measured residuals, $w'_i = v_i - w_i$ (measured visibility, v_i, minus predicted visibility, w_i). However, these residuals are not independent random variables because the uncertainties in the simulated visibilities are correlated through the model parameters. We show below that by making

a linear approximation in our model and analysing the residuals in a rotated basis (i.e. particular linear combinations of the data points) allows us to perform a suitable chi-square test on our data. This test-statistic then bounds the probability that a set of experimentally measured visibilities would have occurred given an underlying model. This allows us to say with quantitative confidence whether our measurements are consistent with a classical or quantum interference model.

The probability to observe the set of m interference visibilities $\boldsymbol{v} = v_1, \ldots, v_m$ given the set of n model parameters $\boldsymbol{\alpha} = \alpha_1, \ldots, \alpha_n$ is constructed from the probability density function (p.d.f) $f(\boldsymbol{v}) \sim \int \mathrm{d}\alpha\, g(\boldsymbol{\alpha}) f(\boldsymbol{v}|\boldsymbol{\alpha})$ where $g(\boldsymbol{\alpha})$ is a p.d.f. for $\boldsymbol{\alpha}$. In our case, $\boldsymbol{\alpha}$ are the 15 characterised circuit elements and a distinguishability factor accounting for the non-perfect spectral overlap of the single photons. Assuming a Gaussian error model for $g(\boldsymbol{\alpha})$, and a Gaussian conditional probability, $f(\boldsymbol{v}|\boldsymbol{\alpha})$, we can write,

$$f(\mathbf{v}) \sim \int \left(\prod_i^n \exp\left(\frac{-\alpha_i^2}{2s_i^2}\right) \prod_i^m \exp\left(\frac{-(v_i - w_i(\boldsymbol{\alpha}))^2}{2\sigma_i^2}\right) \right) d^n\boldsymbol{\alpha}, \tag{3.13}$$

where $w_i(\boldsymbol{\alpha})$ is the calculated i^{th} visibility based on the set on n model parameters $\boldsymbol{\alpha}$, σ_i is the measurement standard deviation in the observed visibility v_i, and s_i is the measurement standard deviation of the characterised model parameter α_i. In order to re-cast this integral into the more conventional multivariate Gaussian form we must approximate each $w_i(\boldsymbol{\alpha})$ to be locally linear in α,

$$w_i(\boldsymbol{\alpha}) \approx w_i^0 + \sum_j^n \frac{\partial w_i}{\partial \alpha_j} \alpha_j. \tag{3.14}$$

The Jacobian matrix, $J_{i,j} = \frac{\partial w_i}{\partial \alpha_j}$ is approximated by calculating the linear gradient of w_i as the parameter α_j is varied over the range $\pm s_j$. We verify the accuracy of this approximation

to be correct within 0.02 over the extended range $\pm 2s_j$. Substituting this in equation (3.13) results in a standard multi-dimensional Gaussian integral which we express in conventional matrix notation in order to yield an analytic solution,

$$f(\boldsymbol{v}) \sim \int \exp\left(-\boldsymbol{\alpha}^T \boldsymbol{A}\boldsymbol{\alpha} + \boldsymbol{J}^T\boldsymbol{\alpha} + k\right) \mathrm{d}^n\alpha \tag{3.15}$$

$$= \sqrt{\frac{(2\pi)^n}{|\boldsymbol{A}|}} \exp\left(\frac{1}{4}\boldsymbol{J}^T\boldsymbol{A}^{-1}\boldsymbol{J} + k\right) \tag{3.16}$$

where,

$$A_{i,j} = \sum_l^m \frac{1}{2\sigma_l^2}\frac{\partial w_l}{\partial \alpha_i}\frac{\partial w_l}{\partial \alpha_j} + \frac{1}{2s_i^2}\delta_{i.j} \tag{3.17}$$

$$J_i = \sum_l^m \frac{w_l}{\sigma_l^2}\frac{\partial w_l}{\partial \alpha_i}, \tag{3.18}$$

$$k = -\sum_l^m \frac{w_l^2}{2\sigma_l^2} \tag{3.19}$$

and we define $w_i = v_i - w_i^0$. We wish to express this as a function of quadratic visibility residuals, so we define $L_i = \sum_l^m \frac{1}{\sigma_l^2}\frac{\partial w_l}{\partial \alpha_i}$ allowing us to write $\boldsymbol{J} = \boldsymbol{L}^T\boldsymbol{w}$ and thus $f(\boldsymbol{w}) \sim \exp\left(-\boldsymbol{w}^T\boldsymbol{B}^{-1}\boldsymbol{w}\right)$ and $\boldsymbol{B} = \boldsymbol{L}\boldsymbol{A}\boldsymbol{L}^T$ is the covariance matrix for the multi-dimensional Gaussian distribution. We have dropped the normalisation constants in front of the p.d.f. since they do not contribute to the calculation of any test statistics we wish to use. Finally, we analyse our data in a rotated basis, $x_j = \sum_i c_{ij}(v_i - w_i^0)$, which diagonalises the covariance matrix. Thus we arrive at a factorable form for $f(\boldsymbol{x})$

$$f(\boldsymbol{x}) \sim \prod_i^m \exp\left(-\lambda_i x_i^2\right), \tag{3.20}$$

where λ_i are the eigenvalues of the covariance matrix $\boldsymbol{B}$. The visibilities $\mathbf{x}$ are thus statistically independent ($\Pr(x_i|x_j) = \Pr(x_i)$) and a typical chi-squared test, $\chi^2_r = 1/m \sum_i^m \lambda_i x_i^2$, can now be applied. If we instead errantly assume that the uncertainties in our simulated visibilities are uncorrelated, the naive probability density function would be

$$f(w_1, \ldots, w_m) \sim \prod_i^m \exp\left(-\frac{1}{2}\frac{w_i^2}{\sigma^2 + \sum_j s_j^2 (\frac{\partial w_i}{\partial \alpha_j})^2}\right) \tag{3.21}$$

which has a built-in assumption of 'independent errors' in the construction of the effective standard deviation in the denominator. Using this incorrect formulation leads to a chi-squared which appears smaller than its true value.

3.4 Multi-photon on-chip interference

Having characterised the individual circuit elements classically and developed a verification protocol, we now investigate the operation of the circuit using quantum input light. We generated three individual input photons from the two spectrally factorable down-conversion pair sources discussed above, using both photons from one pair and using the other as a heralded single photon (see figure 3.2). Detecting all four photons in this way to give four-fold coincidence measurements reduces the effects of noise terms.

3.4.1 Two-photon interference

We first study the quantum interference from two-photon inputs to confirm our classical device characterisation and our single-photon indistinguishability. In these experiments, we injected photons into two of the selected input modes: *cf*, *cd*, or *df* (corresponding to

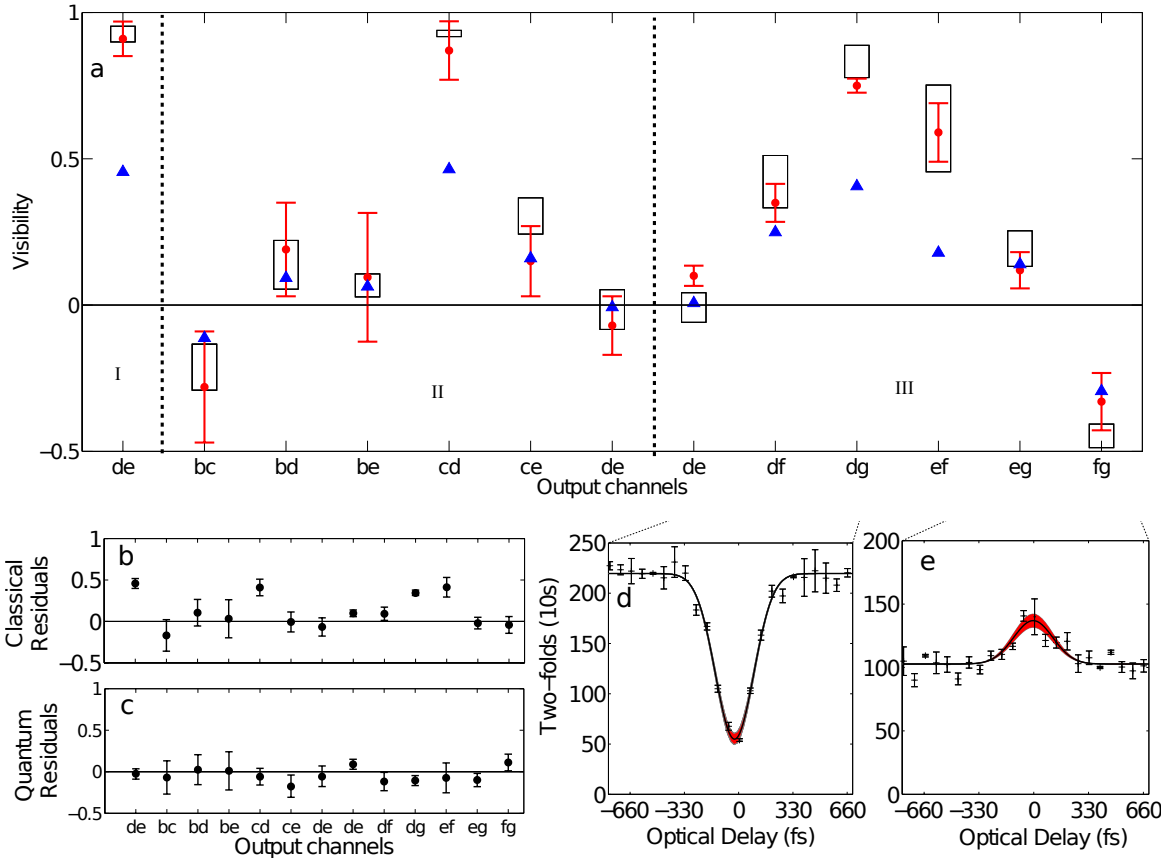


Figure 3.11: Two photon quantum interference. **(a)** Experimentally measured two-photon interference visibilities (red circles) are compared against the quantum (clear boxes) and classical (blue triangles) predictions. Regions I, II and III contain experiments using input modes *cf*, *cd*, and *df*, respectively. The errors shown on the simulated quantum and classical visibilities were calculated by Monte-Carlo simulation (see Supplementary Methods); errors on the classical visibilities are smaller than the marker size due to the decrease in sensitivity of classical interference to the measured circuit parameters. **(b-c)** The residuals between the measurements and the calculated classical and quantum visibilities. **(d-e)** Example two-fold coincidence counts between output channels *dg* and *fg* when two photons are input into modes *df* and the optical delay of mode *d* is varied. The shaded area shows the uncertainty in the determined visibility.

regions I, II and III in Fig. 3.11(a)). We then measured interference visibilities for all possible coincidence outcomes by varying the timing of one photon using an off-chip optical delay stage (e.g. Fig. 3.11(d)-(e)), and compared the observed values to the quantum (boxes) and classical (triangles) predictions (see Fig. 3.11(a)). The expected quantum visibilities were calculated directly from the quantum output state predicted by the classically characterised circuit unitary (see section 3.3). We corrected this simulated interference visibility for the independently measured two-photon distinguishability using a 50:50 fibre beam splitter. The corresponding classical visibilities were calculated by replacing the input Fock states with phase-averaged coherent states as described in section 3.3.

The chi-squared test developed above can now be used to verify that these data are consistent with the quantum predictions whilst in strong disagreement with classical theory [62]. The likelihood for observing a set of interference visibilities is calculated given measurement uncertainties in the observed interference visibilities and the underlying circuit parameters. The residuals from quantum theory (Fig. 3.11(c)) give a reduced chi-squared $\chi_r^2 = 0.9$; a value at least this large will occur with probability $P(\chi_r^2 \geq 0.9) = 0.5$. By contrast, the residuals from classical theory (Fig. 3.11(b)) give $\chi_r^2 = 23$ corresponding to $P(\chi_r^2 \geq 23) < 10^{-16}$. Furthermore, the ultimate classical visibility limit of 1/2 is expected and observed to be exceeded by output mode combinations *de*, *cd*, and *dg*.

The three experiments each used different pair-wise combination of the three chosen input photons, allowing the different input photons to be tested individually against each of the other two. Each pair-wise combination is separated by a dashed line in figure 3.11. In all three cases, the experiments are in good agreement with the expected quantum results. This suggests that our element-wise characterisation technique is a good predictor of device

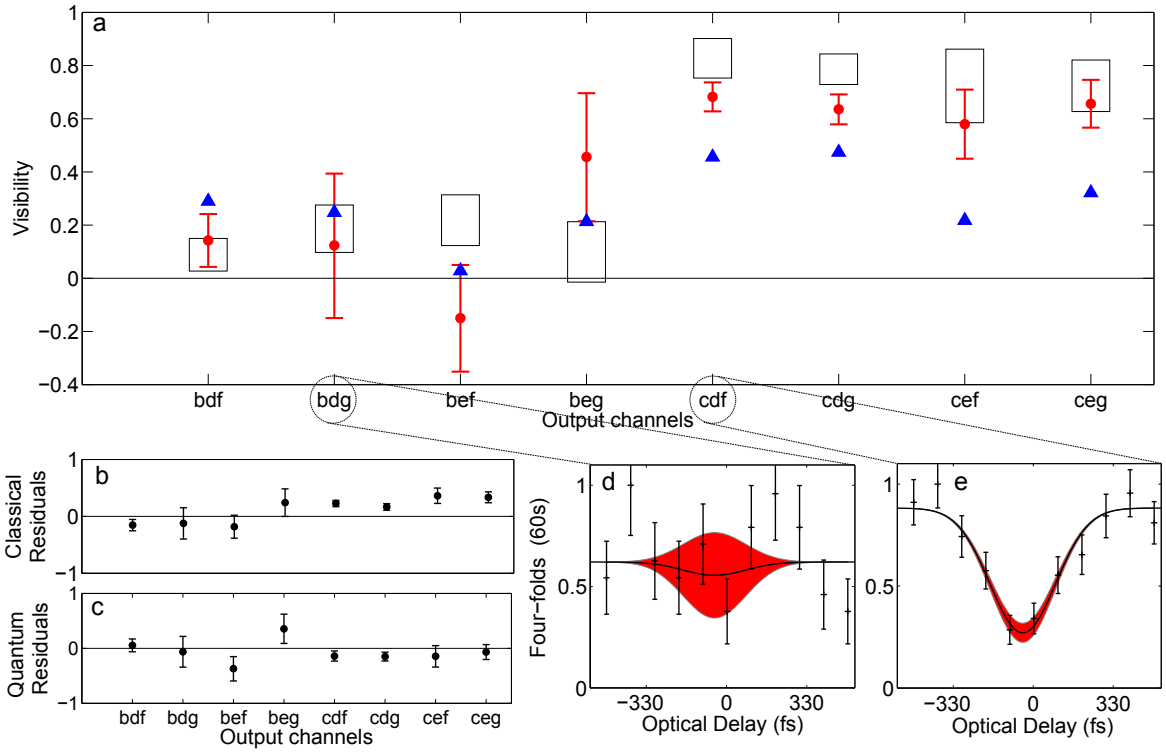


Figure 3.12: Three-photon quantum interference **(a)** The experimentally recorded three-photon interference visibilities (red circles) are compared against the predicted quantum (clear boxes) and classical (blue triangles) results as in Figure 3.11. The errors shown on the simulated quantum and classical visibilities were calculated by Monte-Carlo simulation (see Supplementary Methods); errors on the classical visibilities are smaller than the marker size due to the decrease in sensitivity of classical interference to the measured circuit parameters. **(b-c)** The residuals between the measurements and the calculated classical and quantum visibilities. **(d-e)** Example four-fold coincidence counts of output channels *cdf* and *bdg* when the input photon in mode *d* is temporally delayed. The shaded area shows the uncertainty in the determined visibility.

performance and shows that each of the single photon quantum states has a good fidelity with the others, a critical factor for observing genuine three-photon interference.

3.4.2 Three-photon interference

To demonstrate the complexity of the quantum circuit, we next studied the higher-order non classical interference which arises when the three coupled interferometers are all operating simultaneously and in parallel, and are each injected with quantum light at the input. We observe this via the (heralded) three-photon coincidence counts (e.g. Fig. 3.12(d)-

(e)), with three individual input photons coupled into modes *c*, *d* and *f*. After setting the temporal delay between input modes c and f to maximise their two-photon interference, we varied the delay for input mode d while simultaneously monitoring the eight three-fold coincidence combinations described in Fig. 3.12(a). The combined effect of the various sources of loss within the experiment reduce the rate of measured four-photon coincidence counts, increasing the experimental time required to collect sufficient data. The expected four-photon data rate can be estimated by multiplying the four-photon pair production rate (given by $(\lambda^2)^2$) and experimental losses by the repetition rate of the laser:

$$\begin{aligned} R_{\text{data-rate}} &= (\lambda^2)^2 \times \eta^4_{\text{source-to-fibre coupling}} \times \eta^3_{\text{chip loss}} \times \eta^4_{\text{det. eff.}} \times R_{\text{rep. rate}} \\ &= (0.03)^2 \times (0.4)^4 \times (0.07)^3 \times (0.45)^4 \times (80 \times 10^6) \\ &= 25\,\text{mHz} \end{aligned} \tag{3.22}$$

We observed average four-photon coincidence rates slightly lower than this estimate at around 16 mHz, due to the deterioration in source and chip-coupling over time. In order to collect sufficient date, we measured the heralded three-folds continuously for 294 hours, iterating a full scan of the temporal delays each minute. This method averages out long-term systematic effects, such as drifts in the chip coupling efficiencies and photon source performance[193], and allows an accurate calculation of the statistical error in the counts.

As in the two-photon experiments, the observed three-photon quantum interference visibilities agree with quantum predictions based on the individually characterised circuit elements and are completely inconsistent with the equivalent classical predictions (Fig. 3.12(b)-(c)). The quantum prediction gives $\chi^2_r = 1.5$ with $P(\chi^2_r \geq 1.5) = 0.2$, whereas the classical

prediction gives $\chi_r^2 = 6$ with $P(\chi_r^2 \geq 6) < 10^{-8}$). Moreover, using a global optimisation routine, we determined the maximum classical interference visibility for any possible circuit parameters with this circuit topology to be 0.59. This ultimate limit is exceeded by more that one standard deviation by output channel combination *cdf*. Furthermore, a chi-squared test shows that the circuit parameters which result from this optimisation are strongly inconsistent with our measured values ($P(\chi_r^2 \geq 4.4) = 10^{-8}$). Thus, only the full quantum explanation can plausibly account for the higher observed visibilities. We note that Fig. 3.12(a) includes all measured three-fold coincidence combinations, including several that occur very rarely and which therefore lead to large error bars in the measured visibility. Nevertheless, including all observed data in our analysis, we can exclude classical models with extremely high confidence, despite the low count rates in some channels.

Finally, to verify that the observed interference results from a quantum interaction of all three photons and is not explainable via a bi-separable interaction only, we also simulated the expected quantum visibilities when one of the three photons remains completely distinguishable from the other two at all temporal delays. These quantum bi-separable explanations are also inconsistent with the data, with a probability of $P < 10^{-8}$ of producing the observed interference signature. Thus we may conclude with high confidence that we have achieved genuine three photon quantum interference within our integrated circuit.

3.5 Boson Sampling

"Between bosons and fermions - the bosons got the harder job"

Avi Wigderson, 2000

This chapter has detailed the first experimental demonstration of genuine three-photon quantum interference in a linear network on a photonic chip. It is clear that this achievement is still modest with respect to the requirements of performing full universal quantum computation as laid out in chapter 2. This motivates the search for schemes that can demonstrate quantum-enhanced computation with significantly relaxed experimental requirements. In 2011, a limited-form of quantum-enhanced computation was put forward by Scott Aaronson and Alex Arkhipov (AA), computational complexity theorists working at M.I.T [194]. They described this new computational scheme — Boson Sampling — and proved that whilst it is very likely a classically hard problem, it can be likely solved by a relatively straightforward quantum machine. Demonstrating the central premise of this protocol in the lab shows a quantum computer with a computing capacity beyond any classical computer may be closer to being realised than previously thought.

Below, we briefly review the Boson Sampling problem and report our implementation of it on the small-scale photonic device presented in this chapter. This is not intended to be a thorough report on this work - an accessible, comprehensive treatment of the full theoretical and experimental details can be found in the D.Phil thesis of my graduate colleague, Justin Spring [139]. Rather, this section attempts to highlight that already interesting phenomena may be investigated using the small-scale interference of identical photons in a passive linear optical network pioneered in this chapter.

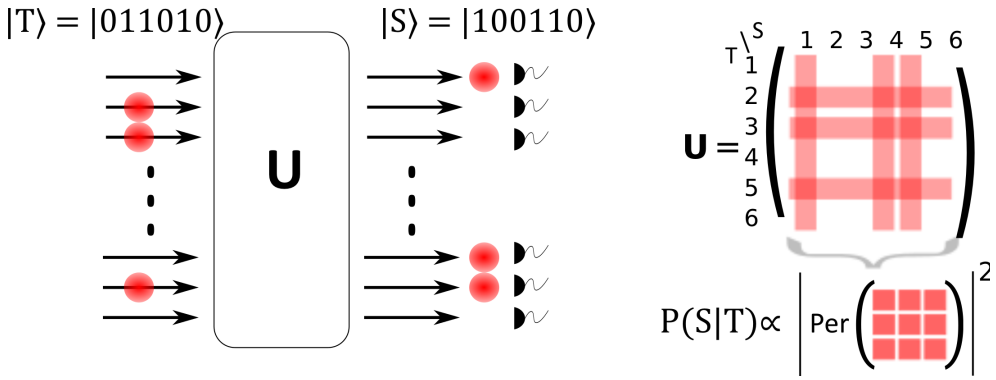


Figure 3.13: Boson Sampling concept. A specified initial distribution of single photons $|T\rangle$ evolve through a linear optical network with an overall unitary transformation U. The aim is to sample from the output probability distribution of possible output configurations $|S\rangle$. The probability of a specific outcome S occurring given an input T, $P(S|T)$ is found by computing the permanent of a specific sub-matrix of the overall linear-optic unitary transformation U as shown on the right of the figure.

3.5.1 Computational complexity of linear optics

The Boson Sampling problem is most easily understood by studying the quantum machine which solves it. We begin by inputing indistinguishable single-photon Fock states into the first n input modes of a passive linear optical interferometer and vacuum states into the remaining $m - n$ input modes, where $m \approx \mathcal{O}(n^2)$ (see figure 3.13). The input state evolves through the complex network of beam splitters and phase shifters and finishes in a highly-entangled output state,

$$|\psi_{\text{out}}\rangle = \sum_S \gamma_S \left| n_1^{(S)}, ..., n_m^{(S)} \right\rangle, \tag{3.23}$$

where s represents a specific output configuration of the n photons distributed across the m modes and γ_S is the associated probability amplitude for that configuration. The aim of the Boson Sampling problem is to draw samples from this probability distribution of possible output configurations. In the quantum machine we have just described this simply requires counting the number of times a specific measurement outcome occurs.

A machine designed to *simulate* sampling from this boson distribution would take as input a description of the linear optical transformation, U, and the initial input state, $|T\rangle$. Drawing samples from the correct probability distribution requires calculation of the probability amplitudes γ_S. It is well known that this requires the computation of the permanent of specific submatrices of the unitary network description,

$$P(\mathbf{S}|\mathbf{T}) = |\langle \mathbf{S}|\Psi_{\text{out}}\rangle|^2 = \frac{\left|\text{Per}(\mathbf{U}^{(\mathbf{S},\mathbf{T})})\right|^2}{\prod_{j=1}^{M'} S_j! \prod_{i=1}^{M} T_i!} \tag{3.24}$$

where the $N \times N$ submatrix $\mathbf{U}^{(\mathbf{S},\mathbf{T})}$ is obtained by keeping S_j (T_i) copies of the j^{th} row (i^{th} column) of $\mathbf{U}$. The permanent, which describes bosonic scattering amplitudes, has long been known to be classically hard to compute[195]. In contrast, fermionic scattering amplitudes are described by the determinant, which is easily classically computable. It is not at all surprising that calculating this full output probability distribution is a computationally hard task - it requires the evaluation of an exponential number of possible output configurations (since the size of $S \sim n!$) each of which requires the computation of an exponentially hard matrix permanent. However, even the quantum machine described above could not produce this probability distribution efficiently as it only outputs samples from the distribution, not the probability distribution itself. It would take an exponential number of samples to actually deduce the distribution.

The key result of Aaronson and Arkhipov was to show that a machine based on the laws of classical physics could not even generate *samples* from this distribution efficiently. Further, and most importantly, they showed that it is very likely not possible for a classical machine to even *approximately* sample from this distribution efficiently. One can thus envision a

race between a classical and a quantum machine to sample the boson distribution, to a given precision, given an input state and **U**. The classical machine would evaluate at least part of the probability distribution, which requires the computation of matrix permanents. The quantum machine instead creates indistinguishable bosons, directly implements **U**, and records the outputs. For a sufficiently large problem size, the quantum machine is expected to win [194].

AA conjectured that only 20-30 photons propagating in 400-900 optical modes would be required to show a clear quantum advantage in the Boson Sampling problem. With no feed-forward, logic gates or ancillary entanglement required, the threshold for performing a true quantum-enhanced computation is dramatically reduced and has spawned a flurry of theoretical and experimental results related to its implementation [127,133,150–153,195–211].

3.5.2 Results

A Boson Sampling device based on 20-30 bosons would likely be able to produce output samples faster than any classical machine could simulate it [194,201]. This would represent the first ever (limited) quantum computation able to outperform a classical computer and would represent a huge advance for the quantum information field.

This motivated a world-wide effort by experimental quantum optic groups [133,152,198], including ours [127], to demonstrate the basic principles of this protocol. It was not the purpose of these experiments to show a computational speed up — experiments using 20-30 photons are still many years from realisation — but rather to to show that an implementation of this protocol was possible in the face of real-world losses, imperfect photon sources and imperfect circuit characterisation.

The Boson Sampling device makes use of the same genuine multi-photon interference effects and multi-port interferometers we have studied in this chapter. Although the experiment was very similar to the three-photon interference data collection presented above, there were a few key differences in the implementation

- The two optical delay stages did not scan the temporal delay to generate 'dips' but were left at the zero delay position in order to collect more data when all three photons were indistinguishable
- All twenty possible 3-fold detector output combinations were recorded as opposed to the 8 considered above
- The pump power of the SPDC sources was reduced ($\lambda^2 \sim 0.01$) to mitigate the effect of higher-order pair emission on the probability distribution
- A 'black-box' characterisation method was used to account for unbalanced propagation and detector losses which affected the probability distribution (the effects of these losses are cancelled out when measuring an interference dip visibility as we did above).

The latter point, in particular, was of great importance in analysing our data. Although the element-wise circuit characterisation method was used successfully in the analysis of the three-photon interference visibility data (and proved critical to the quantum teleportation analysis in chapter 4), extracting accurate fiber-coupling and detector loss figures is difficult. Unbalanced detector losses skew the measured probability distribution away from the classically-hard distribution we are trying to sample from. Around the time of this experiment a characterisation protocol was published [191] which, whilst saying nothing about the individual circuit elements, gave a way to measure the unitary transformation of the entire

liner optical network including the effects of unbalanced loss in the circuit and detectors[3].

The central result of three-boson sampling is shown in figure 3.14. We repeatedly inject the three photons into the 6-mode photonic chip as we already showed in figure 3.8. The same three input modes ($\{c, d, f\}$) were used, as before, resulting in the input state vector $|T\rangle = |011010\rangle$. All output ports were monitored and we collected all possible three-fold coincident events. Shown in blue are the measured relative frequencies for every allowed outcome, along with their observed statistical variation. The corresponding theoretical distribution, calculated using the right-hand side of equation (3.24) is shown in red along with the uncertainties arising from the characterisation of $\mathbf{U}$. We quantify the match of these two distributions ($\mathbf{P}^{\text{exp}}$ and $\mathbf{P}^{\text{th}}$) by calculating the L_1 distance $d(\mathbf{P}^{\text{exp}}, \mathbf{P}^{\text{th}}) = \frac{1}{2}\sum_S |P_S^{\text{exp}} - P_S^{\text{th}}|$ between them.

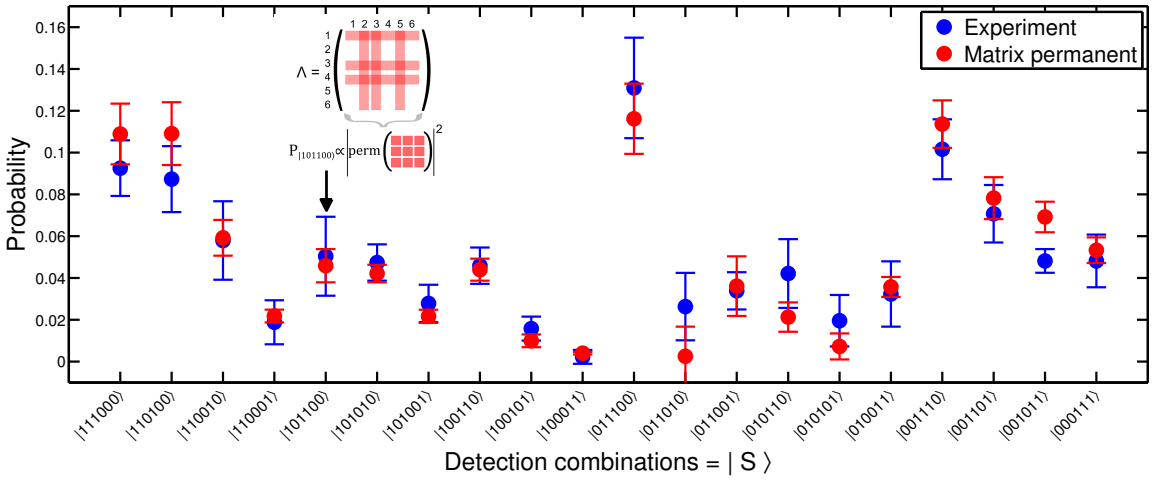


Figure 3.14: Three-photon Boson Sampling results. Three single photons are prepared in the input state $|T\rangle = |011010\rangle$ and the probability for each of the possible $\binom{6}{3} = 20$ output states are plotted for the experimentally measured outcomes (blue). Data was collected for ~ 160 hours and the statistical variations in counts are shown by the error bars. The theoretical distribution (red) is obtained from the permanents of sub-matrices constructed from the full transformation as shown in the inset. The red error bars arise from uncertainties in the characterisation of our linear optical transformation.

[3]Of course exactly because of this loss, the resulting transformation is not actually unitary. However, this changes none of the arguments regarding the analysis of the data in the context of the original Boson Sampling problem [140]

Even in an ideal experiment with perfect state preparation and detection, the statistical variations would result in nonzero d. If we substitute for our experimental data a Monte Carlo sampling of $\mathbf{P}^{\text{th}}$ with sample size equivalent to our experiments, we calculate d=0.043± 0.012. For our actual experiment, we calculate d=0.094 ± 0.014 which suggests experimental imperfections are playing a significant role in our machine. We showed that the main sources of error are partial photon distinguishability and higher-order pair emission from our SPDC photon sources[127,139]. When these error terms are added to our theoretical description, we recover an L_1 distance which agrees with what we would expect given our finite sample size.

It is now an active area of current research to investigate the effect of these additional error terms on the feasibility of out-performing a classical machine[150,196,197,201–209,211]. This experiment quantified the current effect of these errors and showed that there were no other un-explained error sources we had not accounted for. As experiments begin to progress to using more photons in more modes it may become challenging to even verify that our device is working correctly. However a range of theoretical proposals and experiments have already begun to explore this issue[151,153,199,200].

Chapter 4

On-chip Quantum Teleportation

This chapter builds on the results of chapter 3: we now use the three photons distributed over the 6-port on-chip linear network as three qubits in a quantum circuit and perform the first on-chip realisation of the quantum teleportation protocol. In contrast to the results of the previous chapter, the performance of this computational protocol depends critically on the exact values of the fabricated beam splitters and phase shifters. Whilst we successfully achieve teleportation fidelities beyond what is classically possible, the error analysis shows we are limited by the quality of the fabricated circuit. At the end of this chapter we explore the consequences of this analysis in the context of more complex devices. We propose a new modular design for building quantum integrated photonic circuits. The new design includes embedded Bragg gratings for robust element-wise component characterisation and sufficient redundancy and tunability to mitigate inevitable fabricated deviations from design.

4.1 Quantum teleportation

The previous chapter discussed the first genuine quantum interference of three photons on a photonic chip and its application to Boson Sampling. Whilst these experiments relied upon high-visibility quantum interference between all three photons - the integrated circuit provided only the passively stable linear optical network to appropriately mix the optical modes. From chapter 2 we know that for more general forms of linear optical quantum computing, the integrated circuit must be designed to enact specific qubit logic gates and be capable of at least some active control. These additional constraints make on-chip realisations of qubit gates more challenging than the three-photon interference experiments discussed previously.

Integrated circuits implementing two-qubit logic gates are now well established [20,122,125,126,137,212]. However, a genuine three-qubit protocol has not yet been demonstrated on a photonic chip. In this chapter, we implement the quantum teleportation protocol on a photonic chip, thus realising the first on-chip three qubit protocol. Quantum teleportation is a good three-qubit protocol to try and implement, since whilst it is a genuine three-qubit scheme [1] - only *single* qubit tomography is required to verify its essential quantum operation. This significantly simplifies the experimental requirements compared to having to perform full three-qubit process tomography [213]. We also recall from chapter 2 that the teleportation scheme of Gottesman and Chuang elevated this protocol to one of central importance in scalable implementations of QIP [214]. Thus, whilst being comparatively simple to verify, the on-chip realisation of quantum teleportation represents a crucial advance in integrated quantum photonics.

[1] In the sense that there exists entanglement between all three qubits in the circuit. In particular, for teleportation, the three photons exist as a GHZ state before measurement takes place.

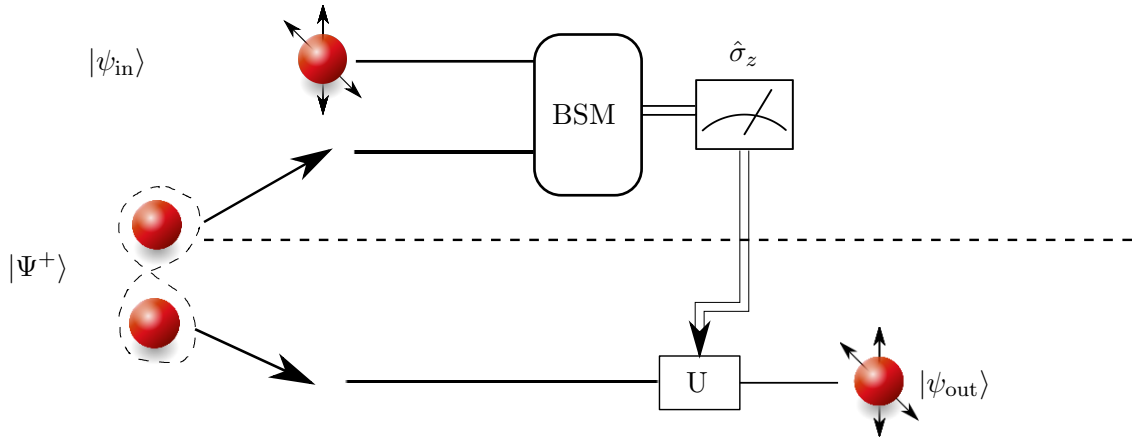


Figure 4.1: Quantum teleportation. An initial maximally entangled resource state coded on Q2 and Q3 is shared between Alice and Bob. It is assumed Alice and Bob can then only communicate classically (the dotted horizontal line delineates a barrier permeable only to classical communications). Alice makes a Bell State Measurement (BSM) between her initial unknown state coded on Q1 and her half of the entangled state Q2. She sends the result of this measurement to Bob, who performs a single-qubit unitary rotation on his half of the entangled state, Q3. It is shown in the text that through an appropriate choice of unitary rotation the final state on Bob's qubit is that of Alice's initial state. The quantum information coded on Q1 has been teleported to Q3.

4.1.1 The teleportation protocol

We will first explicitly layout the quantum teleportation protocol in its original form[71]. Later, we shall derive more general results which deviate slightly from this original prescription, but it will be helpful to always have in mind the protocol as laid out below. The prototypical quantum teleportation circuit, shown in figure 4.3(a), aims to transfer the quantum state Q1, held by Alice, to the target qubit, Q3, held by Bob. The qubit Q1 is initially prepared in an arbitrary quantum state,

$$|\psi_{\text{in}}\rangle_1 = \alpha\,|\mathbf{0}\rangle_1 + \beta\,|\mathbf{1}\rangle_1 \tag{4.1}$$

and the protocol begins by generating a maximally entangled resource, a Bell state, coded on qubits Q2 and Q3,

$$\left|\Phi^{+}\right\rangle_{2,3} = \frac{1}{\sqrt{2}}\left(|\mathbf{00}\rangle_{2,3} + |\mathbf{11}\rangle_{2,3}\right). \tag{4.2}$$

At this stage, the complete quantum state across all three qubits can be straightforwardly written down,

$$\begin{aligned} |\psi_{\text{tot}}\rangle &= \frac{1}{\sqrt{2}}\left(\alpha\,|\mathbf{0}\rangle_1 + \beta\,|\mathbf{1}\rangle_1\right) \otimes \left(|\mathbf{00}\rangle_{2,3} + |\mathbf{11}\rangle_{2,3}\right) \\ &= \frac{1}{\sqrt{2}}\left(\alpha\,|\mathbf{000}\rangle_{1,2,3} + \alpha\,|\mathbf{011}\rangle_{1,2,3} + \beta\,|\mathbf{100}\rangle_{1,2,3} + \beta\,|\mathbf{111}\rangle_{1,2,3}\right). \end{aligned} \tag{4.3}$$

We anticipate making a joint measurement on Q1 and Q2 and so by expressing the two-qubit computational basis states on Q1 and Q2 in the Bell basis,

$$\begin{aligned} |\mathbf{00}\rangle_{1,2} &= \frac{1}{\sqrt{2}}\left(\left|\Phi^{+}\right\rangle_{1,2} + \left|\Phi^{-}\right\rangle_{1,2}\right); \\ |\mathbf{01}\rangle_{1,2} &= \frac{1}{\sqrt{2}}\left(\left|\Phi^{+}\right\rangle_{1,2} - \left|\Phi^{-}\right\rangle_{1,2}\right); \\ |\mathbf{10}\rangle_{1,2} &= \frac{1}{\sqrt{2}}\left(\left|\Psi^{+}\right\rangle_{1,2} + \left|\Psi^{-}\right\rangle_{1,2}\right); \\ |\mathbf{11}\rangle_{1,2} &= \frac{1}{\sqrt{2}}\left(\left|\Psi^{+}\right\rangle_{1,2} - \left|\Psi^{-}\right\rangle_{1,2}\right), \end{aligned} \tag{4.4}$$

we can substitute into equation (4.3),

$$\begin{aligned} |\psi_{\text{tot}}\rangle &= \frac{1}{\sqrt{2}}\Big[\left|\Phi^{+}\right\rangle_{1,2}\left(\alpha\,|\mathbf{0}\rangle_3 + \beta\,|\mathbf{1}\rangle_3\right) \\ &+ \left|\Phi^{-}\right\rangle_{1,2}\left(\alpha\,|\mathbf{0}\rangle_3 - \beta\,|\mathbf{1}\rangle_3\right) \\ &+ \left|\Psi^{+}\right\rangle_{1,2}\left(\alpha\,|\mathbf{1}\rangle_3 + \beta\,|\mathbf{0}\rangle_3\right) \\ &+ \left|\Psi^{-}\right\rangle_{1,2}\left(-\alpha\,|\mathbf{1}\rangle_3 + \beta\,|\mathbf{0}\rangle_3\right)\Big]. \end{aligned} \tag{4.5}$$

It is now clear how the teleportation protocol proceeds. Performing a Bell State Measurement (BSM) on qubits Q1 and Q2 projects the target qubit Q3 onto the original input state via the non-local correlations of the entangled pair, modulo one of four possible local unitary rotations. For example, if the outcome of the BSM is $|\Phi^-\rangle$ then we obtain $|\psi\rangle_3 = \alpha\,|\mathbf{0}\rangle - \beta\,|\mathbf{1}\rangle$. Then applying a local corrective rotation, in this case the phase flip, $\hat{\sigma}_z$, yields,

$$|\psi\rangle_3 = \hat{\sigma}_z \left\langle \Phi^- | \psi_{\text{tot}} \right\rangle = |\psi_{\text{in}}\rangle\,, \tag{4.6}$$

thus completing the teleportation protocol. The specific unitary rotation needed to recover the initial state is uniformly distributed across the four possibilities but can be identified by two classical bits describing the outcome of the Bell measurement. It is straightforward to see that the four appropriate corrective unitary rotations are drawn from an element of $\{\mathbb{1}, \hat{\sigma}_z, \hat{\sigma}_x, \hat{\sigma}_x\hat{\sigma}_z\}$ given the BSM outcomes $\{|\Phi^+\rangle\,, |\Phi^-\rangle\,, |\Psi^+\rangle\,, |\Psi^-\rangle\}$ respectively.

It is clear why teleportation is of interest in quantum communication where quantum states need to be transferred over large distances. For example, we can distribute 'cheap' Bell pairs between two distant parties and then teleport the fragile quantum information using a robust classical communication channel. However, the interest in teleportation goes far beyond quantum communication. It explicitly demonstrates the equivalence of seemingly different resources in quantum mechanics - one unit of maximal bi-partite entanglement plus classical communication is equal to transmitting one qubit worth of information. This is quite remarkable when we consider that a Bell state contains no information about Q1 and yet only 2 bits of classical information are required for the perfect reconstruction of two complex amplitudes. Perhaps, then, it is no surprise that teleportation is at the heart of

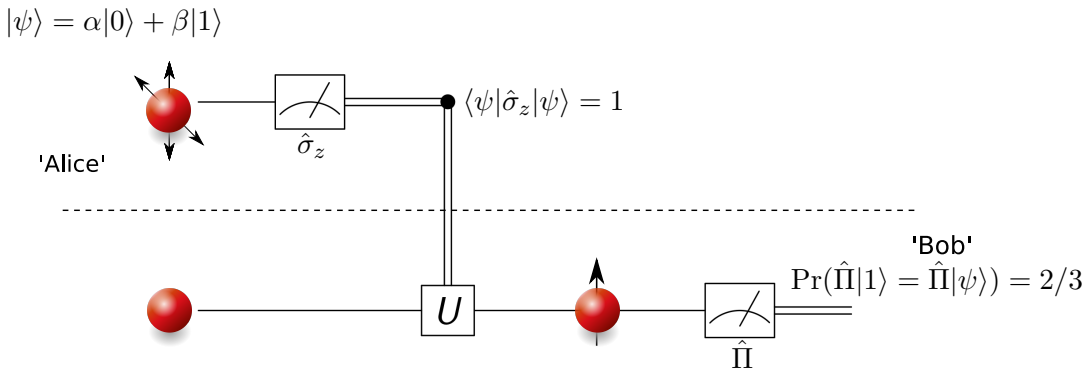


Figure 4.2: The optimal classical teleportation strategy. Bob wishes to reproduce the state belonging to Alice but can only communicate via a classical channel. One strategy which saturates the ultimate classical fidelity limit is shown. Alice makes a projective measurement of her state in the computational basis and sends the result to Bob. Bob prepares his state in one of the two computational eigenstates depending on the result received from Alice. If Bob proceeds to make any projective measurement on his newly prepared state he will obtain the same outcome had he made this measurement on Alice's original state with a probability of 2/3. This is the meaning of the classical teleportation fidelity limit.

quantum error correction schemes, allowing noisy, fragile quantum channels to be replaced by appropriate offline entangled resources and robust classical communication. Indeed, whereas entanglement appears to be a fundamental resource for quantum computing, teleportation can be viewed as the fundamental sub-routine [15].

The classical teleportation fidelity limit

The unit fidelity transfer of the quantum state from Q1 to that of Q3 is a unique feature of quantum teleportation [71]. One may ask what the best achievable fidelity is in the case when there is no prior shared entanglement between the two parties. In the most conservative case, we allow Alice to have perfect knowledge of the original state Q1. Then, transmitting only two bits of classical information allows Bob to reproduce the state with 87% fidelity [215]. However, a more common comparison leaves the initial state unknown, even to Alice. In other words, if there only exists a classical channel between Q1 and Q3, how faithfully

can the initial unknown state be reproduced on Q3 with only local operations and classical communication (LOCC)[216]?

To optimise this classical scheme, we consider the general strategy of performing an arbitrary measurement on Alice's state, followed by the preparation of Bob's state given classical information about the result of Alice's measurement outcome (which in general will require more than 2 bits). We start by observing that the most general measurement Alice can make of her state is described by a certain positive operator valued measure (POVM). In general, the dimension of the Hilbert space of this measuring system is arbitrary and may be much larger than that of the measured system. Without loss of generality, we consider this measurement to have L outcomes, labelled $l = 1, 2, ..., L$ each associated with a projector $|\epsilon_l\rangle\langle\epsilon_l|$, which are required to be complete, $\sum_{l=1}^{L} |\epsilon_l\rangle\langle\epsilon_l| = \mathbb{1}$.

Given the initial state is $|\psi\rangle$, the probability of obtaining the outcome l is given by $|\langle\psi|\epsilon_l\rangle|^2$. Alice sends the information l to Bob who then prepares a state $|\phi_l\rangle$ which is chosen so as to give the best chance of passing a test for the original state. Thus, Bob is left with the state

$$\rho_B = \sum_l |\langle\psi|\epsilon_l\rangle|^2 |\phi_l\rangle\langle\phi_l| . \tag{4.7}$$

We now compute the success of this protocol by computing the fidelity of Bob's state to that of the input state, averaged over all input states,

$$\begin{aligned} \mathcal{F}_{\text{av}} &= \int \mathrm{d}\psi \mathrm{tr}\left(\rho_A \rho_B\right) \\ &= \mathrm{tr} \int \mathrm{d}\psi \rho_A \sum_l |\langle\psi|\epsilon_l\rangle|^2 |\phi_l\rangle\langle\phi_l| \\ &= \mathrm{tr} \int \mathrm{d}\psi |\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi| \sum_l |\epsilon_l\rangle\langle\epsilon_l| \otimes |\phi_l\rangle\langle\phi_l| , \end{aligned} \tag{4.8}$$

where $\mathrm{d}\psi$ is the Haar measure. This integral can be evaluated using Schur's Lemma [217],

$$\int \mathrm{d}\psi \, |\psi\rangle \langle\psi| \otimes |\psi\rangle \langle\psi| = \frac{2}{d(d+1)} \hat{\pi}_{\mathrm{sym}}, \tag{4.9}$$

where d is the dimension of the initial state and $\hat{\pi}_{\mathrm{sym}}$ is the orthogonal projector onto the symmetric subspace. The second term in equation (4.8) can be evaluated by first expanding the state Bob prepares in the $|\epsilon_l\rangle$ basis,

$$|\phi_l\rangle = \sum_s \hat{B}_{ls} \, |\epsilon_l\rangle \tag{4.10}$$

where $\hat{B}_{ls}$ is a general linear transformation and satisfies $\sum_s \hat{B}_{ls}^{\dagger} \hat{B}_{ls} = \mathbb{1}$. Using the cyclic property of the trace and the symmetry of Bob's state we find,

$$\begin{aligned} \mathcal{F}_{\mathrm{av}} &= \frac{2}{d(d+1)} \mathrm{tr}(\hat{\pi}_{\mathrm{sym}} \sum_l |\epsilon_l\rangle \, |\epsilon_l\rangle \langle\epsilon_l| \, \langle\epsilon_l|) \\ &= \frac{2}{(d+1)}. \end{aligned} \tag{4.11}$$

For qubit states, $d = 2$ and the best achievable average fidelity is $\mathcal{F}_{\mathrm{av}} = 2/3$. This result is completely general and defines the limit for any classical channel [216]. Achieving an average fidelity above this limit requires access to a shared quantum resource which must have non-zero non-local correlations.

4.1.2 Previous implementations

Given its utility in distributing quantum information over long distances, it is perhaps no surprise that photonic experiments were the first to realise quantum teleportation [30,86]. The

group led by Francesco de Martini in Rome first teleported a qubit using only two photons - encoding the state to be teleported onto one half of the entangled pair [86]. Whilst this allows a complete deterministic Bell state measurement using only linear optics it is not suitable for concatenation. The group of Anton Zeilinger at Innsbruck then implemented the full three-qubit scheme but at the expense of an incomplete Bell state measurement [30]. Using non-linear optics, teleportation with a complete BSM was realised but at the expense of an extremely low probability of success $\sim 10^{-10}$ [218]. Alternatively, by encoding quantum information into the continuous variable (CV) quadrature states of light the full teleportation protocol can be realised using only linear optics [219]. However, it is impossible to generate maximally entangled CV states (squeezed states) with finite energy and the entanglement which is generated is extremely vulnerable to loss in the quantum channels. However, recent progress has seen the combination of discrete and continuous variable quantum states give rise to complete, deterministic photonic teleportation [220].

Since then, there has been a drive toward realising kilometre-scale teleportation as a means of distributing quantum information over long distances. The Geneva group led by Nicolas Gisin was the first to realise long-distance teleportation using time-bin encoded qubits distributed over 2 km of standard telecommunication fiber [221]. Polarisation qubits are subject to dephasing over long fiber lengths and so, despite one demonstration from the Zeilinger group in Vienna [222], all long-distance teleportation efforts now either use time-bin encoding over fiber or polarisation encoding over the low-loss, low-birefringence channel of free-space. Since the first long-distance free-space teleportation demonstration by the Hefei group of Jian-Wei Pan [223], experiments have become more sophisticated culminating in a recent demonstration of the full teleportation protocol over a 144 km free-space link [224].

Future work in this area now seeks to set-up quantum channels between earth and low-orbit satellites, enabling teleportation of quantum information over ever-greater distances [225].

There have also been numerous 'table-top' photonic teleportation experiments, including a 6-photon implementation of the original Gottesman and Chuang teleportation-based CNOT-gate [104] and a recent experiment using an entangled light emitting diode [226]. Deterministic teleportation has also been realised with matter-based systems including atomic qubits [227,228], superconducting qubits [229,230] and recently with nitrogen-vacancy centres in diamond [231]. However, there has yet to be an *integrated* photonic experiment which demonstrates the full teleportation protocol. Whilst there have been recent efforts toward this goal using a simplified quantum relay circuit [145], no integrated photonic experiments have yet been able to demonstrate teleportation, due to the difficulty in realising three photonic qubits on a sufficiently complex circuit [36]. In particular, integrated components require careful attention to fabricated chips that deviate from design and the effects of increased and potentially unbalanced propagation loss which is present in integrated circuits. Thus, a major motivating factor for this chapter is the experimental verification that integrated circuits continue to perform well as their complexity increases.

4.2 The circuit

The integrated circuit used to perform on-chip teleportation is identical to the one used in the previous chapter. However, we now consider the six central modes to define three dual-rail qubits and recognise that each of the optical components now implements a specific quantum logic operation. To see how we implement the teleportation protocol using this circuit we must first decompose the CNOT-gates into single-qubit Hadamard and two-qubit C-PHASE

gates[232], as shown in figure 4.3(b). We also use the concatenated C-PHASE gate scheme of Ralph[70] where the effective two-photon interactions are induced via post-selection of one photon being detected in each of the three qubit modes (see figure 2.8). This occurs with probability 1/27, which sets the overall success probability of the circuit. We can now identify the beam splitters responsible for each circuit element as highlighted in figure 4.3. In order to for the teleportation circuit to function correctly, these beam splitters must have specific reflectivities of either 1/2 or 1/3. Requiring such specific beam splitting ratios is one reason why realising on-chip qubit logic is more challenging than equivalent quantum interference or Boson Sampling experiments.

4.2.1 Programmable unitary operations

In the previous chapter we made use of the on-chip phase shifters to characterise the phase offsets inside the interferometers. We now utilise these elements to perform programmable unitary operations in order to prepare and measure different quantum states. Whilst the active feed-forward corrective rotation could in principle be implemented using the on-chip phase shifters ϕ_2 and θ_2, in reality these shifters do not have sufficient bandwidth to make this feasible.

Indeed, performing fast, feed-forward operations in quantum photonic experiments remains a key challenge for the field. A common solution is to use long off-chip fibre delay lines[233] to which integrated photonic circuits based on silica are particularly well-suited since they are exceptionally well mode-matched to standard silica-based fibres[140]. Ultimately, such methods will not enable scalable quantum information processing, however. Directly implementing feed-forward in standard integrated optics chips is a major open

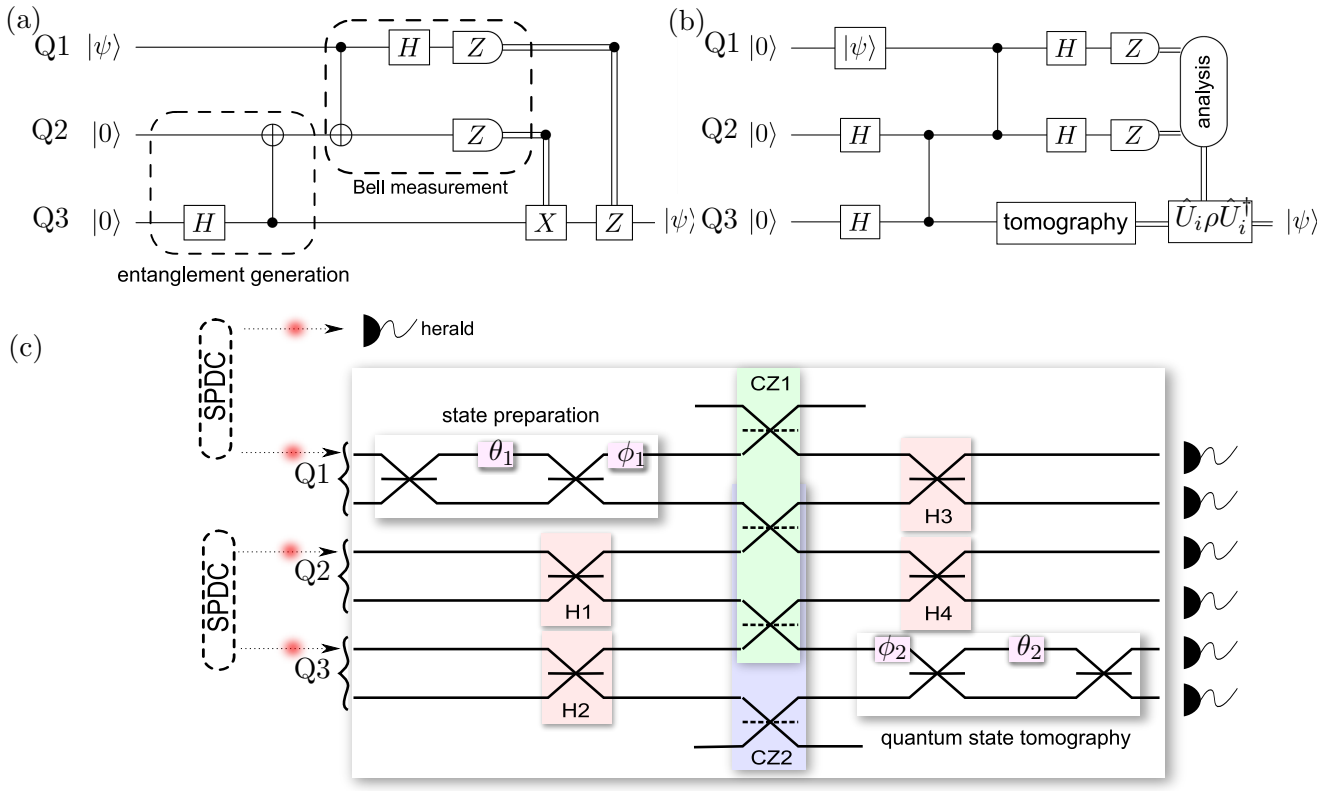


Figure 4.3: Photonic chip realisation of teleportation. (a) Circuit diagram of a general quantum teleportation scheme. (b) In our experiment we replace the two CNOT-gates with two C-PHASE-gates and additional local Hadamard operations. The results of the Bell measurement are used in post-processing as part of quantum state tomography to recover the the teleported qubit. (c) In the on-chip realisation, three qubits are encoded using dual-rail logic in a silica-on-silicon integrated chip. Local Hadamard operations (H1 to H4) are performed using beam splitters of reflectivity 1/2 (solid lines) and the two cascaded C-PHASE gates (CZ1, CZ2) are implemented using four beam splitters of reflectivity 1/3 (dashed lines). State preparation and tomography are performed on chip using thermo-optically controlled phases θ and ϕ.

challenge for integrated-optics QIP, requiring ultrafast on-chip photon detection, classical electronics, feed-forward and phase modulation, all operating at bandwidths approaching the terrahertz regime.

Progress toward this goal is ongoing. Integrated high-speed photon detectors[142], as well as fast phase control based on strain-optic[234] and electro-optic[141] effects have been recently demonstrated. Efficient on-chip optical delay lines, are also now a possibility thanks to advances in lithographic fabrication techniques[235]. Nevertheless, achieving fast enough

speeds for practical use in scalable circuits remains out of reach of current technologies and presents a major technical challenge. Instead, for the purposes of this demonstration, we use these slow thermo-optic phase shifters to perform quantum state tomography (QST) of the teleported qubit Q3 and replace the corrective step by a numerical rotation in post-processing.

We recall from section 2.2.1 that the homomorphism between the group of 2×2 unitary operations, $\mathsf{SU(2)}$, and the rotation group in 3 dimensions, $\mathsf{SO(3)}$, allows us to think of arbitrary single-qubit operations from the computational state as a sequence of two rotations around orthogonal axes,

$$\hat{U}(\theta, \phi) = \hat{R}_z(\phi)\hat{R}_y(\theta) \tag{4.12a}$$

$$= e^{-i\phi\hat{\sigma}_z/2} e^{-i\theta\hat{\sigma}_y/2} \tag{4.12b}$$

These rotations are physically realised in the dual-rail encoding as tunable phase shifters and tunable beam splitters (of reflectivity $\eta = \cos^2(\vartheta)$) as demonstrated in figure 2.1. The former is implemented by selectively retarding the phase of a single mode using a thermo-optic heater deposited directly above the waveguide. The latter can be accomplished by embedding such a phase shifter between two fixed beam splitters in a Mach-Zehnder configuration as shown in figure 4.4. This effectively transforms us into the x-basis where an adjustable phase shifter on one of the rails now performs effective x-rotations. However, in our real experiment the beam splitters which are used to transform us to the rotated basis are not ideal ($\eta \neq 1/2$) or equal to each other (i.e. $\eta_1 \neq \eta_2 \rightarrow \vartheta_1 \neq \vartheta_2$). In order to see the effect of using non-ideal beam splitters in this configuration we write out the transformation

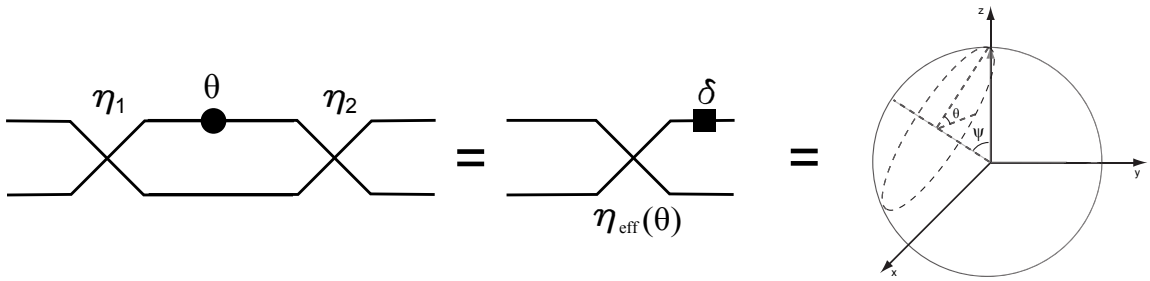


Figure 4.4: A tunable beam splitter based on a Mach-Zehnder Interferometer. To achieve tunable beam splitter rotations we embed a tunable phase shifter between two fixed beam splitters. In the ideal case that both $\eta_{1,2} = \cos^2(\vartheta_{1,2}) = 1/2$, then η_{eff} performs rotations around a great circle on the Bloch sphere with $\psi = \pi/2$ as seen on the right. In the case of matched beam splitters, $\eta_1 = \eta_2$, the state starts along the z-axis and rotates around the dotted circle. If $\eta_1 \neq \eta_2$ then the state is first rotated by an amount $(\vartheta_1 - \vartheta_2)$ around the y-axis. A phase offset δ appears at the output depending on the exact beam splitter reflectivities used and the angle of rotation. In this experiment this additional phase is subsumed into the action of the phase shifters ϕ_1 and ϕ_2.

explicitly:

$$\hat{R}_n(\theta) = \hat{R}_y(\vartheta_1)\hat{R}_z(\theta)\hat{R}_y(\vartheta_2) \tag{4.13a}$$

$$= e^{-i\frac{\vartheta_1}{2}\hat{\sigma}_y} e^{-i\frac{\theta}{2}\hat{\sigma}_z} e^{-i\frac{\vartheta_2}{2}\hat{\sigma}_y} \tag{4.13b}$$

$$= e^{-i\frac{\theta}{2}[\hat{R}_y(\vartheta_1)\hat{\sigma}_z\hat{R}_y(-\vartheta_1)]} e^{-i\frac{(\vartheta_1-\vartheta_2)}{2}\hat{\sigma}_y}, \tag{4.13c}$$

where we have used the Baker-Campbell-Hausdorff lemma in the last line. Using the relation $\hat{R}_y(\vartheta) = e^{-i\vartheta\hat{\sigma}_y/2} = \cos(\vartheta/2)I + i\sin(\vartheta/2)\hat{\sigma}_y$ and the Pauli matrix commutation relation $[\sigma_y, \sigma_z] = 2i\sigma_x$ we can express $\hat{R}_y(\vartheta_1)\hat{\sigma}_z\hat{R}_y(-\vartheta_1) = \cos(\vartheta_1)\hat{\sigma}_z - \sin(\vartheta_1)\hat{\sigma}_x$. Now we can write down the complete rotation enacted by this MZI arrangement,

$$\hat{R}_{tot}(\theta) = e^{-i\theta/2(\cos(\vartheta_1)\hat{\sigma}_z - \sin(\vartheta_1)\hat{\sigma}_x)} e^{-i(\vartheta_1-\vartheta_2)\hat{\sigma}_y}. \tag{4.14}$$

We can immediately see that in the ideal case of using two fixed 50:50 beam splitters (i.e. $\vartheta_{1,2} = \pi/2$) then $\hat{R}_{tot}(\theta) = e^{-i\theta\hat{\sigma}_x/2}$ and we can effect the desired controllable rotations

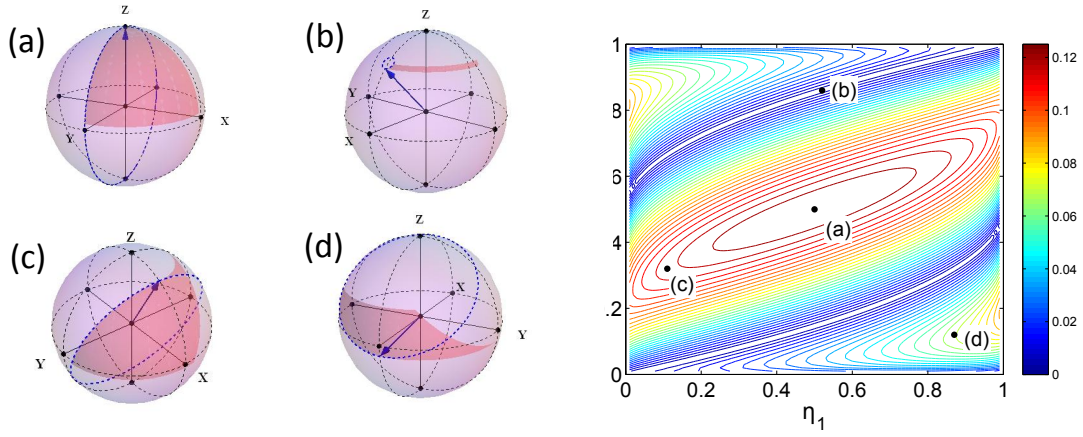


Figure 4.5: Accessible states on the Bloch sphere as a function of phase shifter tuning range and beam splitter reflectivity. The shaded regions on the spheres highlight the space of accessible states. The blue arrow denotes the starting point when both phase shifters are set to zero. As the central phase shifter is scanned through its range we trace out the blue dotted lines on the spheres. The total fraction of the Bloch sphere that can be accessed in this way is shown in the contour plot. The scale ranges from 0 (no accessible states) to 0.125 (can access 1/8th of the total Bloch sphere) since we restrict the tuning range of the two phase shifters to be $\pi/2$ so that they correspond to the achievable states in our experiment. (a-d) The reflectivities of the beam splitters which make up the MZI drastically affect the range of states which may be prepared. The contour plot shows the fraction of possible states which may be prepared or measured as a function of beam splitter reflectivity. For example, for point (a), the two beam splitters have exactly 50:50 reflectivity which allows us to access one octant (0.125) of the Bloch sphere as expected with the restricted phase shifter tuning range.

around the x-axis by adjusting a phase shifter on one of the rails.

Using beam splitters which are not 50:50 (but are matched such that $(\vartheta_1 = \vartheta_2)$ changes the effective axis of rotation that is implemented. This has the immediate effect of restricting the range of states on the Bloch sphere which can be generated since we are no longer rotating around a great circle. In fact, the radius of rotation will shrink by a factor of $\sin(\vartheta_1)$. If the beam splitters are not matched to one another such that $\vartheta_1 \neq \vartheta_2$ then the initial state is first also rotated around the y-axis by a fixed amount $(\vartheta_1 - \vartheta_2)$. Using this rotation picture it becomes clearer to visualise the effect of non-ideal components on the range of states we can implement, as shown in figure 4.5

4.3 Circuit characterisation

In chapter 3 we used an element-wise characterisation method to measure the beam splitting ratios and phase shifter parameters for each of the 17 fabricated on-chip components. Whilst this was successfully used to model the three-photon quantum interference it was found that better results could be obtained in the Boson Sampling experiment using the method outlined in reference [191]. Although this method was able to account for the on-chip losses which remain inaccessible using the element-wise scheme [127,139], it treats the circuit as a single unit and does not reveal any information about individual circuit components.

For this experiment, however, it is vital to be able to separate out specific circuit components to be able to characterise the action of the state preparation and state tomography stages of the chip. This is only possible using an element-wise characterisation method. Previous experiments using simpler two-qubit protocols, have been able to prepare and measure states without individually characterising circuit elements [122]. However, as integrated circuits become more complex it will be necessary to have more robust characterisation procedures in place in order to best mitigate inevitable fabricated deviations from design. For example, if we know which precise components are responsible for which errors, it is possible to then correct for these using active control (see section 4.8), or adapt the analysis routine (see section 4.5) to increase the achievable process fidelity.

We consider the results of our element-wise characterisation in the context of the above discussion on programmable unitary operations. Using the beam splitter ratios and phase shifter parameters for the state preparation and state tomography stages of the chip we can work out the range of states we can prepare and measure.

Experimentally prepared input states

The two beam splitters in the state preparation Mach-Zehnder have reflectivities $\eta_1 = 0.54$ and $\eta_2 = 0.47$. Although this in principle allows for a wide range of states to be prepared, we are limited by a $\approx \pi/2$ tuning range in the phase shifters θ_1 and ϕ_1. This restricts us to preparing states over only a single octant of the Bloch sphere as seen in figure 4.5(a). In order to demonstrate the teleportation of any general input state, we must prepare a set of three linearly independent states to send into our circuit. The phase settings used to prepare each of the three states along with the quantum state generated by each setting are summarised in figure 4.6.

	PS Settings		$\|\psi_{\text{in}}\rangle = \alpha\|H\rangle + \beta\|V\rangle$	
Input	θ_1	ϕ_1	α	β
$\|V'\rangle$	0.435	3.48	-0.154-0.101i	-0.983
$\|D'\rangle$	-1.12	4.72	-0.558-0.196i	-0.806
$\|R'\rangle$	-1.12	3.72	-0.005-0.592i	-0.806

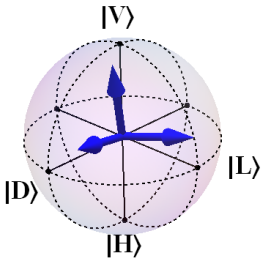


Figure 4.6: Experimentally prepared input states. Three linearly independent input states are prepared by setting the phase shifters θ_1 and ϕ_1 appropriately. Combined with the characterisation of the beam splitters which form the MZI we can work out the quantum states generated on Q1. These three states are shown on the Bloch sphere for clarity.

Measurement Operators

Similarly, state tomography of the output state requires us to measure the state in at least three linearly independent bases[236,237]. Avalanche photodiodes (APDs) monitor the output ports of the waveguide and, in the ideal case of zero loss and no higher-order photon emission, performs projective measurements on the qubit state in the computational basis. Measurement along the three orthogonal axes of the Bloch sphere minimises the error in the

state tomography routine[236] and so we set the circuit to perform measurements as close to these three axes as possible.

Measurement in these linearly independent bases is achieved using an analogue of the system used to prepare the initial quantum states. We use a programmable unitary to rotate the action of the APDs to perform projective measurements from the computational- to any other desired basis. Exactly the same analysis follows in that the range of measurement operators we can implement is determined by the fabricated beam splitting ratios which make up the MZI on Q3 and the tuning range of the phase shifters. In this case the relevant beam splitters have reflectivities of 0.35 and 0.55. The phase shifter settings and measurement operators which are realised are presented in figure 4.7.

	PS Settings		$\hat{\Pi} = \begin{pmatrix} \lvert\alpha\rvert^2 & \alpha\beta^* \\ \alpha^*\beta & \lvert\beta\rvert^2 \end{pmatrix}$	
Operator	θ_2	ϕ_2	α	β
1	-5.65	2.55	-0.311+0.208i	0.927
2	-4.63	2.35	-0.5261+0.533i	0.663
3	-5.65	3.55	-0.343-0.149i	0.927

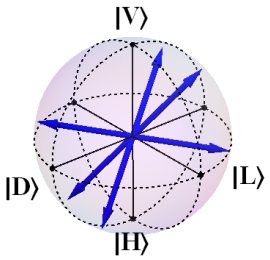
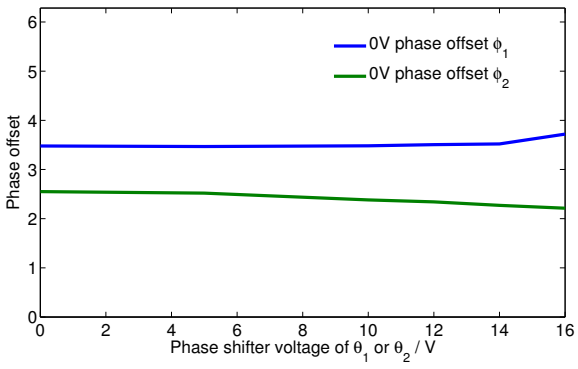


Figure 4.7: Measurement operators. Three linearly independent measurement bases are implemented by setting the phase shifters θ_2 and ϕ_2 appropriately. Combined with the characterisation of the beam splitters which form the MZI we can work out the measurement operators on Q3, shown on the Bloch Sphere as three linearly independent measurement axes.

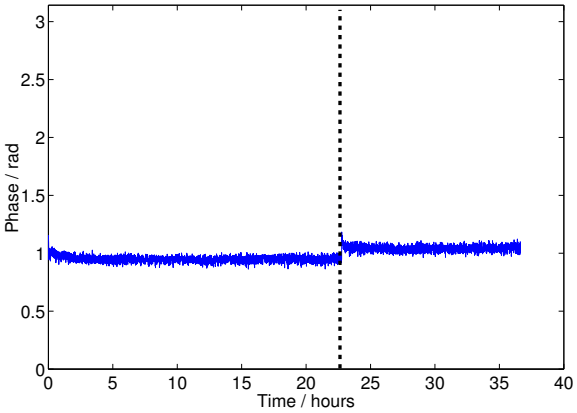
4.3.1 Phase shifter cross-talk

The two pairs of phase shifters,(θ_1, ϕ_1, and θ_2, ϕ_2) responsible for implementing the programmable unitary operations are only 1 cm apart on the chip. The relatively small thermo-optic coefficient of silica necessitates using high temperatures to enact the required phase shifts. In addition, the nichrome heaters had a limited maximum power rating. Operating beyond

this rating for extended periods can cause the thin films to burn out and render the phase shifter useless. In order to prepare input states and measurement operators which were as close to orthogonal as possible, the phase shifters had to be run close to this upper limit. This had the consequence of causing some cross-talk between nearby interferometers. For example, as shifter θ_1 was tuned toward its maximum power rating we see a small shift in the phase difference between the arms of the MZI in which ϕ_1 is embedded. This is seen as a small upturn in the blue line in figure 4.8(a). To account for this effect, at each of the two settings of θ_1 we performed a full scan of ϕ_1 to obtain two different zero-voltage phase offsets for the MZI in which ϕ_1 was embedded. This procedure was repeated for the phase shifters ϕ_2 and θ_2 on the measurement operator stage of the chip. No cross-talk between the phase shifters on Q1 and Q3 was noted during these characterisation runs, suggesting that it is only a localised heating effect. These effects have been included in the tables above when calculating the correct phase shifts for a given set of voltages applied to the chip.



(a) Phase shifter cross talk



(b) Phase shifter repeatability

Figure 4.8: Non-ideal behaviour of phase shifters. (a) Effect of cross-talk on phase shifters in close proximity. As the phase shifters θ_1 and θ_2 are run with higher voltages - the heat dissipation causes a slight change in the phase of neighbouring interferometers on the order of 10° as seen in the shift in zero voltage phase offset in the interferometers in which ϕ_1 and ϕ_2 are embedded. (b) Continuously monitoring the phase of ϕ_1 with an applied voltage of 20 V. At the dotted line, the voltage was switched to 0 V for 1 minute and then bought back to 20 V.

The repeatability and stability of the phase shifters was also measured. The temperature of the silicon substrate was temperature stabilised using a thermo-electric element at 19° C. Applying 20 V to ϕ_1 we record the ratio of output intensities on the Q1 output modes and calculate the phase inside the MZI as a function of time. This was performed using the loss-independent ratiometric method with an automated switching system to alternately couple light into both input ports. The measured phase as a function of time is plotted in figure 4.8b and shows a stability of < 35 mrad over 24 hours. Switching the voltage between 0 V and 20 V we find a repeatability error of < 90 mrad as shown at the step at the dotted line. These errors are comparable to the uncertainty due to the characterisation and are added in to the experimental error on the calculated fidelities.

4.3.2 Unbalanced detector losses

The element-wise characterisation of the beam splitters can also be used to determine the relative coupling and detector losses between two output modes. In order to accurately reconstruct the quantum state on Q3 we must quantify the different propagation losses suffered by each of the two output modes. With the output v-groove array butt-coupled and the APDs connected we sent single photons into the circuit on the upper mode of Q2. From the element-wise characterisation, we could predict the relative intensities of the photons emerging from each of the ports of Q3. Comparing with the measured counts we inferred a relative loss between the modes of 0.9 — the bottom mode having 90% of the transmission of the upper mode. This was used to correct the raw collected click statistics before we performed quantum state tomography.

4.4 Circuit simulation

The circuit characterisation allows us to simulate the expected behaviour of our actual circuit and to verify we will be able to show quantum teleportation beyond the classical limit. In chapter 3 we used a linear optical network simulation to predict the visibility of multi-photon correlation functions and compared it to our results to confirm genuine three-photon interference. However, extending this model to account for higher-order multi-photon emission and loss soon becomes computationally expensive to run. This occurs because we calculate the full output state vector which grows exponentially in the number of photons and modes. However, evaluating each of these possible terms is not only computationally demanding but is unnecessary. We eventually want to find the expected click statistics of successful teleportation events by applying the appropriate projection operators to the output state. We anticipate which terms will contribute to this projection knowing that we only care about those terms which contain at least one photon in each of the three qubit rails. This restriction dramatically cuts down the number of terms which we must calculate. We show below how the matrix permanent is used to incorporate the effects of multi-photon emission and loss into the simulation whilst keeping the computational time manageable. We then use this new model to explore the effects of experimental imperfection on the expected teleportation fidelity.

4.4.1 Simulation based on matrix permanents

We consider a general $L \times L$ unitary linear network, $\hat{U}_{\text{circ}}$. There are a total of L modes in the system with the first M accessible to the experimenter and the remainder describing loss modes.

A vector $|T\rangle$ describes the input distribution of Fock states (total photon number, N) among the M accessible modes,

$$|\psi_{\text{in}}\rangle = \bigotimes_{i=1}^{M} \frac{(\hat{a}_i^\dagger)^{T_i}}{\sqrt{T_i!}} |0\rangle \otimes \bigotimes_{i=M+1}^{L} |0\rangle_i \,, \tag{4.15}$$

and in general this state is transformed into an entangled state at the output of the circuit which can be written as a superposition of D terms in which the N photons can be distributed over the L modes,

$$|\psi_{\text{out}}\rangle = \sum_{i=1}^{D} c_i \left(\prod_{j=1}^{L} \frac{(\hat{b}_j^\dagger)^{S_{ij}}}{\sqrt{(S_{ij})}} \right) |\text{vac}\rangle \tag{4.16}$$

where S_{ij} gives the number of photons in output mode j in the i^{th} superposition term. The amplitudes for each of these possible terms c_i can be computed by evaluating a matrix permanent,

$$c_i = \frac{\text{Per}(\hat{U}^{(\boldsymbol{S}_i, \boldsymbol{T})})}{\sqrt{\prod_{j=1}^{M} S_{ij}! \prod_{i=1}^{M} T_i!}}, \tag{4.17}$$

where $\mathbf{U}^{(\mathbf{S_i},\mathbf{T})}$ is a sub-matrix of $\mathbf{U}_{\text{circ}}$ formed by keeping the columns listed by $\boldsymbol{S}_i$ and the rows listed by $\boldsymbol{T}$. The method used in chapter 3 is computationally equivalent to evaluating all of the terms in equation (4.16). However, we choose to only calculate terms where there is at least one photon in each qubit as these are the only terms that contribute to our statistics.

In the ideal case of three Fock state inputs there are only 8 output terms we are interested in. Since we post-select on detecting three photons, any terms containing a photon in a loss mode does not contribute to a successful event and we are left only with the 4 possible

combinations of successful Bell state measurements - each allowing a photon to be detected in one of the two computational modes of Q3. Using an SPDC source, however, means there is some probability of injecting more than three photons into the circuit. Now terms in which these 'extra' photons appear in the loss modes will contribute an error to the click statistics. Since we also use APD detectors, which cannot resolve photon number, terms in which more than one photon occupies a relevant mode will also cause additional error counts. To account for this 'bucket' nature of our detectors we formulate the POVM elements on mode j as:

$$\hat{\Pi}_j^{(\mathrm{APD})}(\text{'no click'}) = \left(\bigotimes_{i=1}^{j-1} \mathbb{1}_i\right) \otimes \left(\sum_{n=0}^{N} (1-\eta_j)^n \left|n\right\rangle_j\!\left\langle n\right|_j\right) \otimes \left(\bigotimes_{i=j+1}^{L} \mathbb{1}_i\right) \tag{4.18a}$$

$$\hat{\Pi}_j^{(\mathrm{APD})}(\text{'click'}) = \mathbb{1} - \hat{\Pi}_j^{(\mathrm{APD})}(\text{'no click'}) \tag{4.18b}$$

where η_j is the loss associated with the detector on mode j. These projectors are applied to the partially calculated output state to obtain the probabilities of coincident clicks being recorded in our experiment.

Considering only the 6 accessible modes in our experiment we represent the perfect Fock state input to our chip as $|\boldsymbol{T}\rangle = |011010\rangle$ where it is implied the remaining $L-6$ loss modes all start in the vacuum state. The dominant error term corresponds to one of the SPDC sources emitting two pairs of photons whilst the other emits a single pair. This corresponds to the input states $|\boldsymbol{T}\rangle = |022010\rangle$ and $|\boldsymbol{T}\rangle = |011020\rangle$ depending on which crystal emits the two pairs. The click probabilities calculated for these input states are added to those calculated using the ideal three-photon input. We must experimentally determine how often our SPDC sources emits these higher order terms to accurately weight these extra

contributions relative to the ideal three-photon case. The state emitted by our two sources that includes the lowest order terms that contribute to our post-selected data is

$$|\psi_{\text{SPDC}}\rangle = \mathcal{N}\left(|1111\rangle + \lambda^2(|2211\rangle + |1122\rangle)\right), \tag{4.19}$$

where $\mathcal{N}$ is a normalisation constant. The $|1111\rangle$ term is the desired emission for the teleportation experiment, with three of the photons coupled into the circuit and the other used as a herald. The latter two terms are the lowest order errors from higher order SPDC emission. We measure $\lambda^2 \approx 0.03$ for our SPDC sources (see section 3.1.1). We use this figure in our simulation to predict the effect of these higher order terms on the achievable fidelity of the on-chip teleportation.

Partial distinguishability

Photon distinguishability will reduce the achievable teleportation fidelity. We introduce the distinguishability parameter, t, such that the input state to our circuit can be written,

$$|\psi_{\text{ini}}\rangle = \hat{a}^\dagger_{1,t_1}\hat{a}^\dagger_{2,t_2}\hat{a}^\dagger_{3,t_3}|\text{vac}\rangle = |t_1, t_2, t_3\rangle\,. \tag{4.20}$$

We are clearly hoping that our photons are not completely distinguishable such that, in general, we have $\langle t_i|t_j\rangle \neq 0$. However, for calculations, it is convenient to work with states which satisfy $\langle t_i|t_j\rangle = \delta_{i,j}$ - i.e. states which are either completely identical or completely distinguishable[238]. A Gram-Schmidt orthogonalisation provides us with a such a basis, $\{t'_1, t'_2, t'_3\}$, in which to expand our input state so we can write:

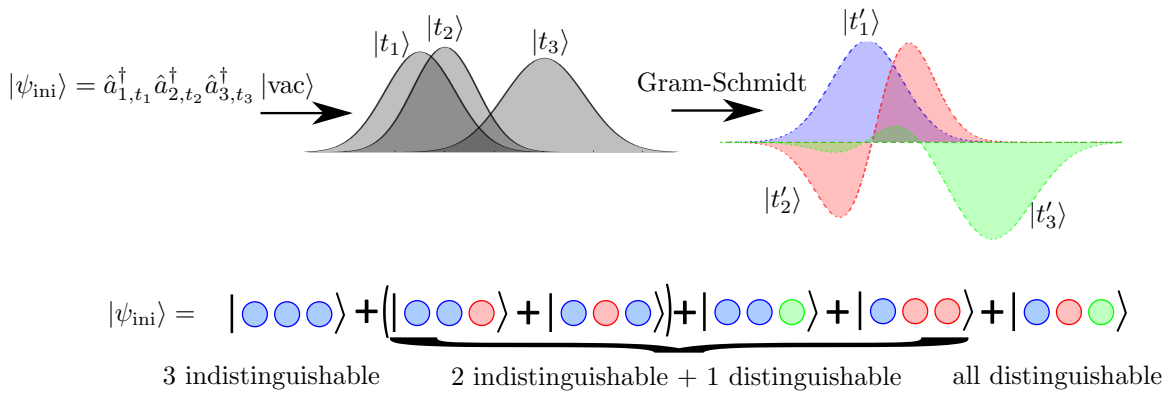


Figure 4.9: Partial distinguishability. The input state to our circuit is composed of three partially distinguishable photons represented by the overlap of the gray wavefunctions (top left). A Gram-Schmidt orthogonalisation (top right) allows us to represent this input state as a superposition of six terms consisting of fully distinguishable or indistinguishable states (bottom). The weighting of these terms is determined by the mutual distinguishability of each photon with the other two. The two terms in brackets at the bottom give rise to indistinguishable output states and hence the interference of these terms must be included in the calculation.

$$|t_1\rangle = |t'_1\rangle \tag{4.21a}$$

$$|t_2\rangle = |t'_1\rangle \underbrace{\langle t_1|t_2\rangle}_{=:c_{2,1}} + |t'_2\rangle \underbrace{\sqrt{1-|\langle t_1|t_2\rangle|^2}}_{=:c_{2,2}} \tag{4.21b}$$

$$|t_3\rangle = |t'_1\rangle \underbrace{\langle t_1|t_3\rangle}_{=:c_{3,1}} + |t'_2\rangle \underbrace{\langle t_2|t_3\rangle}_{=:c_{3,2}} + |t'_3\rangle \underbrace{\sqrt{1-|\langle t_1|t_2\rangle|^2-|\langle t_2|t_3\rangle|^2}}_{=:c_{3,3}} \tag{4.21c}$$

$$\rightarrow |\psi_{\text{ini}}\rangle = \left[\hat{a}^\dagger_{1,t'_1}(c_{2,1}\hat{a}^\dagger_{2,t'_1} + c_{2,2}\hat{a}^\dagger_{2,t'_2})(c_{3,1}\hat{a}^\dagger_{3,t'_1} + c_{3,2}\hat{a}^\dagger_{3,t'_2} + c_{3,3}\hat{a}^\dagger_{3,t'_3})\right]|\text{vac}\rangle \tag{4.21d}$$

These six terms are shown in figure 4.9 and are weighted by the three independent distinguishability parameters $c_{2,1}$, $c_{3,1}$, and $c_{3,2}$ (the remaining terms are fixed by requiring the states be correctly normalised).

In our experiment, the dominant source of distinguishability is given by the spectral mismatch of the photons. As discussed in chapter 3 we must consider the distinguishability between the two narrow-band photons (α_{nn}) and the narrow- and broad-band photons (α_{nb}). Photons in qubits 1 and 2 are narrow-band whilst the photon used for qubit 3 is broad-band.

Hence, we consider only two distinguishability parameters and identify $c_{2,1} = \alpha_{nn} = 0.984$ and $c_{3,1} = c_{3,2} = \alpha_{nb} = 0.969$ as calculated in chapter 3. In general, these coefficients, $\langle t_i | t_j \rangle$, are complex and contain information about the phase difference between the three particles. It can be shown [239] that for the most general three-photon pure-state input there are four independent parameters which characterise the interference — $|\langle t_1|t_2\rangle|^2$, $|\langle t_1|t_3\rangle|^2$, $|\langle t_2|t_3\rangle|^2$ and $\mathcal{R}e(\langle t_1|t_2\rangle \langle t_1|t_3\rangle \langle t_2|t_3\rangle)$.

To calculate the probability of an output event we must sum over all permutations of distinguishable photons which can occur for a given output distribution. For example, if we wish to compute the probability of the output state $|S\rangle = (1, 0, 1, 0, 1, 0)$ occurring, we must consider the contribution of all ways this can occur such as $|S\rangle = (t'_1, 0, t'_1, 0, t'_1, 0)$ or $|S\rangle = (t'_1, 0, t'_2, 0, t'_2, 0)$. There are 18 such permutations which are available for our circuit and we must sum over all of these combinations:

$$P_{\text{dist}}(S) = \sum_{\psi_{\text{final}}} |\langle \psi_{\text{final}}(S)| U |\psi_{\text{ini}}\rangle|^2. \tag{4.22}$$

The squared probability amplitudes in the above expression are now calculated using an incoherent sum of the permanents of $(N-1) \times (N-1)$ or $(N-2) \times (N-2)$ matrices when there are 1 or 2 distinguishable particles respectively. We see that this computation becomes ever more complex as the number of particles increases since the number of terms which now contribute to the output grows as $n!$. Indeed, modelling the effects of partial distinguishability of many particle systems is a rich theoretical field [238,240].

For simplicity, because both the effect of higher order terms and the photon distinguishability are small in our experiment we consider only the first order effect where *either*

there was higher order emission of identical photons or there were exactly three single photons, one of which was distinguishable from the other two. In this sense we do not include the effect when the higher-order terms are also distinguishable or when there is more than one distinguishable photon.

4.5 Corrective unitary rotations

4.5.1 Teleportation with non-maximally entangled states

In general, the effect of circuit deviations away from the designed parameters leads to the generation of non-maximally entangled states. We quantify the effect of using both non-maximally entangled (NME) resource states and NME two-qubit measurements on the achievable teleportation fidelity in our circuit [241]. The derivation is similar to that performed in section 4.1.1 but now allows general two-mode entangled states and measurements.

Let the initial state Alice wishes to teleport to Bob be $|\psi_{\text{in}}\rangle$, which we will expand in the computational basis $|i\rangle$,

$$|\psi_{\text{in}}\rangle_1 = \sum_i I_i \, |i\rangle_1 \,, \qquad i = \{0, 1\} \tag{4.23}$$

and we let the NME two-qubit state shared by Alice and Bob now take the general form

$$|E\rangle_{2,3} = \sum_{j,k} E_{j,k} \, |j\rangle_2 \, |k\rangle_3 \,, \qquad j, k = \{0, 1\} \tag{4.24}$$

where the coefficients E_{jk} define a matrix $\boldsymbol{E}$. We can write this state as a superposition of

two orthogonal basis vectors using the Schmidt decomposition. Taking the singular value decomposition (SVD) of $\boldsymbol{E}$ we can write $\boldsymbol{E} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}$, for unitary matrices $\boldsymbol{U}$ and $\boldsymbol{V}$ and the diagonal matrix $\boldsymbol{D}$. Then $|E\rangle_{2,3} = \sum_{j,k,l} U_{jl} D_{ll} V_{lk} |j\rangle_2 |k\rangle_3$ and we can define the rotated bases $|l'\rangle = \sum_j U_{jl} |j\rangle$ and $|l''\rangle = \sum_k V_{lk} |k\rangle$ and re-write

$$|E\rangle_{2,3} = \sum_l D_{ll} |l'\rangle_2 |l''\rangle_3 . \tag{4.25}$$

As before, we can now write the initial state of the composite system of Alice and Bob as

$$|\psi_{\text{tot}}\rangle_{1,2,3} = |\psi_{\text{in}}\rangle_1 \otimes |E\rangle_{2,3} = \sum_{i,l} I_i D_{ll} |i\rangle_1 |l'\rangle_2 |l''\rangle_3 \tag{4.26}$$

and allow Alice to make a general 2-qubit projective measurement on the states Q1 and Q2. There are four possible outcomes for this measurement, which in the computational basis are, $|\mathbf{00}\rangle, |\mathbf{01}\rangle, |\mathbf{10}\rangle$ and $|\mathbf{11}\rangle$. Each measurement outcome projects the 2-qubit state into one of four entangled eigenstates depending on the measurement basis, which we write in the parameterised form,

$$\left|B^{(1)}\right\rangle = M(|\mathbf{00}\rangle + m |\mathbf{11}\rangle); \quad \left|B^{(2)}\right\rangle = M(m^* |\mathbf{00}\rangle - |\mathbf{11}\rangle); \tag{4.27a}$$

$$\left|B^{(3)}\right\rangle = N(|\mathbf{01}\rangle + n |\mathbf{10}\rangle); \quad \left|B^{(4)}\right\rangle = N(n^* |\mathbf{01}\rangle - |\mathbf{10}\rangle). \tag{4.27b}$$

for $n, m \in \mathbb{C}$ and where $M = 1/\sqrt{1+|m|^2}$ and $N = 1/\sqrt{1+|n|^2}$ are normalisation factors. In this form it is convenient to see that when $n = m = 1$ we perform maximally entangling Bell state measurements and when $n = m = 0$ we make completely separable measurements in the computational basis. These states can be written compactly as

$\left|B^{(r)}\right\rangle = \sum_{i,j} B_{ij}^{(r)} \left|i\right\rangle_1 \left|j\right\rangle_2$, where $r = 1, 2, 3, 4$ and $(i, j) = \{0, 1\}$. The coefficients $B_{ij}^{(r)}$ are the matrices

$$B^{(1)} = M \begin{pmatrix} 1 & 0 \\ 0 & m \end{pmatrix} \qquad B^{(2)} = M \begin{pmatrix} m^* & 0 \\ 0 & -1 \end{pmatrix} \tag{4.28a}$$

$$B^{(3)} = N \begin{pmatrix} 0 & 1 \\ n & 0 \end{pmatrix} \qquad B^{(4)} = N \begin{pmatrix} 0 & n^* \\ -1 & 0 \end{pmatrix} \tag{4.28b}$$

and in the limit of $m = n = 1$ they become $\{\mathbb{1}, \hat{\sigma}_z, \hat{\sigma}_x, \hat{\sigma}_x\hat{\sigma}_z\}$ which, by no coincidence, are the unitary, corrective rotations needed to recover the initial state on Q3 which we derived in equation (4.5). We also see the effect of making measurements in a NME basis, i.e. where $|m|, |n| \neq 1$. In this case the matrices, $B_{ij}^{(r)}$ are not unitary - suggesting that there is no physical (i.e. unitary) corrective step which can be taken in final processing to perfectly recover the initial state on Q3. We now proceed to show that when $|m|, |n| \neq 1$ there still exists an optimal unitary corrective rotation but we will no longer recover the initial state with unit fidelity.

After Alice has made her measurement with outcome r the state with Bob is now

$$|\text{Bob}\rangle^{(r)} = \mathcal{N} \left\langle B^{(r)} | \psi_{\text{tot}} \right\rangle = \mathcal{N} \sum_{i,j,l} B_{ij}^{(r)^*} I_i U_{jl} D_{ll} \left|l''\right\rangle_3 \tag{4.29}$$

with $\mathcal{N}$ a normalisation factor. Defining $\boldsymbol{J^{(r)}} = \boldsymbol{U^T B^{(r)\dagger} I}$ and $\boldsymbol{I} = (\boldsymbol{B^{(r)\dagger}})^{-1} \boldsymbol{U^* J^{(r)}}$ we

can express the initial state of Alice and the final state of Bob as

$$|\psi_{\text{in}}\rangle = \sum_{i,k,l} J_l^{(r)} U_{kl}^* B_{ki}^{'(r)^*} |i\rangle_1 \quad \text{where} \quad \boldsymbol{B'} = \boldsymbol{B^{-1}} \tag{4.30a}$$

$$|\text{Bob}\rangle^{(r)} = \mathcal{N} \sum_l J_l^{(r)} D_{ll} \left|l''\right\rangle_3 \tag{4.30b}$$

and by inspection we can now see that the appropriate corrective rotation on Bob's state is $X^{(r)} = \sum_{i,k,l} U_{kl}^* B_{ki}^{'(r)^*} |i\rangle \langle l''|$ which leaves us with the final teleported state

$$|\psi_{\text{final}}\rangle^{(r)} = X^{(r)} |\text{Bob}\rangle = \mathcal{N} \sum_l J_l^{(r)} D_{ll} \left|l'''\right\rangle_3 \tag{4.31}$$

where we have defined the new basis $|l'''\rangle = \sum_{i,k} U_{kl}^* B_{ki}^{'(r)*} |i\rangle$. Computing the fidelity between this teleported state and the initial input state $\mathcal{F}^{(r)} = |\langle \psi_{\text{in}} | \psi_{\text{final}} \rangle|^2$ and including the normalisation factor we arrive at our final result

$$\mathcal{F}^{(r)} = \frac{\left(\sum_l |J_l^{(r)}|^2 D_{ll}\right)^2}{\sum_l |J_l^{(r)}|^2 D_{ll}^2}. \tag{4.32}$$

We discuss a few features of this result. In general, the teleportation fidelity is dependent on both the specific BSM outcome, r, and also on the input state through the coefficients $J_l^{(r)}$. In the specific case of using a maximally entangled resource state ($D_{ll} = 1$) and making a maximally entangling measurement on Q1 and Q2 ($\boldsymbol{B}\boldsymbol{B}^\dagger = \mathbb{1}$) then we can simplify the expression in equation (4.32) to find $\mathcal{F}^{(r)} = \sum_l |J_l^{(r)}|^2 = \text{tr}(|\boldsymbol{J}^{(r)}|^2) = \text{tr}(\boldsymbol{U}^T \boldsymbol{B}^{(r)\dagger} \boldsymbol{I} \boldsymbol{I}^* \boldsymbol{B}^{(r)} \boldsymbol{U}^*) = \text{tr}(\boldsymbol{B}^{(r)} \boldsymbol{B}^{(r)\dagger} |\boldsymbol{I}|^2) = \text{tr}(|\boldsymbol{I}|^2) = 1$ which is independent of the input state and measurement result as expected for perfect teleportation. However, this perfect teleportation is an exceptional case, if either the resource state *or* the joint measure-

ment are not maximally entangled then the average teleportation fidelity will be less than unity and will exhibit both input state and measurement result dependence.

4.5.2 Numerically calculated unitary rotations

Our fabricated circuit is not ideal and will not perform maximally entangling operations. We must calculate the optimal unitary transformations $X^{(r)}$ using the characterised circuit parameters. The analysis in the previous section was valuable because it rigorously showed that circuit imperfections lead to Bell State dependent teleportation fidelities that are no longer unity. However, in practise, it is not straightforward to calculate $\boldsymbol{E}$ and $\boldsymbol{B}^{(r)}$ separately as a function of the circuit parameters due to the compiled nature of the linear optical circuit. For example, the beam splitters responsible for creating the initial entangled state are also shared by those responsible for the bell state measurement. The central interferometer is also shared by both entangling operations, making it difficult to assign a definite phase to either part on its own.

We therefore perform a numerical analysis to approximate the optimal Q3 rotations and refer back to equation (4.32) to help understand the results. We simulate the output click statistics of the fabricated device minus the state preparation and state measurement stages. This is possible due to our element-wise characterisation of the device. This allows us to numerically prepare any input state we like on Q1 and to easily perform perfect state tomography on Q3. It is necessary to perform state tomography in this simulation since we specifically avoid calculating the full output state.

We prepare 10,000 input states on Q1 and simulate the expected clicks in the output modes. These input state are chosen uniformly according to the Haar-measure using a simple

Matlab script [2]. The Haar-measure ensures the states we generate are distributed uniformly across the surface of the Bloch sphere and are not biased toward the poles. Since we are interested only in the effect of the circuit we use perfect indistinguishable Fock states as our input to the simulation. Post-selecting on each of the four possible Bell measurements, for each input state we prepare, we generate four sets of simulated clicks on the output modes of Q3. Numerically implementing the three measurement operators, we use state tomography to recover the simulated quantum state on Q3. This leaves us with 10,000 input states and four sets of 10,000 output states corresponding to the four measurement outcomes. Our aim is to find the four unitary rotations, which when applied to the output states, will maximise the average fidelity between the set of all input and output states.

After generating all of the input and output states for our circuit we use a least-squares minimisation method in Matlab to find these unitary rotations. Each unitary is parameterised as,

$$\hat{U} = e^{i\hat{\sigma}_x\beta}e^{i\hat{\sigma}_y\gamma}e^{i\hat{\sigma}_z\delta}, \tag{4.33}$$

where the three parameters β through δ are varied in order to maximise the average fidelity between the input and output states. The resulting four unitary matrices are plotted in figure 4.10 as process matrices. The process matrix, χ, expresses the unitary rotation in the Pauli basis — $\hat{U}_i\rho\hat{U}_i^\dagger = \sum_{j,k}\chi_{j,k}\hat{\sigma}_j\rho\hat{\sigma}_k$ — and is therefore real and diagonal in the ideal case of perfectly entangling operations.

Figure 4.11 plots histograms showing the teleportation fidelity associated with each input state when the optimal unitary is used to recover Q3. As expected from equation (4.32), they show the fidelity of teleportation is strongly input-state dependent. Although the

[2]The Matlab script `randPsi.m` is available from `http://www.dr-qubit.org`

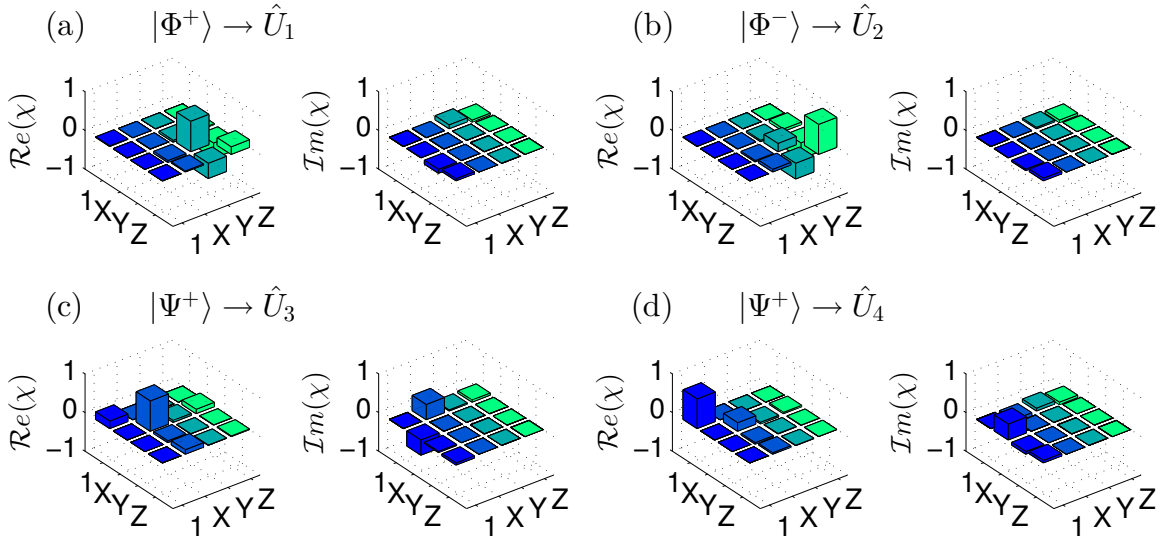


Figure 4.10: Numerically computed unitary rotations. Using the characterised circuit parameters we find the four optimal unitary rotations which are plotted above as the equivalent process matrix (χ) in the Pauli basis. In this basis, the original unitary rotations for a perfect circuit would correspond to one of the four diagonal elements with no imaginary components.

average fidelity still exceeds the classical limit the effect of the non-ideal circuit elements is quite striking. We note that the fidelity distributions can be clearly divided into two sets. Those with BSM outcomes $|00\rangle$ or $|10\rangle$ have a centrally peaked distribution whilst outcomes $|01\rangle$ or $|11\rangle$ show a peak at the highest fidelity. This reflects the fact that a non-maximally entangled BSM can be parameterised using two complex numbers resulting in the two classes of measurement outcome given in equation (4.27). These histograms show that one may be able to increase the average teleportation fidelity at the expense of lower success probability by only post-selecting on the $|01\rangle$ or $|11\rangle$ BSM outcomes.

4.6 Simulation results

The predicted teleportation fidelities for this experiment, using the above simulation methods, are show in table 4.1. Each row depicts a different set of model parameters with

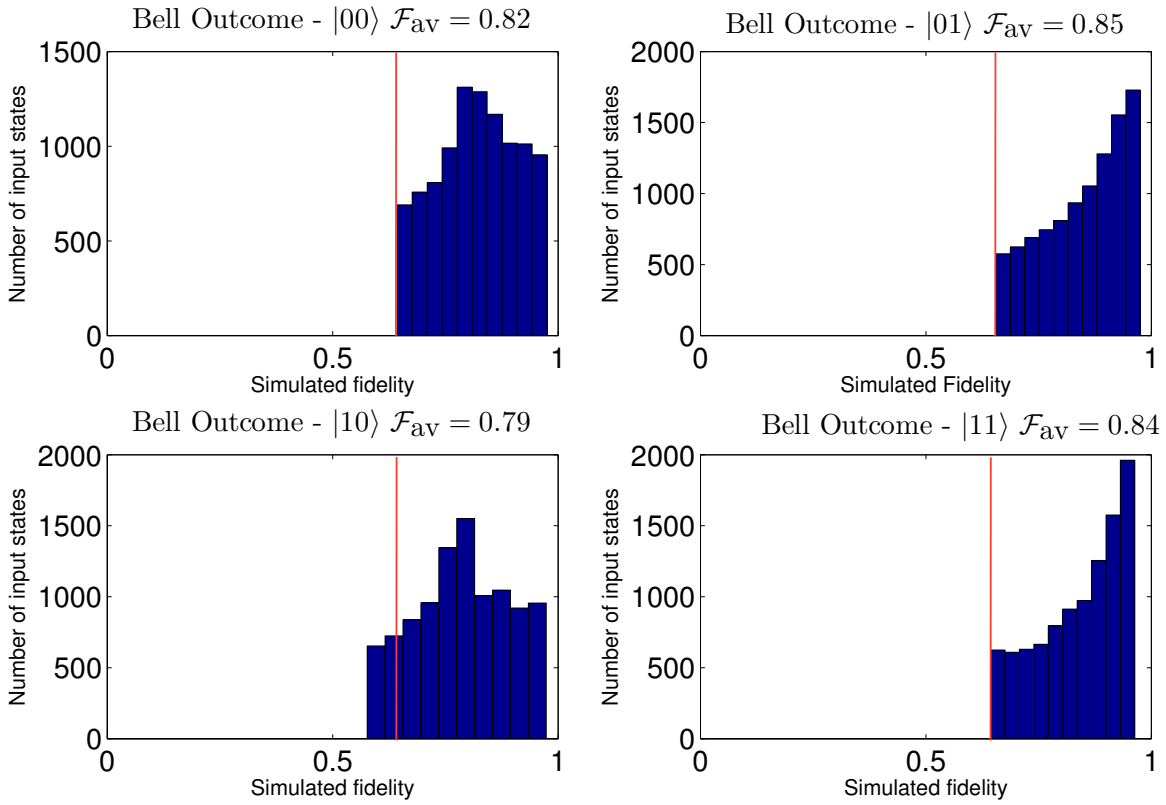


Figure 4.11: Histograms of simulated teleportation fidelity for 10,000 Haar-random input states using partially distinguishable squeezed state inputs. The effect of using the experimentally non-ideal photon sources as input to the circuit is seen to decrease the achievable average fidelity but still remains above what is classically possible (red line). The fidelity is seen to have strong input-state dependence as well as different characteristic shapes for each Bell state measurement outcome.

each column depicting the expected average teleportation fidelity for each of the four Bell state measurement outcomes. The table starts with a perfectly fabricated circuit and perfect photon sources resulting in unit average fidelity across all BSM outcomes as expected. Substituting the characterised circuit parameters into the model causes a collapse of the expected fidelity to below the classical limit. From the above discussion in section 4.5, this is expected and we must calculate new corrective unitary rotations once the circuit parameters are no longer ideal. All subsequent rows of the table for which the characterised circuit parameters are used include these new unitary rotations and we recover an average fidelity above the classical limit once again. Examination of the final row shows that even when we

take into account the imperfect circuit parameters, photon loss, photon distinguishability and multi-photon noise terms we still expect to be able to exceed the classical teleportation fidelity limit.

	BSM Outcome			
Model Parameters	$\|00\rangle$	$\|01\rangle$	$\|10\rangle$	$\|11\rangle$
Ideal circuit, ideal sources	1	1	1	1
Actual circuit, ideal sources, standard unitary rotations	0.14	0.49	0.17	0.22
Actual circuit, ideal sources, optimised unitary rotations	0.92	0.97	0.89	0.97
Actual circuit, incl. multi-photon emission, no circuit loss	0.87	0.96	0.83	0.94
Actual circuit, incl. multi-photon emission, incl. circuit loss	0.85	0.90	0.80	0.90
Actual circuit, incl. multi-photon emission, incl. circuit loss, incl. residual distinguishability	0.82	0.85	0.79	0.84

Table 4.1: Simulation results. The simulated fidelity of the teleportation circuit when different model parameters are taken into account. All simulations were performed using the optimised unitary rotations unless otherwise stated.

4.7 Experimental results

4.7.1 Experimental method

We now describe the main experiment showing on-chip quantum teleportation. Three single photons generated via two parametric down-conversion sources were coupled into the UV-written, silica-on-silicon photonic chip in the same fashion as the experiments in chapter 3. The arrival time of the three photons at the chip were adjusted using two motorised translation stages and the zero delay (temporally indistinguishable) point was found by performing two, two-photon HOM interference dips between the photons encoding Q1 and Q2, and those encoding Q2 and Q3, shown in figure 4.12. Taking these interference dips not only identifies the zero delay position required for the experiment but serves as a stringent check that the

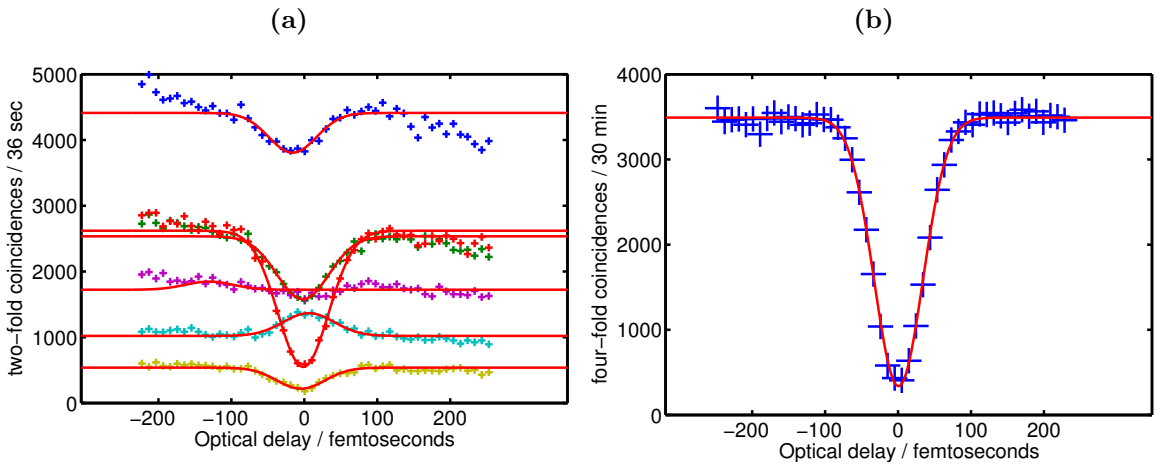


Figure 4.12: Temporal alignment of the photons. The three photons are made indistinguishable in time by adjusting two optical delay stages. (a) Photons input to Q2 and Q3 originate from the same SPDC source allowing an unheralded dip to be taken. There are multiple indistinguishable paths these two photons can take leading to the collection of 6 interference dips and anti-dips seen in the left panel. (b) Photons input on Q1 and Q3 originate from separate sources and thus we must take a heralded dip, collecting 4-fold coincident counts. In this case there is only one indistinguishable path - when both photons interfere on beam splitter η_{10} - leading to the single dip seen on the right. The visibility of all of these dips are computed using a Gaussian fit and compared with the expected values to verify the photons are indistinguishable.

photons remain truly indistinguishable within the chip [3]. An array of six avalanche photodiodes (APDs) monitored all output modes. We identified successful teleportation events as trials in which one photon is detected for each of the three qubits, as well as in the ancillary heralding arm. Four-fold detection coincidences were registered using an FPGA. Data was collected for three linearly independent input states, generated on Q1 by appropriately setting θ_1 and ϕ_1 - $|\psi_{\text{in}}\rangle = \hat{U}(\theta_1, \phi_1)\,|0_L\rangle$. For each input state, θ_2 and ϕ_2 are adjusted to perform projective measurements of Q3 in three different bases, allowing us to reconstruct the teleported state via maximum-likelihood tomography. In order to verify successful teleportation, we apply the corrective unitary rotations found above in section 4.5.2, to the reconstructed state of Q3, $\hat{\rho}_{\text{out}}$, in post-processing and calculate the fidelity to the input states, $\mathcal{F} = \langle\psi_{\text{in}}|\,\hat{U}_i'\hat{\rho}_{\text{out}}\hat{U}_i'^{\dagger}\,|\psi_{\text{in}}\rangle$.

[3]For example, the spectral overlap of the photons can shift through different waveguide dispersion, higher-order spatial modes can be present in the x-couplers and polarisation rotation can occur

There were a number of challenges in running such an experiment. The insertion loss of the teleportation chip was -10 dB. The probability of all three photons making it through the experiment was thus $\sim$ 0.1%. Of the photons which are transmitted through the circuit, the probabilistic nature of the protocol means only 1/27 of these trials is recorded as a successful event. This low probability is compounded by our use of probabilistic SPDC photon sources which means our final data rate for successful teleportation events is only 2-5 mHz. For each setting of the phase shifters there are 4 possible outcomes and we require of order 100 events in each outcome in order to keep the Poissonian counting error on our results low enough to be useful. This requires of the order of 3 days continuous data collection. With three input states each requiring measurement in three different bases this amounts to around a month's worth of continuous running time.

Keeping the experiment running for this long period required regular interventions to correct slow drifts in the coupling of the source to the circuit. Two-fold coincidence rates were monitored remotely by computer during the experiment. When the rates dropped below a certain threshold, the data collection was stopped and adjustments made to the fiber collection optics and SHG stage of the source and to the VGA butt-coupling from the source to the chip. This was required at least once a day due to changes in lab temperature. Approximately every 5 days, the experiment was completely re-optimised and new HOM dips found to ensure the photons remained indistinguishable. Occasionally, this resulted in very slight adjustments of the narrow-band interference filters in order to maintain the spectral overlap between all three photons. To try and average-out these systematic drifts in performance, the measurements operators were changed every 2 hours by automatically adjusting the voltages applied to phase shifters θ_2 and ϕ_2. The stability of the phase shifters

was monitored by periodically blocking all photons Q1 and Q3 using servo-controlled beam-blocks and recording the ratio of single photons measured in each pair of qubit rails due to Q2. For a set phase this ratio should remain constant and this served as a useful check that the measurement operators were being cycled correctly during unattended operation.

4.7.2 Teleportation results

The teleportation results for the Bell measurement outcome $|01\rangle$ are summarised in figure 4.13. The experiment was performed using the three linearly independent input states detailed in section 4.3. We label them $|H'\rangle$, $|D'\rangle$ and $|L'\rangle$, according to the closest lying axis state on the Bloch sphere (figure 4.13(a-c)). The initial input states are recovered with an average fidelity of $89 \pm 3\%$ which exceeds the classical limit of 2/3 by more than 6 standard deviations. Since any state on the Bloch sphere can be generated using only linear combinations of these three states, these results are sufficient to conclude that this chip is capable of teleporting general quantum states with high fidelity. The full set of reconstructed density matrices for all four BSM outcomes are plotted in Appendix A as figures A.1 and A.2.

We summarise the results from all four Bell measurement outcomes in figure 4.14. In each case, the average fidelity for the three input states (shaded red boxes) is higher than the classical limit (dashed red line), irrespective of which classical outcome is obtained. Unbalanced propagation loss in combination with beam splitter ratios which deviate from their designed reflectivity mean that the four Bell outcomes occur with non-equal probability resulting in variations in the number of coincidence events recorded for each measurement. As a result, some measurements exhibit a larger error than others due to the fewer events recorded for those outcomes.

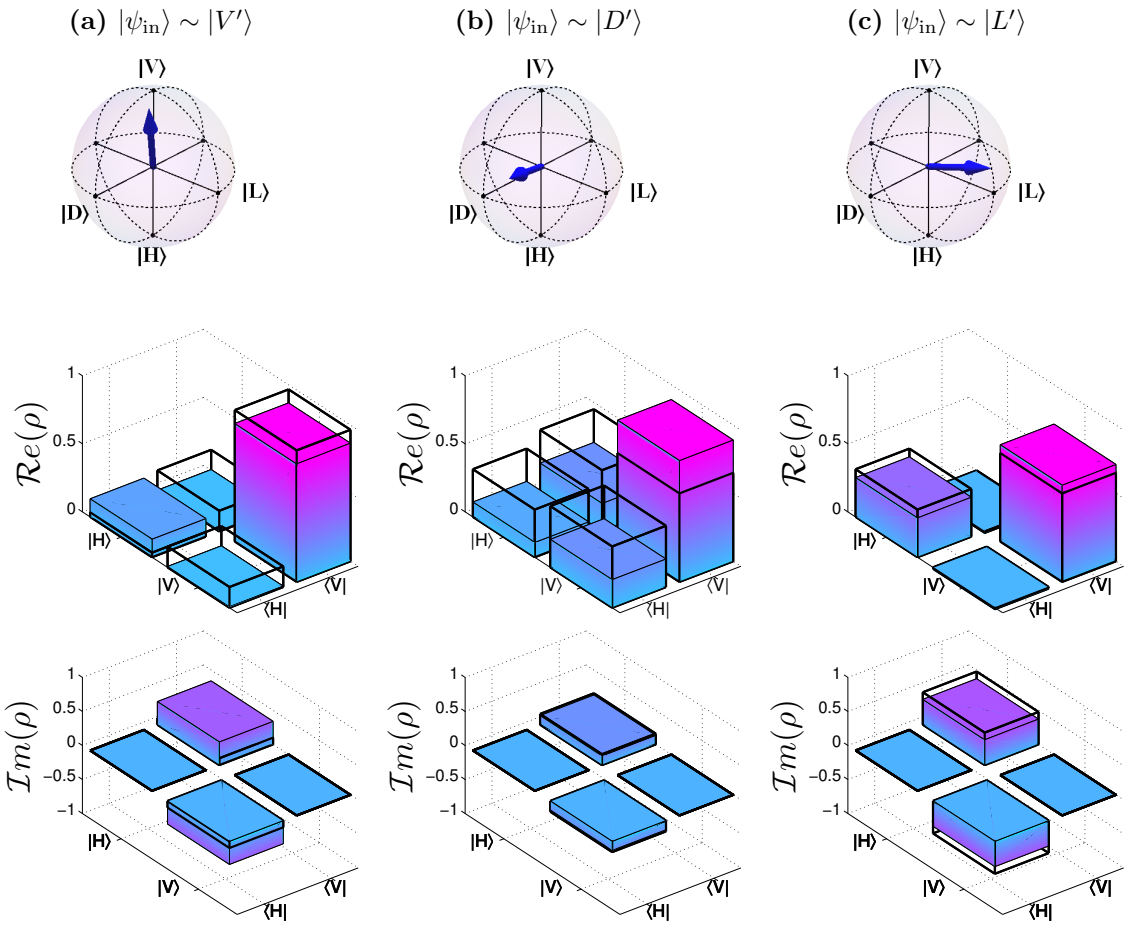


Figure 4.13: Reconstructed density matrices of the teleported states. The initial qubit states on Q1 for each of three trials (columns) are depicted on the Bloch sphere (top) and as real and imaginary parts of a density matrix (black wire frames, middle). The final teleported states on Q3 are reconstructed using on-chip quantum state tomography and then transformed by optimal state-independent rotations in post-processing (coloured bars - shading is related to bar height). The fidelity, $\mathcal{F}$, between the initial and final state shown is calculated (bottom). Representative data here are for experiments with a $|01\rangle$ Bell measurement outcome.

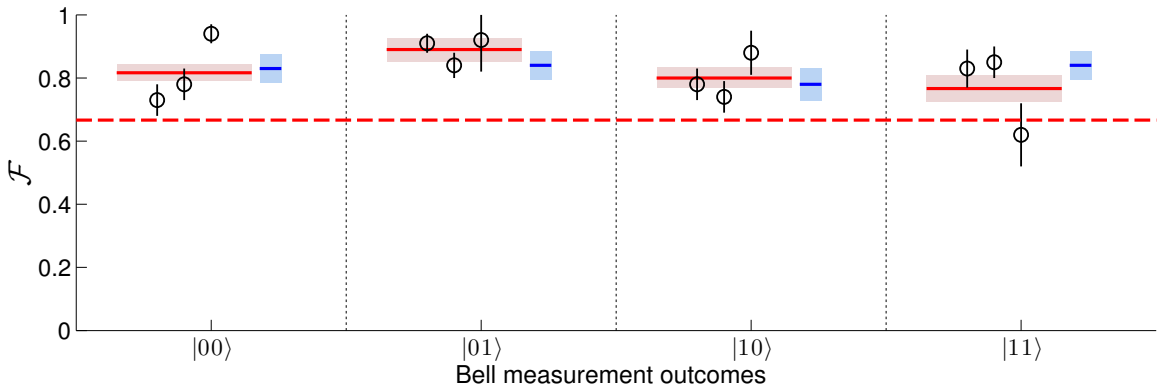


Figure 4.14: Measured and simulated fidelity of on-chip quantum teleportation. The measured fidelity of the three teleported states ($|V'\rangle$, $|D'\rangle$ and $|L'\rangle$ from left to right respectively) for each outcome of the Bell measurement are plotted with circles. Errors are calculated using a Monte-Carlo method taking account of Poissonian counting statistics and uncertainties in the characterised measurement operators and input states. The red shaded areas show the measured fidelity averaged over all three input states and clearly exceeds the best average classical fidelity given by the red dotted line. The blue shaded areas show the predicted average fidelity taking account of imperfect circuit fabrication, higher order photon emission and residual photon distinguishability. The data used to generate this plot is given in Appendix A.

We now compare our results to the simulated performance of our circuit carried out in section 4.6 taking into account the experimentally characterised beam splitter ratios and interferometer phases, higher-order photon emission, photon distinguishability and propagation losses. The predicted teleportation fidelity averaged over all input states are summarised in figure 4.14 as shaded blue boxes. The error on these predictions are estimated using a Monte-Carlo method over a range of different input states and model parameters. We find that the overlap between the shaded red and blue regions is high suggesting that the model is well able to explain the observed average fidelities. This gives us some confidence in using the model to quantify the reduction in achievable fidelity which can be attributed to different parts of the experiment.

The non-ideal beam splitting ratios reduce the quality of the entanglement in the circuit and immediately reduce the achievable fidelity to around 90% - a figure already below the $> 99\%$ fidelity which is thought to be required as part of a fault-tolerant quantum

computer[18]. The non-ideal nature of the photon sources is found to be responsible for the remaining observed reduction in measured fidelity, since the protocol relies critically on high-visibility quantum interference between all three photons. From our theoretical model we find that residual photon distinguishability and the higher-order terms from the down-conversion contribute in roughly equal proportion to the fidelity.

4.8 Future devices

Aside from the errors due to the imperfect photon source as stated above, the teleportation experiment was especially affected by errors due to the integrated circuit itself. The compact footprint and inherent stability of integrated photonic circuits allows more complex linear optical networks to be built than is possible using only bulk optical components. However, integrating these optical components presents additional problems for scaling circuits to meaningful levels of complexity. First, the increased propagation loss in these circuits exponentially reduces the achievable data rate and amplifies the negative effect of higher-order photon emission[242].

Second, we have seen that the fabricated beam splitters and phase shifters will also inevitably show some deviation from their designed parameters. The work in this chapter highlighted that non-ideal beam splitters and phase offsets not only require modified unitary rotations to be carefully calculated but introduce fundamental limits to the performance of the circuit. Further, the fragility and restricted tuning range of the thermo-optic phase shifters limited the number and quality of the experiments we could perform.

4.8.1 Reasons for poor device performance

In previous work by our group[121], we have seen both lower propagation losses, better phase-shifter tunability and more accurate beam splitting ratios from the UV-written silica-on-silicon waveguide platform. However, this was only a simple device consisting of a single Mach-Zehnder interferometer. The more complex design used in this work took significantly longer to fabricate (of order 90 minutes) during which, there was significant out-gassing of hydrogen from the waveguide core layer. This leads to unknown changes in the photosensitivity of the core layer making it difficult to correct the fluence of the UV writing beam. This made achieving a desired beam splitting ratio very difficult since the chip parameters were constantly evolving.

The limited tuning range of our phase-shifters that we observed was due to two factors: (1) There was a 17 μm silica overclad between the core layer and the heater. In order to achieve the required temperature change of the core layer a large amount of heat had to be generated by the phase shifter. (2) Because the heaters were run at very high temperatures, they tended to oxidise and undergo moisture-induced degradation. This reduced the maximum safe operating temperature of the heaters before they failed. Ultimately these two factors meant we could apply up to 200 mW of power to each phase shifter before they would fail. However, due to the thick silica overclad, this only generated phase shifts of $\approx \pi/2$.

4.8.2 Proposed modular chip design

Based on the findings of the past two chapters we propose a new *modular* device design which will enable more complex photonic chips to be constructed whilst keeping loss and

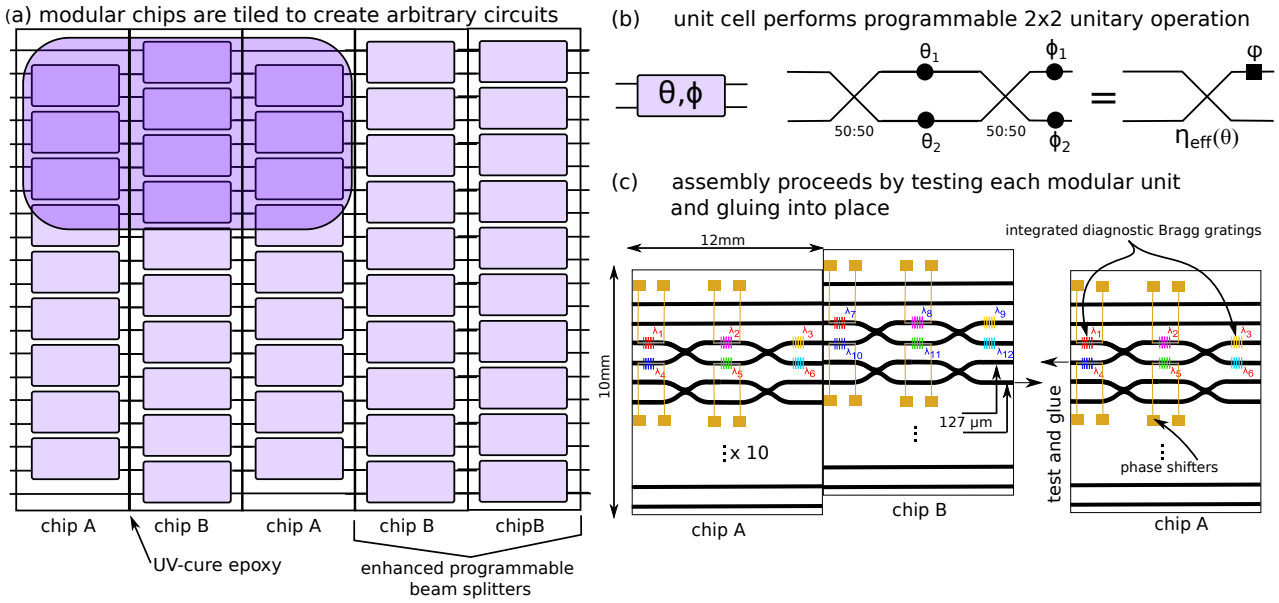


Figure 4.15: Modular chip design. (a) Schematic of an assembled circuit consisting of interconnected fully tunable two-mode unitary gates. Each modular chip contains up to 10 MZI unit cells in two offset chip designs, A and B. UV-cure epoxy is used to glue the individual modules together. Finer control of the two-mode unitary operation can be achieved by placing two modular units in series as shown on the right of the chip. We have highlighted in the top left a section which could implement all of the experiments discussed in the previous two chapters. (b) Each unit cell is composed from a MZI and four tunable thermo-optic phase shifters. The phase shifters occur in pairs to reduce the total amount of power each individual shifter must consume. Each unit cell is equivalent to a fully tunable beam splitter and phase shifter as shown on the right. (c) Schematic of the actual chip designs showing the phase shifter bond pads and integrated diagnostic Bragg gratings for in-situ characterisation. These three chips would enable the construction of the circuit used in this thesis but with arbitrary control over all of the beam splitting ratios.

component errors to a minimum. The basic design idea is shown in figure 4.15. The unit cell of this modular design consists of a series of Mach-Zehnder interferometers made using two 50:50 beam splitters. A set of four thermo-optic phase shifters are used to implement tunable two-mode unitary operations. Tiling many of these individual modular chips together, it is possible to realise any $N \times N$ linear-optical circuit [57]. In addition, weakly reflecting out-of-band diagnostic Bragg gratings are patterned along the waveguides in each unit cell. These can be used to accurately characterise the beam splitting ratios, calibrate the phase shifters and determine the (possibly unbalanced) waveguide propagation loss [243].

The advantages of this design are three-fold: (1) Fabrication efforts can be concentrated

on making a single optimised design. Arbitrary beam splitter values and phase offsets can be 'dialled-in' during runtime rather than having to be optimised for each fabrication run, mitigating individual component errors. (2) Each unit cell can be quite short and be fabricated quickly, resulting in less hydrogen out-gassing and leading to more consistent waveguide writing and lower propagation loss. (3) Each unit cell can be tested and characterised before being added to the circuit using the integrated Bragg gratings. This allows poor quality chips to be discarded separately from the rest of the circuit. Of course this approach also presents several challenges, including achieving low loss chip-to-chip coupling, mechanical weakness at the joins and the propagation of errors from imperfect element-characterisation.

However, this design will allow us to move beyond the current generation of circuits and help mitigate any component deviations from design. The Bragg gratings can be used to quickly and accurately characterise the fabricated circuit enabling us to verify quantum behaviour. Because the circuit is fully reconfigurable, it is possible to implement a number of different experiments using the same chip, including on-chip state-preparation and tomography.

In order to realise the full potential of this design, we must ensure the thermo-optic phase shifters are capable of implementing a 2π phase rotation. We are working to implement a number of changes in order to overcome the limitations we have mentioned above. First, we will reduce the overclad thickness to only a few microns, greatly reducing the heat needed to achieve a given temperature change. We will also deposit thicker NiCr heating elements over the waveguides. Currently these elements were only 300-400nm thick, and potentially 2-3 microns can easily be achieved. Thicker NiCr elements are not only more resilient to oxidation but also have reduced resistance that makes a more effective heater at

lower voltage. Finally, the modular chips will also feature a thin capping layer of silica, a few microns thick, which will be deposited over the NiCr heater elements to prevent them oxidising. This will enable us to run the heater elements hotter than before.

Testing chip-to-chip assembly

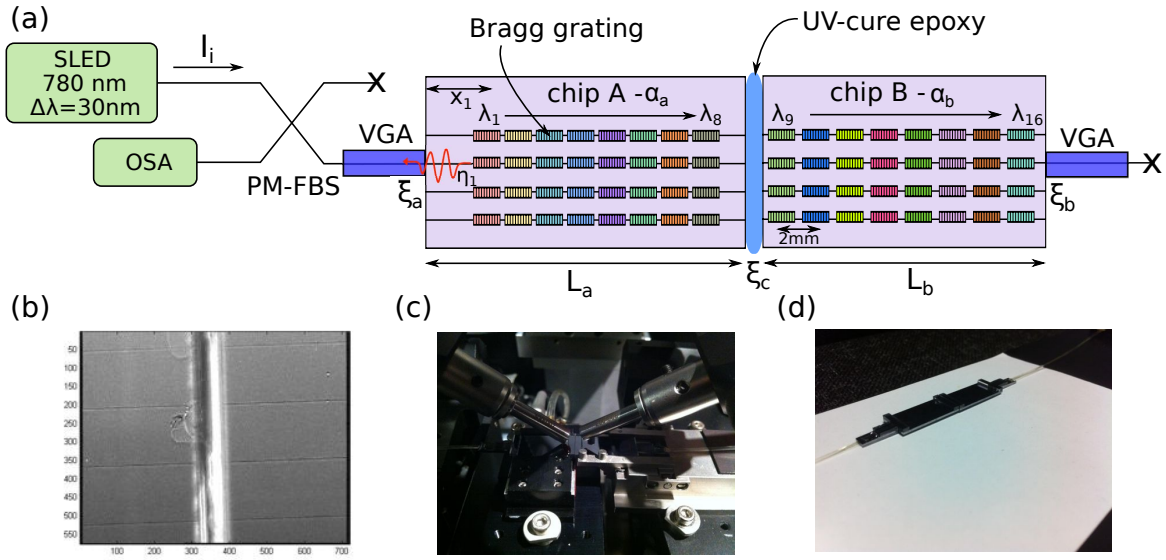


Figure 4.16: Chip-to-chip gluing and characterisation with Bragg gratings. (a) A super-luminescent LED (SLED) is used as a broadband light source to interrogate the diagnostic Bragg gratings. The SLED is coupled into the chip through a polarisation-maintaining fibre beam splitter (PM-FBS) and all free fibre ends are carefully terminated to prevent unwanted reflections. An optical spectrum analyser (OSA) records the reflections from the chip. Each chip is $L_i \approx 20\,\text{mm}$ long, has a fibre coupling efficiency, ξ_i a propagation loss, α_i and contain 8 weakly reflecting Bragg gratings of reflectivity, η_i, at wavelengths, λ_i, positioned at x_i from the chip face. It is shown in the text how this can be used to determine the chip-to-chip transmission, ξ_c, independently of all other losses. (b) Microscope image of the chip-to-chip joint just after gluing, showing four waveguides. (c) Photograph of the UV light source ready to cure the UV epoxy. (d) Photograph of the final assembled chip-to-chip test device.

This new modular design relies on being able to assemble many individual chips into one monolithic circuit. This needs to be done with very low loss [4]. This should be possible since the waveguide modes of the chips are exceptionally well matched. To test this, we fabricated

[4]The exact definition of 'low loss' will depend on the complexity of the circuit we wish to implement. However, if the loss is $< 0.1\,\text{dB}$ per interface this should be comparable to the additional loss we would have without the modular design.

a series of dummy chips consisting of straight waveguides patterned with diagnostic Bragg gratings and glued them together. The Bragg gratings allowed us to make a measurement of the chip-to-chip coupling loss we achieved after the glue had cured. We used computer-controlled piezo-driven 6-axis alignment stages (Thorlabs NanoMax 609) to accurately align the two chips and used an axial force sensor (Thorlabs FSC103) mounted under the chip to make sure the two chip faces were parallel. Index-matched, low-shrinkage UV cure epoxy (Dymax OP-29) was carefully deposited between the chip faces and a laser was coupled through the device to monitor the transmitted power. A high-power UV lamp was used to cure the epoxy once optimal coupling was achieved. Monitoring the transmitted power revealed no measurable change in the transmission during the cure cycle.

Using the weakly reflecting diagnostic Bragg gratings we can now characterise the interface loss across the glue joint independent of propagation and coupling loss [243]. A schematic of this experiment is detailed in figure 4.16. A broadband light source (super-luminescent LED) is coupled into the chip through a polarisation maintaining fibre beam splitter. Reflections from the Bragg gratings propagate back through the beam splitter and are recorded on an optical spectrum analyser (OSA). Figures 4.17(a-b) show the reflected spectra from the guide when coupling from the forward and backward direction. Assuming that all the losses are wavelength independent and the reflectivity of the Bragg gratings are the same when coupling in each direction we can write down an expression for the reflected spectra for each of the 16 Bragg gratings,

$$I_i^{(\mathrm{f})} = \begin{cases} \xi_{\mathrm{f}}^2 e^{-2\alpha_a x_i} \eta_i I_i & \text{for} \quad i = [1,8] \\ \xi_{\mathrm{f}}^2 e^{-2\alpha_a L_a} e^{-2\alpha_b (x_i - L_a)} \xi_c^2 \eta_i I_i & \text{for} \quad i = [9,16] \end{cases}, \tag{4.34}$$

where $I_i^{(\mathrm{f})}$ denotes the reflection from Bragg reflector i at position x_i with reflectivity η_i. The coupling efficiency is ξ_f , the two chips have length $L_{a,b}$ with propagation loss $\alpha_{a,b}$ and the chip-to-chip transmission is ξ_c. Similarly, the reflection spectra when coupling in the backward direction, $I_i^{(\mathrm{b})}$, can be written as

$$I_i^{(\mathrm{b})} = \begin{cases} \xi_{\mathrm{b}}^2 e^{-2\alpha_b L_b} e^{-2\alpha_a (L-x_i)} \xi_c^2 \eta_i I_i & \text{for} \quad i = [1,8] \\ \xi_{\mathrm{b}}^2 e^{-2\alpha_b (L_a+L_b-x_i)} \eta_i I_i & \text{for} \quad i = [9,16] \end{cases} . \tag{4.35}$$

The logarithm of the ratio $R_i = I_i^{(\mathrm{f})}/I_i^{(\mathrm{b})}$ is then taken, giving two straight line fits

$$\ln R_i = \begin{cases} 2\left(\ln \frac{\xi_{\mathrm{f}}}{\xi_{\mathrm{b}}} + \alpha_a L_a + \alpha_b L_b - \ln \xi_c\right) - 4\alpha_a x_i & \text{for} \quad i = [1,8] \\ 2\left(\ln \frac{\xi_{\mathrm{f}}}{\xi_{\mathrm{b}}} - \alpha_a L_a + \alpha_b (L_b + 2L_a) + \ln \xi_c\right) - 4\alpha_b x_i & \text{for} \quad i = [9,16] \end{cases} . \tag{4.36}$$

The gradient of each of these straight lines provides an estimate of the propagation loss of each of the chips independent of the coupling efficiency. The difference in the intercepts of the line, $\Delta c = 4(L_a(\alpha_a - \alpha_b) - \ln \xi_c)$, can be used to work out the chip-to-chip coupling loss, again independent of the propagation and coupling losses. Example results are shown in figure 4.17.

Once we had optimised the gluing procedure, we managed to achieve a chip-to-chip interface loss of less than 0.2 dB (4.5%) across four guides simultaneously. We expect to be able to better this with an improved polishing routine as well as moving to work at 1550 nm. However, this loss is already small enough to enable at least 10 modular chips to be connected together which will provide an excellent platform on which to realise more complex quantum circuits in the next few years.

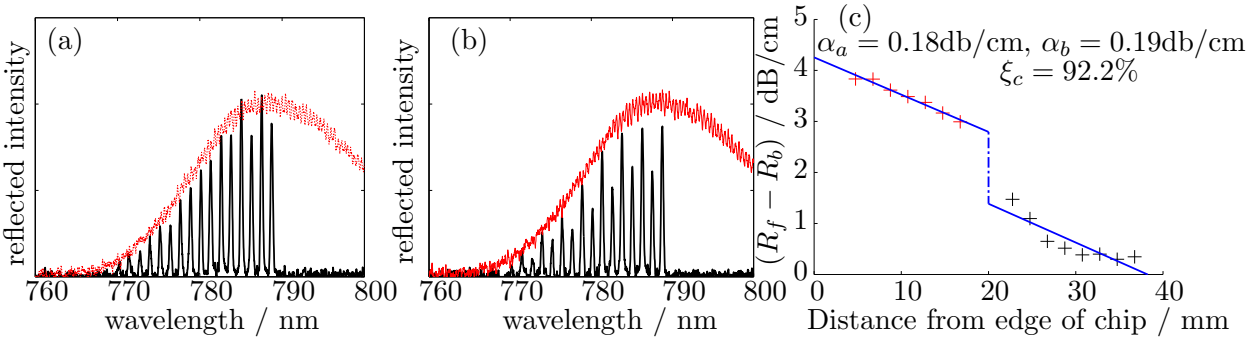


Figure 4.17: Reflection spectra from glued chip modules. (a) Reflected Bragg peaks from the forward coupled direction showing the response from the 16 Bragg reflectors patterned along the waveguide. The reference spectrum from the SLED is shown in red. (b) The reflected Bragg peaks from the backward coupled direction. The larger difference in peak heights is because one chip has a larger propagation loss than the other. (c) Plot of the logarithm of the ratio of the two sets of peak heights. The gradients give the propagation loss of each chip and the difference in the intercepts give the interface loss between the chips.

In a similar fashion, these diagnostic Bragg gratings may be used to characterise each module before it is glued into place. The two beam splitters, the central phase shifters and the propagation loss can be calibrated within each module, with the final phase shifters and coupling loss being calibrated after it is glued into place. This provides us with a fully reconfigurable low-loss waveguide platform which can be accurately calibrated and characterised.

The construction and manipulation of larger photonic quantum states continues to be ultimately limited by the absence of a suitable on-demand single photon source. Whilst this modular system will likely only enable us to realise linear optical circuits consisting of a few tens of spatial modes, combined with the high-efficiency telecom-band photon detectors discussed in the next two chapters, it is likely that this system will remain more than sufficient to carry out a wide variety of interesting experiments in the next few years given the limited number of photons which can be feasibly generated.

Chapter 5

Transition-Edge Sensors

THE previous two chapters have explored how one may use integrated photonic circuits to controllably implement quantum logic operations on single photons. However, so far we have not discussed single photon detectors in any detail, despite the crucial role they play in quantum optics experiments[17,244–246]. In this chapter, we review the state of currently available photon detectors and discuss the installation and characterisation of superconducting transition edge sensors (TESs) as part of our experimental setup here in Oxford. The high efficiency and intrinsic number-resolution provided by this detector technology will help enable a new generation of experiments using more photons.

5.1 Single Photon Detectors

Photo detection describes the process of photons being absorbed by some medium and converted into a large enough electrical signal to be measured. This consists, first, of the interaction of light and matter, converting each photon into an excited electronic state, followed by low-noise amplification of this electrical signal. In typical room-temperature

devices, the energy of a single photon at visible wavelengths is of the same order as the fundamental electronic excitation gap ($\sim 1\,\mathrm{eV}$) of the material. This means only one photo-electron is generated per incident photon. Typically, the current is increased by subsequent amplification, introducing additional noise. The subject of this chapter concerns photon detectors based on superconducting materials. In the superconducting state, the relevant energy gap is orders of magnitude smaller ($\sim 1\,\mathrm{meV}$), leading to the creation of many thousands of excited quasi-particles (broken cooper-pairs) per absorbed photon. This requires less amplification, leading to lower noise as well as spectrally broadband operation.

Returning to the conclusions from chapter 2, we know that single-photon detection plays a fundamental role in LOQC. It is the inherent non-linear nature of photo-detection (we either measure a photon or we don't) which enabled KLM to generate measurement-induced non-linearities and synthesise effective two-photon interactions. In particular, for LOQC the detection efficiency, η, — the probability that we record a count given that a photon arrives at the detector — is an absolutely critical parameter [75]. Failing to detect a photon within a LOQC protocol constitutes an error which necessitates using complex loss-tolerant encoding schemes in any practical realisation [247–249]. For every probabilistic gate within a particular LOQC scheme, *all* of the photons entering the circuit must be detected at some point for the gate to succeed. It is clear that the detection efficiency enters exponentially into this success probability as the number of photons in the protocol increases and is a particular problem for most current experiments which rely on post-selection. In fact, we not only have this exponential scaling, but there are often hard limits on the minimum detection efficiency, below which it is not possible, even in theory, for certain schemes to work such as loop-hole free Bell tests [250,251], quantum enhanced precision metrology [174,252] and LOQC [75,253,254].

5.2 Available single-photon detector technology

There are a wide-variety of performance parameters that characterise available photon detectors, including: detection efficiency, the spectral range over which they operate, the dead time of the device, the timing jitter on the generated electrical pulse, the dark count rate and, for number resolving detectors, the energy or electron resolution.

Of these parameters, we have already argued that detection efficiency, η, is the single most important one in the context of LOQC. Experiments are usually conducted at a single wavelength so that spectral range is not an issue. Dark counts are a potential source of noise in experiments and can lead to errors in certain photon counting experiments. However, the single photons used in most LOQC experiments are highly time-correlated allowing narrow time-gating to be used to almost completely remove the effects of dark counts. Dead time and timing jitter simply limit the clock rate, R, at which the experiment can be run. Since the success rate of an experiment involving the detection of N photons scales as $P_{\text{success}} \sim R \cdot \eta^N$, it is generally advantageous to optimise η, even at the expense of R. Clearly, η is the more important parameter for scalability.

Multiplexed-detection

We recall from chapter 2 that being able to discriminate photon numbers is crucial in almost all LOQC schemes as well as being essential for other technologies such as quantum-enhanced metrology[252]. True photon number resolution requires that the detector itself is inherently energy resolving. The TES's discussed here fall into this category. However, one can multiplex conventional 'binary' detectors either spatially[255,256] or temporally[257,258] to create a *pseudo* number-resolving detector — each photon now has some chance of being

detected by a separate detector. To achieve high fidelity number-resolution not only is it necessary to have very efficient detectors but also to have many more spatial pixels (or time bins) than there are incident photons to reduce the possibility of two or more photons being absorbed by only one detector. The probability that such a multiplexed detector will correctly output n clicks given that n photons were incident is given by,

$$P(n|n) = \frac{N!}{N^n(N-n)!}\eta^n, \tag{5.1}$$

where N is the number of spatial pixels or time-bins used in the multiplexing strategy[257–259]. This means that even for a perfectly efficient binary detector, 10 pixels or time-bins are required to achieve the same two-photon detection fidelity as a 95% efficient number-resolving detector ($P(n|n) \approx 0.9$). The number of pixels or time-bins required to maintain a given detection fidelity grows as the number of incident photons grows, increasing the complexity of the read-out electronics. Keeping this fidelity as high as possible is critical for LOQC since it avoids the amount of costly error-encoding which has to be performed[94]. Many multiplexing strategies also suffer from decreased overall detector efficiency, since splitting the incoming signal cannot be performed losslessly. Detectors which are intrinsically able to resolve photon number with high fidelity are thus highly desirable for LOQC.

5.2.1 Conventional detectors

Photomultiplier tubes

In these photoelectric devices, an incoming photon strikes a low work-function material and liberates an electron. This single free electron is accelerated through a large electric field before striking a secondary emitter with around 100 times more energy, liberating more electrons. Cascading these gain stages can produce $> 10^6$ electrons at the final anode for each initial photo-electron - a current which is easily amplified and measured. Despite being expensive, fragile, requiring high voltages and vacuum components, PMTs continue to find extensive use today due to their low noise and easy construction of large area devices [245]. For typical guided-wave quantum optics applications however, since all photons are tightly confined to the $\sim 5\,\mu$m core diameter of a single-mode fibre, large area devices are not of huge importance and their use has been supplanted by their solid-state replacement — the avalanche photodiode (APD).

Avalanche photodiodes

In a solid-state photodiode, instead of photons liberating an electron from a metallic surface, they are absorbed in the depletion region of a $p-n$ junction semiconductor. So long as the photon has an energy greater than the band-gap of the material, the absorbed photon promotes an electron from the valence band into the conduction band to create an electron-hole pair. The APD is reverse-biased above its breakdown voltage using a power-supply which can source effectively unlimited current to operate in what is known as *Geiger*-mode. In this regime, the photo-generated electron-hole pairs multiply by impact ionisation faster, on average, than they can recombine or be extracted by the charge collection circuitry.

The subsequent current grows exponentially in time to a size where it can be detected by conventional electronics. The achievable detection efficiencies are around 50-60% at 800 nm for silicon-based devices. The wide availability of highly pure silicon provides good quality detectors with low dark counts ($D < 500\,\text{Hz}$) and good timing jitter ($\Delta t \sim 500\,\text{ps}$) at a reasonable price which, coupled with the simple, room-temperature, bench-top operation, has ensured they remain the single photon detector of choice for the majority of quantum optics experiments world-wide.

Unfortunately, the band-gap of silicon is too wide to enable its use at the lower-energy telecommunication-band around 1550 nm. Instead, detectors made using indium gallium arsenide (InGaAs) as the active semiconductor can be used to detect photons over the wavelength range $\sim 1.1\,\mu\text{m} - 1.7\,\mu\text{m}$. However, InGaAs has a direct band-gap which means the lifetime of the photo-generated electron-hole pairs is too short for enough of them to diffuse to the collection electrodes. Instead, InGaAs detectors are grown on a Indium Phosphide (InP) substrate which provides much higher charge transport [260]. The band-structure discontinuity at the interface of these two materials provides a small energy well where electron-hole pairs can be become trapped. This reduces the achievable detection efficiency to $< 25\%$ and contributes extra noise terms which necessitate long dead times of $> 10\,\mu\text{s}$.

Moving toward 1550 nm

The long-distance, high-bandwidth networks which make today's internet infrastructure a reality all rely on the ultra low-loss transmission of optical signals pioneered in the 1970's [175]. Doped silica-based fibres exhibit exceptionally low propagation loss around 1550 nm. Because of this, all of the telecommunication infrastructure is based on this standard, and as

a result there has been hundreds of billions of dollars invested into making a vast array of optical devices — switches, circuits, modulators, couplers, filters, etc. — which are compatible with this standard. Working at this longer wavelength reduces propagation loss in waveguide circuits and reduces the tolerances needed to achieve good coupling. It would appear to make enormous sense to leverage this technology in the construction of future *quantum* optical devices.

Yet, the vast majority of quantum optics experiments performed to date have been carried out in the wavelength range for which conventional silicon APDs display good detection efficiency (400 nm-900 nm). The poor performance of InGaAs APDs, in terms of their low efficiency, high dark counts and long dead times have so far out-weighed the benefits of working at telecommunications wavelengths. The emergence of superconducting detectors which exhibit high-efficiency and low-noise at 1550 nm has started the gradual transition of moving toward telecomm wavelengths. Coupled with the higher available detection efficiencies, these superconducting detectors will enable more complicated future quantum experiments by decreasing loss and taking advantage of the many mature telecomm components that are easily available.

5.2.2 Superconducting detectors

Superconducting detectors do not suffer from band-gap or work-function limited cutoff since the relevant superconducting energy gap is only ~ 1 meV, far smaller than any photon energy of interest [261,262]. This intrinsically gives them a very broad-band spectral response and means they provide a similar high performance at both visible and telecommunications wavelengths. Further, because these devices are necessarily operated at cryogenic tem-

peratures they can take advantage of the reduced thermal and electrical noise and smaller material heat capacities at low temperatures [262]. This allows these detectors to significantly out-perform room temperature devices in several measures (e.g. dark counts). I discuss four such devices below.

Superconducting tunnel junction detectors - STJs

Superconducting tunnel junctions (STJs) were among the first photon detectors to make use of a superconducting material [263]. An STJ consists of two superconducting electrodes separated by a thin insulating layer placed in a small magnetic field to prevent any super-current from flowing. An incoming photon is absorbed by the material and breaks many Cooper pairs, creating a large number of normal electrons. These electrons diffuse across the junction under a small voltage bias, inducing a normal current transient that can be measured with a low-noise FET amplifier. Because the size of the current is proportional to the energy of the incoming photon, they can be used to resolve photon number. Whilst STJs remain widely used for x-ray detection [264], there is currently little interest in developing them for use in low energy quantum optics experiments due to the difficulty of fabricating ultra-low leakage tunnel junctions and the lack of an easy multiplexed readout scheme [265].

Microwave Kinetic inductance detectors - MKIDs

MKIDs are another detector which makes use of broken Cooper-pair quasi-particles to detect the absorption of a single photon. In contrast to the STJs, this sensitivity to the number of quasi-particles is measured by the change in kinetic inductance[1] of the material. In an MKID detector, the superconducting absorber material is part of a microwave resonator circuit,

[1] kinetic inductance is essentially a measure of the 'inertial mass' of the charge carriers in a material

for which a shift in resonant frequency corresponds to a change in kinetic inductance.

The magnitude of the frequency shift can be used to infer the energy of the incoming particle and so provides intrinsic photon number resolution. This read-out scheme also lends itself to large-scale frequency domain multiplexing, enabling large detector arrays to be read-out in parallel with simple microwave components [266]. Because of this, since their invention in 2003 [267], MKIDs have quickly developed into a major detector technology for astronomical instruments from sub-millimetre to x-ray wavelengths [265]. As such, current devices are optimised for broad spectral coverage and long integration times. However, there has recently been some interest in optimising these detectors for use in quantum optics-type photon counting applications [268] and it will be very interesting to see how this work progresses in the next few years.

Superconducting nanowire single photon detectors - SNSPDs

An SNSPD consists of a nanopatterned strip of superconducting metal of typical width $\sim 100\,\text{nm}$ and thickness $\sim 5\,\text{nm}$ biased with a current just below the critical current at which it becomes resistive [269]. A single absorbed photon can induce a temperature change large enough to create a localised resistive hotspot. This causes the current density around the sides of the hotspot to exceed the critical current causing a resistance to develop along the nanowire. The voltage drop across this resistance constitutes the detection signal which can be read out with a simple room-temperature amplifier.

The hotspot formation is extremely quick because of the small size of the wire, leading to exceptionally low timing jitters ($\sim 30\,\text{ps}$ [270]). Many of the original nanowire devices were fabricated using niobium nitride (NbN) wires with detection efficiencies as high as 70%

reported [271] and recent results based on a new tungsten silicide (WSi) devices have achieved detection efficiencies of 93% at 1550 nm, which is surpassed only by TES detectors [272,273].

It is interesting to note that SNSPDs are unique among the superconducting detectors in *not* being able to intrinsically resolve photon number. However, their nano-scale size and meander-like fabrication allows several separate nanowires be interleaved together all within the mode field diameter of a single-mode fibre [255,256,270,271,274]. This lets SNSPDs achieve low-loss spatial multiplexing and devices have been made with 4 bins with individual detector efficiencies of around 30%, mainly due to the difficulties of fabricating SNSNPDs with high detection efficiencies with high yield [256].

The high-speed and high-efficiency of SNSPDs along with their simple room-temperature read-out electronics has made them a popular choice in enabling quantum optics experiments at telecomm wavelengths in recent years [275]. However, if ultimate detection efficiency and high-fidelity number resolution are required then the TES still remains the only available choice.

Transition edge sensors - TESs

The final cryogenic detector we will describe is the subject of the remaining two chapters of this thesis - the transition-edge sensor (TES).

The TES is unique among the other detectors mentioned here in that it operates *within* the superconducting transition region such that it acts as a high dR/dT thermometer. This means that, unlike the SNSPD, the small temperature change associated with a single absorbed photon is able to generate a measurable signal without the full resistive breakdown of the material. The change of resistance is related to the incident photon energy and so

the TES retains photon-number resolution. A major breakthrough in 1995 was the realisation that negative electro-thermal feedback (ETF) could be used to automatically bias the TES within its superconducting transition [276]. Coupled with the adoption of SQUIDs as a convenient, low-impedance amplifier, TESs have become widely popular for astronomical applications [277–282]. Starting in 2003, the group of Sae Woo Nam at NIST began developing the TES as a tool for photon-counting in the near-IR for metrology and quantum optics experiments [283]. There has been rapid progress since with TESs showing the highest ever reported single-photon detection efficiencies [284,285] enabling a number of ground-breaking experiments in quantum physics (conclusive quantum steering [286], violation of Bell's inequality without the fair sampling assumption [287,288]) which would not have been possible with any other detector.

One drawback is that TESs suffer from relatively poor timing performance. The thermal coupling between the TES and the substrate must necessarily be small so that the TES can be heated into its superconducting transition region. However, this small thermal coupling means it takes a long time for the film to cool down again after the absorption of a photon. The large inductance of the SQUID readout electronics have historically limited the timing jitter of TESs to be around 100 ns. However, recent progress in reducing this inductance has seen timing jitters fall toward 1 ns [289]. As we have already argued, this timing performance will only limit the effective clock-rate of the experiment which does not fundamentally affect the scalability of the system.

We present a comparison of the detector technologies so far reviewed in table 5.1. We have not included any results pertaining to STJs or MKIDs since neither of these devices have ever been used for quantum optics photon-counting experiments. Note that for many

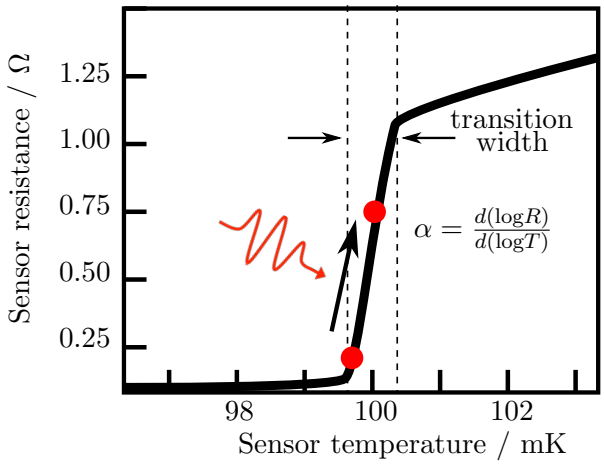


Figure 5.1: The superconducting-to-normal transition region. A TES makes use of the large sensitivity of the superconducting transition to convert minuscule changes in the temperature of the film into measurable changes in its resistance. The logarithmic derivative, α, is a dimensionless measure of this slope. By resolving how much the resistance has changed we can infer how many photons were absorbed by the detector.

of the detectors listed there are several different results. This is because detector design almost always involves some compromise on certain parameters. What is important is what detector parameters may be achieved *simultaneously* for each technology.

5.3 Transition edge sensors

Throughout this section we will introduce a range of new symbols. Although all relevant symbols are defined as they are needed, a convenient reference of all symbols used in this section is provided in Appendix B.

5.3.1 Operating principle

A TES is a micro-calorimeter composed of four components: an efficient absorber, a small heat capacitance, a sensitive thermometer, and a weak thermal link to restore the device back to its initial temperature. One of the more significant discoveries in the history of TES

	Measured detector parameters							
Detector type	Material	η	Δt_{FWHM}	τ_{dead}	D.C.R.	PNR?	$T(K)$	ΔE
PMTs								
PicoQuant PMA	-	45% @500 nm	50 ps	100 ns	100 Hz	No	250	N/A
APDs								
Laser Components, Count NIR	Si	65% @810 nm	750 ps	45 ns	50 Hz	No	250	N/A
idQuantique id220	InGaAs	25% @1550 nm	250 ps	10 μs	1 kHz	No	200	N/A
MPD InGaAs SPAD[290]	InGaAs	35% @1550 nm	150 ps	30 μs	15 kHz	No	233	N/A
TESs								
A. J. Miller et al. (2003)[291]	W	20%† @1550 nm	100 ns	15 μs	0.01 Hz	15	0.125	0.12 eV
A. E. Lita et. al. (2008)[284]	W	95% @1550 nm	100 ns	800 ns	-	7	0.178	0.29 eV
B. Calkins et. al. (2011)[292]	W/Au	> 90%†† @1550 nm	-	460 ns	-	4	0.150	0.24 eV
A. L. Linares et al. (2013)[289]	W	> 90%†† @1550 nm	4 ns	759 ns	-	4	0.150	0.33 eV
D. Fukuda et al. (2011)[293]	Ti	98.4% @850 nm	25 ns	1.7 μs	0.6 Hz	8	0.272	0.39 eV
SNSPDs								
V. .B. Verma et al. (2014)[273]	WSi	78% @1310 nm	191 ps	-	530 Hz	No	2.5	N/A
F. Marsilli et al. (2013)[294]	WSi	93% @1550 nm	150 ps	40 ns	1 kHz	No	0.12	N/A
D. Rosenberg et al. (2013)[271]	NbN	68% @1550 nm	80 ps	5 ns	2 kHz	4†††	2.5	0.02†††
E. Dauler et al. (2009)[270]	NbN	25% @1550 nm	30 ps	-	-	-	2.7K	

Table 5.1: Comparison of detector technologies. A list of the state-of-the-art for various single-photon detector technologies currently in use in quantum optics experiments. Achieving optimum performance across a large array of parameters is often not possible and some compromise must usually be made in optimising for a specific application. Therefore, we list several different results so it is possible to see what can currently be achieved simultaneously. NOTES: † - Experiment performed before the adoption of AR-coatings and optical cavities. †† - although detector efficiency is not quoted in these papers, private conversations with the authors suggest they should be > 90% @1550 nm. ††† - A 4-bin spatially multiplexed detector formed using interleaved meanders. The energy resolution is the calculated detection fidelity for four photons based on the quoted detection efficiency.

development was finding that tungsten (W) was able to provide all of those functions in a single device[276,282].

In a W-TES, the film and substrate are first cooled below the film's transition temperature and the electron system of the W is then heated back up to the transition through voltage biasing (figure 5.2). In this state, the energy from a single absorbed photon, E_γ, causes a temperature change related to the electronic heat capacity $\Delta T = E_\gamma / C_e$. The minuscule material heat capacity at low temperature ($C_e \propto T$) results in an incredible temperature sensitivity to the small absorbed energy pulses. The ultra-steep resistance curve ($d(\log R)/d(\log T) \sim 100$) of the W in its transition region ensures that this temperature change manifests itself as a large resistance change which can be measured (figure 5.1). The detector then dissipates this heat back into the substrate via the thermal conductance of the link, g, with time constant $\tau_0 = C_e/g$. In a W-TES device, the dominant thermal impedance is due to the strong electron-phonon decoupling of tungsten (see figure 5.2) and typical numbers make $\tau_0 \approx 100\,\mu$s.

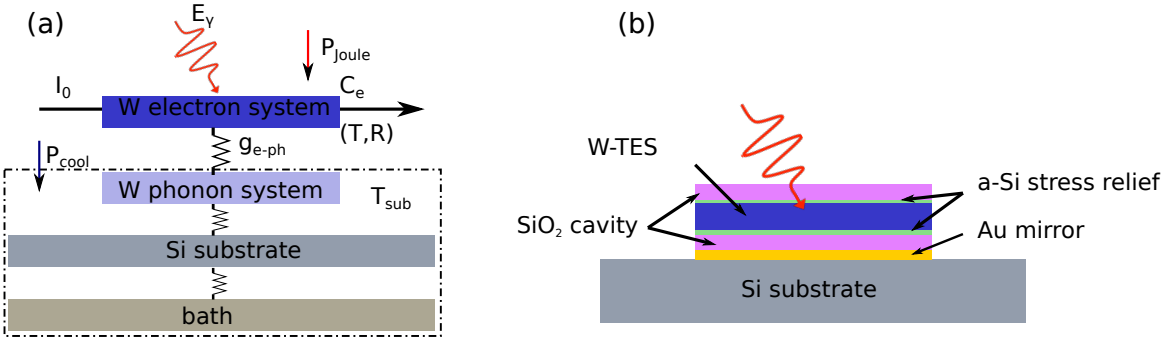


Figure 5.2: Basic operation of a Transition Edge Sensor. (a) The electronic system of tungsten is biased in its transition region via Joule heating, P_{Joule}. An incoming photon heats the electron system by an amount determined by the heat capacity, C_e. The resulting resistance change (see figure 5.1) causes a measurable change in the current flowing through the device, I_0. The sensor is de-coupled from the rest of the substrate via the weak thermal link provided by the strong electron-phonon de-coupling in tungsten. (b) To enhance the detection efficiency, the tungsten is embedded within an optical cavity with a gold backing mirror. The cavities are optimised for specific operating wavelength ranges and can achieve upward of 99% absorption.

5.3.2 Electro-thermal feedback

A major advancement in TES design came with the ability to position the W within the very narrow (few mK) transition region using a self-regulating electro-thermal feedback (ETF) effect. The ETF technique actually serves two very important purposes, the first allows us to very simply bias the TES at the correct temperature by injecting power directly into the electronic system (and not the phonon system which we wish to use as a heat sink) and the second helps to dramatically reduce the thermal recovery time from $\tau_0 = C_e/g$. The ETF technique uses the sensor's own Joule heating to raise its temperature. If the sensor is held under a constant voltage bias, the Joule heating, $P_{\text{Joule}} = V^2/R$ provides negative feedback. If the temperature of the film rises, the resistance increases which causes a reduction in heating power. Conversely, if the temperature were to drop, then the resistance will fall and the increased Joule heating will act to push the detector back toward the quiescent point. This self-regulating feedback mechanism means there is no need for precise thermal control via external means and simplifies device operation greatly [277].

A simple model

In the steady state, the temperature of the sensor is determined by the competition between Joule heating and the cooling power of the substrate:

$$\frac{dE_{\text{sensor}}}{dt} = -P_{\text{cool}} + P_{\text{Joule}} \tag{5.2a}$$

$$C_e \frac{dT}{dt} = -\kappa(T^n - T_s^n) + \frac{V^2}{R(T)} \tag{5.2b}$$

where T is the temperature of the sensor (around 150 mK) and T_s is the temperature of the substrate (around 50 mK), n is a variable exponent which is approximately 5 for the systems considered here[279] and κ is the thermal coupling constant, $\kappa = g/(nT_0^{n-1})$. The steady-state quiescent temperature of the device, T_0, is given by the condition $V^2/R_0 = \kappa(T_0^n - T_s^n)$. When an incoming photon of energy, E_γ, is absorbed by an electron in the tungsten, the temperature of the entire electron system system rapidly (few ps) rises by $\Delta T = E_\gamma/C_e$. We can expand equation (5.2b) to first order in ΔT,

$$\begin{aligned} C_e\frac{d\Delta T}{dt} &= \frac{V^2}{R_0} - \kappa(T_0^n - T_s^n) - \frac{V^2}{R_0^2}\frac{dR}{dT}\Delta T - n\kappa T_0^{n-1}\Delta T \\ \rightarrow \frac{d\Delta T}{dt} &= -\left(\frac{V^2}{T_0 C_e R_0}\alpha + \frac{g}{C_e}\right)\Delta T, \end{aligned} \tag{5.3}$$

where we have defined $\alpha = d(\log R)/d(\log T)$ as the logarithmic 'steepness' of the transition curve and $g = n\kappa T_0^{n-1}$ as the thermal conductance. This differential equation has a simple exponential solution with decay time constant,

$$\tau_{etf} = \tau_0\left(1 + \frac{\alpha}{n}\left(1 - \frac{T_s^n}{T_0^n}\right)\right)^{-1}. \tag{5.4}$$

Under ETF the sensor temperature is typically around twice that of the substrate and so $T_s^n/T_0^n \ll 1$ and α is typically around 100 allowing us to simply write,

$$\tau_{etf} \approx \frac{n\tau_0}{\alpha} \tag{5.5}$$

from which it is apparent that the effect of ETF is to reduce the recovery time of the sensor by a factor of around 20 or more. This works because the Joule heating from the ETF is

the dominant power source into the sensor and anything which causes a reduction in this leads to a rapid fall in temperature.

Small changes in temperature lead to small changes in current,

$$\Delta I = \Delta\left(\frac{V}{R}\right) = -\frac{V}{R_0^2}\frac{dR}{dT}\Delta T = -\frac{\alpha I_0}{T_0}\Delta T, \tag{5.6}$$

so we can write the full solution for the exponential decay in terms of the initially deposited photon energy, E_γ, and the resulting change in current through the sensor,

$$\Delta I(t) = -\frac{\alpha I_0}{T_0}\frac{E_\gamma \epsilon}{C_e} e^{-t/\tau_{etf}} \tag{5.7}$$

where ϵ is the fraction of the photon energy which is transferred to the electron system (as opposed to the phonon system) and is between $0.4 < \epsilon < 0.9$[289] for W-TES films (Note that this is not related to the detection efficiency of the device). Since the decay time is independent of photon energy (at least in the small signal limit we are considering here) it is possible to measure the input photon energy by recording the reduction in current allowing simple threshold electronics to determine the photon number.

5.3.3 Readout

Detection and amplification of this current signal requires a suitable amplifier which has a noise floor below that of the TES, enough dynamic range to provide sufficient signal output for further data processing, and can be impedance-matched to the sensor. Traditional low-noise cryogenic amplifiers based on FETs are not well-suited to the low-impedance output of a W-TES[282]. Instead, to detect this current change through the sensor we inductively

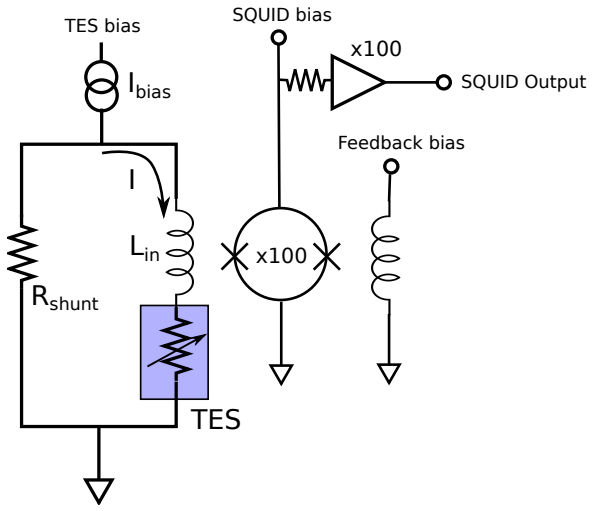


Figure 5.3: Schematic of the TES bias circuits. The TES is held under a stiff voltage bias by placing it in parallel with a bias resistor. This parallel circuit is supplied by a current source which controls the voltage bias of the TES. The TES is in series with an input coil which wraps around a SQUID array and converts changes in current through the detector into a changing magnetic field through the SQUID. The SQUID is biased through a bias-tee circuit and the output voltage is subsequently amplified and filtered. The feedback bias circuit adjusts the local magnetic environment of the SQUID to ensure it operates within its most sensitive region.

couple it to a Superconducting QUantum Interference Device (SQUID) as seen in figure 5.3.

SQUIDs

A SQUID is a very sensitive magnetic flux to voltage converter [295]. It consists of a superconducting ring with two Josephson junctions in parallel and can be thought of the electronic analogue of the Mach-Zehnder interferometer. Because a superconducting loop may only be threaded by integer multiples of magnetic flux quanta, when the external magnetic flux threading the loop is varied a supercurrent flows through the SQUID to oppose this change. If the loop is biased close to its critical current, this extra screening current is sufficient to make the two junctions resistive and develop a voltage drop. The voltage drop across the SQUID varies sinusoidally with the magnetic flux threading the superconducting loop. Each period of oscillation corresponds to one magnetic flux quantum entering or leaving the

loop which is allowed to happen when the supercurrent is sufficient to make the entire loop normal for a small time.

Because the magnetic flux quantum is small, SQUIDs serve as incredibly sensitive magnetometers. Great advances have been made in SQUID design over the past 20 years allowing large series arrays of 100 SQUIDs to be fabricated. This increases the flux-to-voltage gain that can be achieved, facilitating easy matching to further amplifiers [296]. It should be noted that because of the periodic nature of SQUIDs, the range over which the voltage varies linearly with flux is rather limited. For applications which require a large dynamic range, the SQUIDs are operated in a flux-locked loop (FLL) mode, meaning the signal from the SQUID is fed back to a coil nearby, keeping the SQUID in a state of constant flux in the most linear part of the response curve [297]. However, such feedback loops severely limit the bandwidth of operation which is often detrimental for photon-counting work [298]. For all of the work carried out here, we operated in 'open-loop' mode which not only simplifies the detection electronics but allows us to collect data at far higher bandwidths with the caveat that we expect some non-linear response as the number of absorbed photons increases.

The TES is placed in series with an 'input coil' which is wrapped around the SQUID (see figure 5.3). This efficiently couples the changing current through the TES into a measurable change in magnetic flux threading the SQUID. In order to achieve ETF the TES must be held under constant voltage bias which is achieved by using a room-temperature current-source driving a small cold shunt resistor placed in parallel with the input coil and TES. Because the input coil is part of the monolithic SQUID PCB and is superconducting it does not have any DC voltage drop across it. If we also ensure that $R_{\mathrm{shunt}} \ll R_{\mathrm{sensor}}$, then we can create a stiff (i.e. stable) voltage bias across the TES. In practise, the room temperature

current source is a voltage source connected to a large series resistance. Adjusting this room temperature resistance sets the voltage bias across the TES and thus the operating point of the detector. A schematic of this bias circuit is shown in figure 5.3 and forms one of three bias circuits which are necessary to run the detectors. The second circuit is used to bias and read out the SQUID signal itself. We apply a current through the SQUID via a room-temperature bias-tee which is set close to the critical current. The voltage drop across the SQUID series array is capacitively coupled to a two-stage low-noise FET pre-amplifier (20 dB total gain) before passing through a selection of filters and further amplification and finally into a computer-interfaced DAQ system described in section C.3. The final bias circuit controls what is known as the 'feedback-coil'. This is another coil wrapped around the SQUID and can be used to provide a feedback signal to linearise the array if operated in FLL mode. In our system it is used to fine-tune the local magnetic environment of the SQUID such that we are operating close to the linear (most sensitive) part of the voltage-flux response curve.

Limits to signal rise-time

Once a photon is absorbed by the detector, the resistance increases on a timescale of only a few picoseconds[277,282]. This causes extra current to be shunted through R_{shunt} causing a change in magnetic field threading the SQUID. However, the finite inductance of the input coil prevents this change in current taking place instantaneously, limiting the rise time of the current pulse. This rise time is set by $\tau_{\text{rise}} \sim L/R$ and for typical values of input coil inductances ($L \sim 100\,\text{nH}$), this sets rise times on the order of $\tau_{\text{rise}} \sim 100\,\text{ns}$. Coupled with noise on this rising edge, this limits the achievable jitter on the TES as discussed in

section 5.4.3.

In fact, this finite input coil inductance can also lead to unstable oscillatory behaviour in the circuit. Energy is able to oscillate between the magnetic field of the inductor and the heat capacitance of the TES, making a resonant LC circuit. If the TES is biased too close to the edge of the transition region, τ_{etf} and τ_{rise} become similar and we can get under-damped oscillations in the pulse signal eventually leading to full self-sustaining 'electro-thermal' oscillations which ruin the detection signal[297].

Room temperature amplifiers

The final data acquisition is performed by a PCI-based 14-bit digitiser (Alazar 9440). This digitiser requires signals of the order of millivolts compared to the $\sim 10\,\mu$V level signals which are generated by the SQUID array. At this level, there is significant excess RF-noise from the environment which could couple into our signal and degrade the signal-to-noise ratio. To prevent this, there is a low-noise pre-amplifier placed on the same circuit as the bias-electronics. Ideally, the measurement cable is extremely well shielded up until this point to ensure that we minimise any RF-pickup before amplification. After the bias electronics, the signal is further amplified and filtered by a dedicated ultra-low noise amplifier (LeCroy DA1855A) to a level which can be passed directly to the DAQ electronics.

5.3.4 Noise

In this section we consider the effect of noise on the achievable energy resolution of a TES. We will not consider any 'excess' noise sources due to imperfect experimental conditions such as noisy amplifiers, RF pickup, microphonics in the leads, contact resistance fluctuations, Johnson noise in the leads or bath temperature fluctuations etc. Even without these

excess noise terms, the energy resolution of a TES is limited by the thermodynamic energy fluctuations in the detector (thermodynamic fluctuation noise - TFN or phonon noise), and the Johnson noise of the film. For simplicity, we first derive the results in the absence of electro-thermal feedback and then argue that ETF actually improves the energy resolution from what one would naively expect. Because different noise sources tend to manifest themselves in different ways as voltage, current or power fluctuations both before and after amplification, we follow convention and express everything referenced to an equivalent current noise on the SQUID input coil as shown in figure 5.4. This makes it straightforward to compare the effects of each noise source.

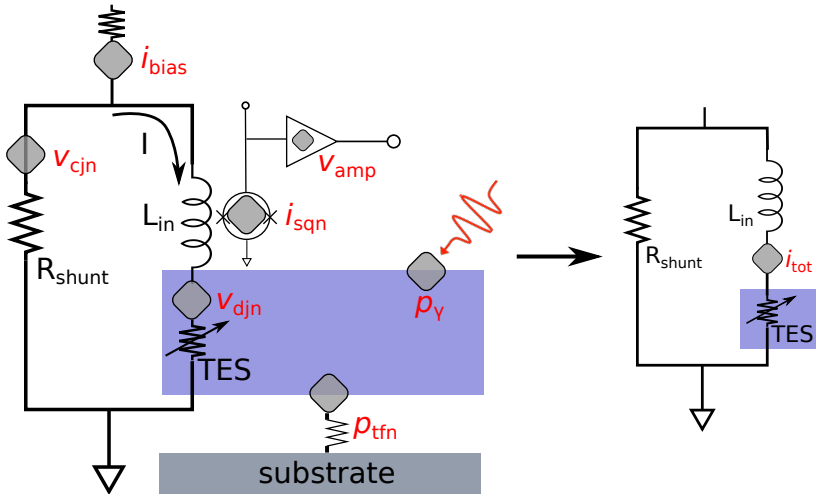


Figure 5.4: Noise sources in a TES. There are two intrinsic noise sources in the TES, the Johnson noise associated with the operating resistance, v_{djn} and the thermodynamic fluctuation noise due to the power fluctuations across the weak thermal link, p_{tfn}. As can be seen a variety of other voltage, current and power noise sources exist in other places. In order to easily compare the different noise sources we use the equivalent circuit on the right where all of the noise is generated by a single current source on the input coil. All of the noise figures quoted for the rest of this chapter are referenced to this position.

Let us first consider the thermodynamic fluctuation noise (TFN) due to phonon fluctuations in the weak thermal link between the biased film and the colder substrate which contribute changes in the overall resistor power. In the limit of no ETF, Joule heating is

negligible and the film temperature is close to that of the substrate, so that we may write the power flow equation for the sensor simply as,

$$C_e \frac{dT}{dt} = -g(T - T_s) + P_{\mathrm{TFN}}(t), \tag{5.8}$$

where $P_{\mathrm{TFN}}(t)$ is the power flow from the substrate due to phonon noise. Looking at small variations around the operating temperature and performing the Fourier transform we find,

$$i\omega C_e \Delta T_\omega = -g\Delta T_\omega + P_\omega, \tag{5.9}$$

where the ω subscript specifies a particular spectral component. We can convert to current fluctuations using equation (5.6) to achieve,

$$\Delta I_\omega = -\frac{I_0 \alpha}{T_0} \frac{P_\omega}{g + i\omega C_e}. \tag{5.10}$$

The power spectral density of phonon noise is given by[278] $P_\omega^2 = 4k_B T_0^2 g$ which means we can write the spectral distribution of current fluctuations through the detector due to TFN as,

$$\left\langle |\Delta I_{\mathrm{TFN}}(\omega)|^2 \right\rangle = \frac{4k_B \alpha^2}{R_0} \frac{(T_0 - T_s)}{1 + \omega^2 \tau^2}, \tag{5.11}$$

where $\tau = C_e/g$ is the intrinsic time constant. This noise is flat at low frequencies but rolls off at higher frequencies with a knee at $\omega = 1/\tau$.

Recall from equation (5.7) the form of our signal pulse decays as $\sim e^{-t/\tau}$ which, after

performing the Fourier transform, gives the power spectral density of our signal,

$$\left\langle |\Delta I_\gamma(\omega)|^2 \right\rangle = \frac{\alpha^2 I_0^2 E_\gamma^2 \epsilon^2}{T_0^2 C_e^2} \frac{\tau^2}{1+\omega^2\tau^2}, \tag{5.12}$$

which has a very similar form to that of the phonon noise since both are governed by the weak thermal link. However, if this were the only noise source, the signal-to-noise ratio would be white and, in theory, arbitrarily small energy resolution could be obtained by simply increasing the measurement bandwidth. In reality, of course, we have an additional noise term from the Johnson noise associated with the resistance of the detector. Without ETF, Johnson noise is well approximated as white and has a constant power spectral current density given by the usual expression,

$$\left\langle |\Delta I_{\text{Johnson}}(\omega)|^2 \right\rangle = \frac{4k_B T_0}{R_0}. \tag{5.13}$$

Since these two noise sources are uncorrelated, we can simply add them together in quadrature to get the total noise as plotted in figure 5.5.

A more involved analysis to the one above includes the effect of strong-ETF, where the power spectral densities of the input signal and phonon noise now take the modified form[277],

$$\left\langle |\Delta I_{\text{TFN}}(\omega)|^2 \right\rangle = \frac{4k_B T_0}{R_0} \frac{n/2}{1+\omega^2\tau_{etf}^2} \tag{5.14a}$$

$$\left\langle |\Delta I_\gamma(\omega)|^2 \right\rangle = \frac{I_0^2 E^2 g^2}{T_0^2 n^2} \frac{1}{1+\omega^2\tau_{etf}^2} \tag{5.14b}$$

which are very similar to the equations (5.11) and (5.12) already derived. However, their

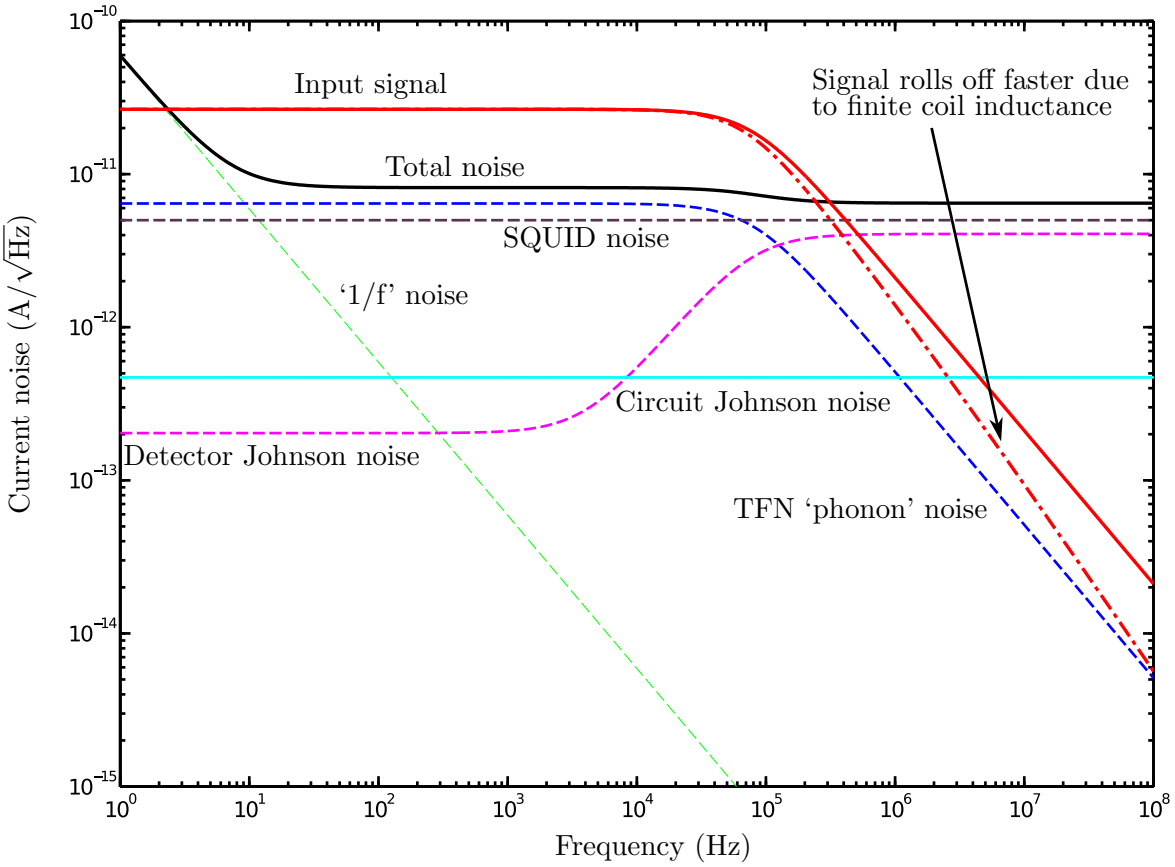


Figure 5.5: Noise contributions for our TES detector. Using estimated detector parameters and the equations derived in section 5.3.4 we plot the frequency response of the various noise sources within a TES detector. All the noise is referenced as a current noise applied to the input coil of the SQUID. The '1/f' pink noise is estimated based on measurements for our device. The SQUID noise is estimated based on the specifications provided by NIST. The size of the input signal is taken from characterisation results of section 5.4. The dashed line for the signal noise shows the effect of the finite input coil inductance in limiting the rise time of the detector. We can clearly see the suppression of detector Johnson noise due to ETF at low frequencies. We apply a low-pass filter at 1 MHz in our experiments to reduce the impact of the higher frequency noise whilst retaining the important parts of the signal.

magnitude is no longer dependent on α and we have replaced $\tau \rightarrow \tau_{etf} \approx C_e n/g\alpha$. It can also be shown[282] that ETF suppresses the detector Johnson noise by a factor of n^2/α^2 at low frequencies, which can be several orders of magnitude for high-α films. Thus, it is often assumed that Johnson noise is negligible for TESs operated with strong ETF. At high frequencies the ETF cannot respond fast enough to the thermal fluctuations and the Johnson noise approaches its white noise value.

We can now work out the achievable energy-resolution for such a noisy detector. It can be shown[277] that this is related to the current noise we have just derived by,

$$\Delta E_{RMS} = \left(\int_0^\infty \frac{4df}{[NEP(f)]^2} \right)^{-1/2} \tag{5.15}$$

where $NEP(f)$ are the noise equivalent power fluctuations that would cause the current noise given in equations (5.11 — 5.13). After performing the integral in equation (5.15) we arrive at the thermodynamic fundamental energy resolution of the detector,

$$\Delta E_{\mathrm{FWHM}} = \frac{2.36}{\epsilon} \sqrt{4k_B T_0^2 C_e \sqrt{n/2}/\alpha}. \tag{5.16}$$

We can express this in more convenient terms by defining the maximum amount of energy that ETF can remove as the saturation energy of the detector, $E_{\mathrm{sat}} = P_0 \tau_{etf}$. Knowing that the maximum sensor power occurs when the detector is at the transition temperature and with $T_0^n \gg T_s^n$, we have $E_{\mathrm{sat}} = C_e T_0 \alpha^{-1}$, so that we can write,

$$\Delta E_{\mathrm{FWHM}} = \frac{2.36}{\epsilon} \sqrt{4k_B \sqrt{n/2} T_0 E_{\mathrm{sat}}}. \tag{5.17}$$

This result is interesting when compared to other superconducting detectors such as STJs or MKIDs. For these devices, the fundamental energy resolution scales like $\sqrt{E_\gamma}$. This means that the fractional resolving power, $E_\gamma/\Delta E$, scales better for higher photon energies in a TES than it does in an STJ or MKID. It turns out that the crossing point for these detectors occurs around 0.8 eV. For photon energies higher than this, the TES has a fundamentally better resolving power making them better-suited for near-IR and visible-light applications than other comparable technologies. For approximate detector parameters of $T_0 = 150\,\mathrm{mK}$, $E_{\mathrm{sat}} = 11\,\mathrm{eV}$ and for $n = 5$ and $\epsilon \approx 0.4$ we expect $\Delta E_{\mathrm{FWHM}} = 0.18\,\mathrm{eV}$. If we could achieve this limit, photon number peaks at 1550 nm (0.8 eV) would be separated by more than 10 standard deviations, providing exceptional number-resolving capability.

5.4 Characterisation of TES detectors

In Appendix C we describe the setup of the cryogenic system and subsequent installation of the TES detectors in our lab. With the fridge wired up and the detectors installed we could begin characterising our new detection system. First, the bias points must be set correctly which can be done as follows: (1) Drive the 'sweep' input of the feedback bias line with a triangular wave ($1\,\mathrm{V}_{p-p}$, 100 Hz) which has the effect of modulating the SQUID output voltage — figure 5.6(a). (2) Adjust the SQUID bias current to maximise the modulation depth of this oscillation [2]. This sets the bias current to just below the critical current. (3) Apply the triangle-wave sweep voltage to the TES input and observe that the SQUID oscillations occasionally disappear as the TES is driven through its superconducting transition — figure 5.6(b). (4) Adjust the feedback bias current to position the transition region on

[2] The maximum modulation depth is achieved when all the individual SQUIDs in the series array add in phase.

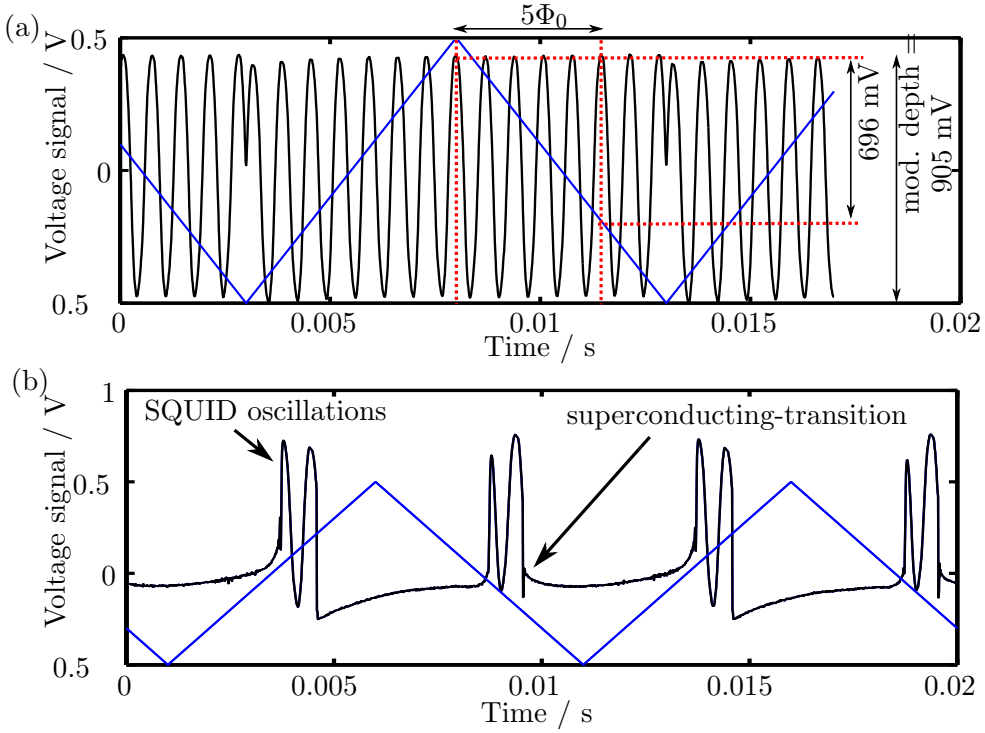


Figure 5.6: Tuning the SQUIDs. (a) Applying a triangle wave to the feedback coil, the SQUID response oscillates between many magnetic flux quanta. Using the known series resistance in the feedback bias line, we can estimate the current required to pass through one magnetic flux quanta. The modulation depth can is maximised by adjusting the SQUID bias current and can be used to estimate the overall SQUID sensitivity in its linear region. (b) A similar triangle driving pulse applied to the TES input coil also causes the SQUID to undergo oscillations. However, when the bias current is enough to drive the TES normal, the current is shunted through the parallel bias resistor and the oscillations disappear. This plot allows us to see where the transition region is and by adjusting the feedback bias whilst looking at this trace we can position the transition region at the linear part of the SQUID response curve.

the most sensitive (linear) part of the SQUID response. (5) Switch off the triangle-wave and slowly increase the TES bias current until you go past the electro-thermal oscillations and a TES pulse signal appears. (6) Fine tune the SQUID and Feedback bias points to maximise the signal-to-noise ratio.

5.4.1 Characterisation of detector noise

Using the information contained in the SQUID response shown in figure 5.6, we can derive the input-referenced current noise from measurements made of the voltage noise using the

data acquisition electronics. Driving the feedback coil with a triangular wave, we see the SQUID response pass through many magnetic flux quanta, Φ_0. From figure 5.6(a) we obtain a current sensitivity on the feedback coil of,

$$\frac{dI}{d\Phi} = \frac{(696\,\mathrm{mV}/5\,\mathrm{k\Omega})}{5\Phi_0} = 25.4\,\mu\mathrm{A}/\Phi_0 \tag{5.18}$$

We would like to know this figure for the input coil but this is made more difficult because we only pass through one flux quanta before the TES is driven normal as seen in figure 5.6(b). However, we know that the input coil has 2.5 times more turns than the feedback coil and so should be 2.5 times more sensitive. Looking at the trace in figure 5.6(b), it appears that this is not the case and the sensitivity looks the same. This is because we have some parasitic resistance in series with the TES on the SQUID electronics board. Estimates based on the amount of copper wiring used in the circuit board put this parasitic resistance at $\sim 20\,\mathrm{m\Omega}$. Although this is not enough to cause a problem during normal operation, when the TES is superconducting, it results in around a third less current flowing through the input coil than one would otherwise assume, reducing the apparent sensitivity. We consider the 2.5 turns ratio to be more accurate than the estimated parasitic resistance, so we assume an input coil sensitivity of $\xi_{\mathrm{in}} = dI/d\Phi = 10.2\,\mu\mathrm{A}/\Phi_0$. This means we must apply $10.2\,\mu\mathrm{A}$ to the SQUID input coil to pass through one magnetic flux quantum.

The modulation depth of the SQUID then lets us work out the $\partial V/\partial I$ of the combined SQUID-amplifier combination at the most sensitive part of the curve. It is always worth checking the modulation depth one can achieve when tuning the SQUIDs up. If it is smaller than expected, it is worth 'zapping' them momentarily to free any trapped flux which may

Parameter	Value
T_c	$150\,\mathrm{mK}^\dagger$
T_s	$50\,\mathrm{mK}$
α	$100^\dagger$
τ_{etf}	$2\,\mu\mathrm{s}$
n	$5^\dagger$
R_{TES} - quiescent	$500\,\mathrm{m}\Omega^\dagger$
R_{TES} - normal	$4\,\Omega^\dagger$
R_{bias}	$10\,\mathrm{m}\Omega$
$R_{parasitic}$	$20\,\mathrm{m}\Omega$
Agilent E3620	$350\,\mu\mathrm{V}/\sqrt{\mathrm{Hz}}$ (20 Hz - 20 MHz)
LeCroy DA1855A	$1\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$ (1 Hz-100 MHz)
Bias-box pre-amplifier	$\sim 1\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$

Table 5.2: Estimated detector parameters. Some of the parameters we were able to measure for our particular TES. However the parameters marked with a † we have estimated based on results from other similar devices.

be there and then re-optimising the bias current. From the results in figure 5.6(a) we have,

$$\begin{aligned} V(\Phi) &= \frac{906\,\mathrm{mV}}{2}\sin\left(2\pi\frac{\Phi}{\Phi_0}\right) \\ \rightarrow \left.\frac{\partial V}{\partial \Phi}\right|_{\Phi=0} &= \frac{2\pi(453\,\mathrm{mV})}{\Phi_0} \\ \rightarrow \left.\frac{\partial V}{\partial I}\right|_{\Phi=0} &= \frac{2\pi(453\,\mathrm{mV})}{\xi_{in}} = 280\,\mathrm{mV}/\mu\mathrm{A} \end{aligned} \tag{5.19}$$

with this relation we now know that for every microamp of current flowing through the SQUID input coil we will measure a voltage signal of 280 mV.

From the simple thermodynamic considerations of section 5.3.4 we can make estimates on the expected noise performance of our detectors. We list the parameters used in the model in table 5.2, and calculate the input current referenced noise in table 5.3. These figures are expressed as a white noise equivalent, which is a good approximation for frequencies below 1 MHz as seen in figure 5.5. We assume all of the noise sources are independent so that they add in quadrature. We also derive the expected total RMS voltage noise measured

over a 1 MHz bandwidth using the result from equation (5.19) so that we can compare it to our measurements. From the table we see that the intrinsic thermodynamic noise of the TES is expected to be dominant, as desired, but there are significant contributions from the SQUID and room temperature amplifiers also.

To test this noise model we measure the RMS voltage fluctuations we see from the TES signal. We make sure to bias the TES correctly in the transition region and set the feedback bias to position the SQUID in the linear regime. Over a 1 second integration time we measure a noise of $V_{RMS} = 2.8\,\mathrm{mV}$ which is slightly higher than that expected from table 5.3. We can isolate the effect of the intrinsic thermodynamic TES noise by driving the detector normal by increasing the bias voltage. In its normal state, the large sensor resistance causes very little current to flow through that side of the bias circuit, removing the noise contributions from the TES. In this state we measure a noise of $V_{RMS} = 2.1\,\mathrm{mV}$, which implies that the noise due to the TES is $6.6\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$. This is quite close to the value we predicted using our simple model and suggests our estimated detector parameters from table 5.2 are quite good. Modifying our expected noise based on this revised TES figure we find we should now expect 2.5 mV RMS voltage fluctuations when the TES is biased in its transition. The remaining difference between this value and the measured noise can be explained using an excess noise term of $4.2\,\mathrm{pA}/\sqrt{\mathrm{Hz}}$. This suggests we have sufficiently suppressed our external noise sources so that the TES and SQUID noise still dominate the total.

We estimate the saturation voltage by shining a bright laser beam onto our detector such that we drive it normal. This is around 100 mV for our device, and knowing (i.e. being told by NIST) that the detectors should saturate with about 14 photons at 1550 nm (0.8 eV)

we estimate a signal per photon of ∼7 mV (giving $E_{\mathrm{sat}} = 11.2\,\mathrm{eV}$). This gives a signal-to-noise ratio of about 2.5 which corresponds to an estimated energy resolution of $\Delta E_{\mathrm{FWHM}} = (1/2.5) * 0.8 = 0.31\,\mathrm{eV}$. A more accurate calculation using a revised equation (5.15) to account for all of our noise sources yields $\Delta E_{\mathrm{FWHM}} = 0.27\,\mathrm{eV}$. This should allow us to easily resolve photon numbers at 1550 nm. The relative contributions of the amplifier and excess noise can be reduced further by using a SQUID array with a larger $\partial V/\partial I$ sensitivity. In fact, the previous generation SQUID chips from NIST had around a 3x higher sensitivity which would help reduce the impact of the SQUID, amplifier and excess noise on our achievable energy resolution.

5.4.2 Counting photons from coherent states

Most of the pertinent properties of our detectors can be conveniently measured using a pulsed laser. Using a pulsed system makes it easier to analyse the data since we can trigger the DAQ system and it allows us to avoid the problems caused by pulse pile-up by setting the repetition rate to be slower than the recovery time of the device. All of the measurements in the following section were performed at 1550 nm using the TES optimised for that wavelength. Although we had installed a TES suitable for use at 850 nm, during one of the cool-downs, the wire bond to the detector snapped, making it unavailable for characterisation. Because of the higher energy of 850 nm photons we expect the results for that detector to be better than the ones quoted below.

Experimental setup

The experimental setup is shown in figure 5.7. We use a pulse generator to drive a laser diode ($\lambda = 1546.6\,\mathrm{nm}$) at a frequency $R_{\mathrm{rep}} = 50\,\mathrm{kHz}$ with 20 ns pulses which gives a good

Noise Source	Current noise ($pA/\sqrt{Hz}$)
Bias Circuit	
Bias resistor	0.31
Parasitic resistance	0.44
Agilent E3620A voltage source	0.26
Room-temperature amplifiers	
Bias-box pre-amplifier	3.57
LeCroy DA1855A	0.36†
SQUID	
Input current noise	5
TES	
Johnson noise	0.2
Phonon noise	6
TOTALS	
With TES biased	8.6 (2.4 mV_{rms})
With TES normal	6.2 (1.7 mV_{rms})

Table 5.3: List of noise contributions. Each noise source is referenced to a current noise on the TES input coil. The total noise is the quadrature sum of all of the terms. When the TES is driven normal, there is no contribution from the intrinsic TFN or Johnson noise allowing us to separate out these effects. † - Although the two amplifiers have similar noise floors, they contribute differently to the input current referenced noise because they occur at different points in the amplifier chain.

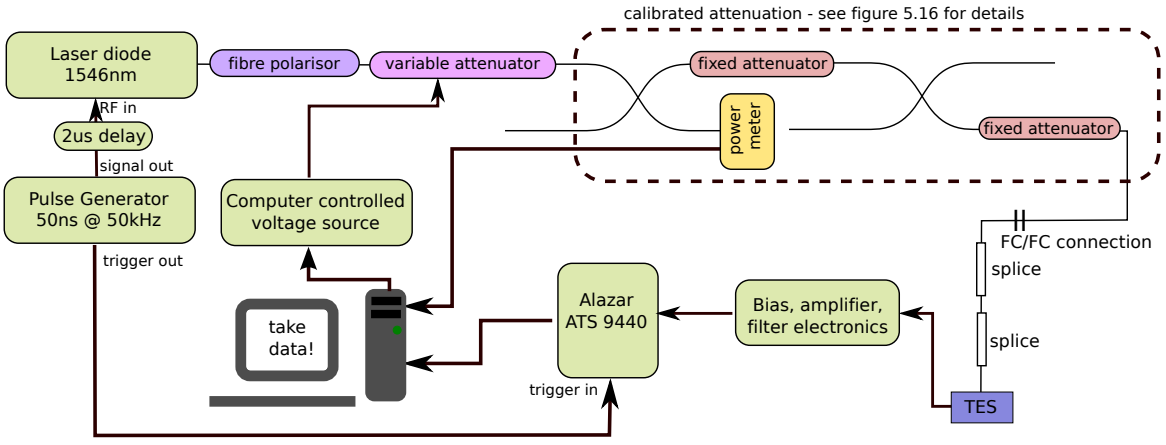


Figure 5.7: Experimental setup. We characterise our detector using pulses of attenuated coherent states from a laser diode. A variable attenuator is used to set the input power before passing through a calibrated attenuation (see figure 5.14) and being coupled to the TES. The data acquisition is triggered 2 μs before the pulse arrives at the TES to ensure we capture the whole of the pulse signal.

trade-off between average power (so we can monitor the power using a photo-diode) and not affecting the timing performance of the detector (we expect the intrinsic jitter to be $\sim$100 ns). We control the power using a computer-interfaced lithium niobate electro-optic modulator. When the insertion loss of the components is taken into account this allows us to control the average output power between $P_{in} = 0.25\,\mathrm{nW} - 50\,\mathrm{nW}$. Adding an additionally calibrated attenuation (see section 5.4.4 for calibration details) of about $\eta_{\mathrm{att}} = -55\,\mathrm{dB}$ allows us to set the average input photon number per pulse to,

$$\langle n\rangle_{in} = \frac{\lambda}{hc}\frac{P_{in}}{R_{\mathrm{rep}}}\eta_{\mathrm{att}} \tag{5.20}$$

$$\longrightarrow \qquad 0.1 \lesssim \langle n\rangle_{in} \lesssim 15 \tag{5.21}$$

This calibrated coherent state was sent to the TES using an FC/FC connector on the external fibre connector box. There are two additional fibre splices within the fridge before we reach the detector. The detector bias points were optimised and we then took several rounds of characterisation data. The attenuation was set using the computer-controlled variable attenuator and the input power to the calibrated attenuation measured using a calibrated power meter (Newport 918D-IG-OD3R). The Alazar DAQ card was triggered using a pulse generator to record the voltage signal from the TES for 8 μs, and after a 2 μs delay, it would pulse the laser diode. At each input average photon number, 16834 voltage traces were taken and stored by the Alazar DAQ card for later analysis. A sample voltage trace for an average input photon number of $\langle n\rangle = 1.1$ is shown in figure 5.8.

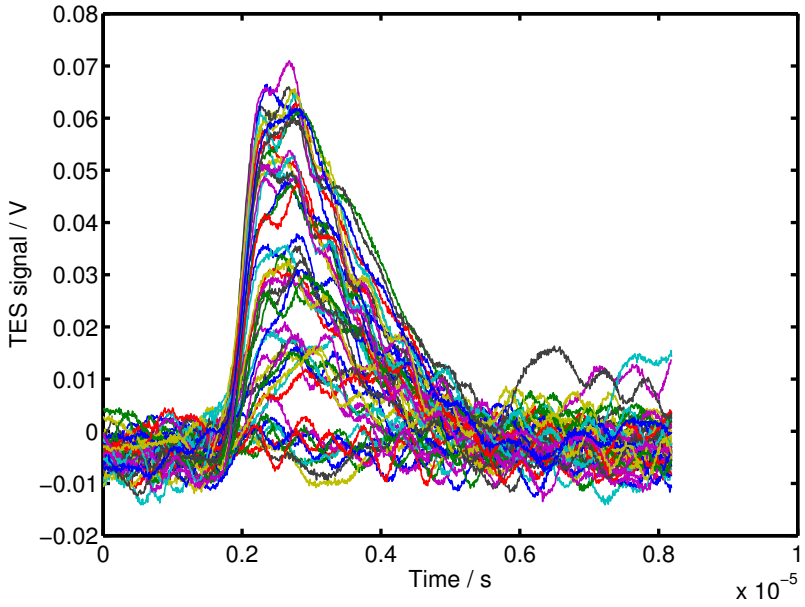


Figure 5.8: Example TES voltage trace. A sample of 150 individual traces have been plotted to the response of the detector to an input coherent state ($\langle n \rangle = 1.1$). The high-frequency oscillations within each trace are due to ringing between the SQUID-wiring-amplifier system. Despite the presence of noise on each trace, it is still possible to clearly see the effects of different photon numbers impinging on the detector.

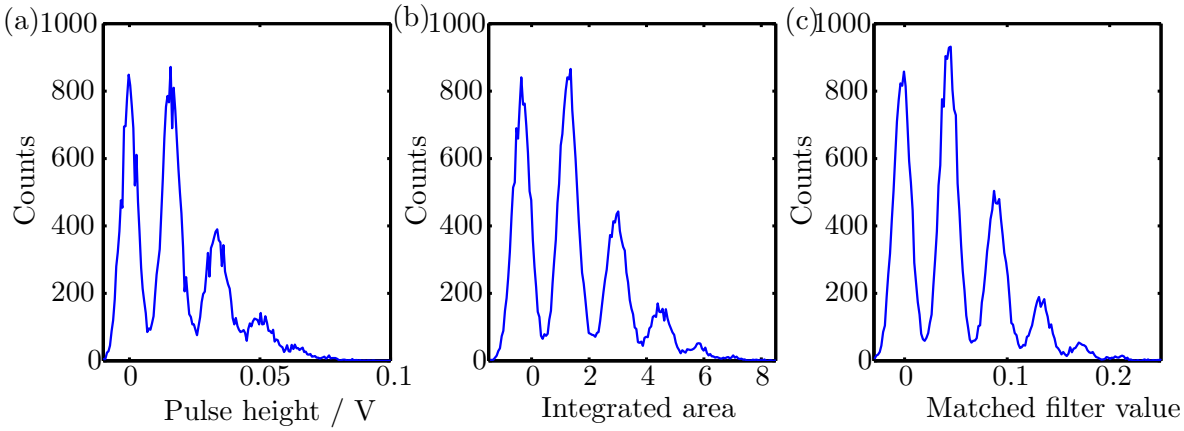


Figure 5.9: Comparison of pulse filtering methods. The signal traces plotted in figure 5.8 are filtered according to one of three methods with the results plotted in the histograms above. (a) filtered using the maximum pulse height (b) filtered by integrating the signal trace in a region around the rising edge (c) filtered using the 'matched-filter' method. In all three cases we can see clear peaks corresponding to well-distinguished photon numbers. The matched-filter method clearly provides the better resolution because the minima in the histogram are closest to zero, indicating fewer traces have ambiguity of peak assignment.

Pulse filtering

For each recorded TES trace, we seek a method to reduce the large dimensional raw data set (a 14-bit voltage signal sampled over 1024 time bins), to a single number specifying the relative energy of the absorbed pulse. From our measured noise spectra, our energy resolution should be good enough to allow us to then assign each trace a photon number. The method of going from the raw voltage signal to a given photon number is called pulse filtering and the aim is to achieve the best possible photon number discrimination for a given data set [299].

We recall from section 5.3.2 that the pulse height should be directly proportional to the absorbed photon energy. The simplest method of pulse filtering is thus to simply find the maximum height of the signal during the rising edge of the pulse. This is a common method used if real-time pulse filtering is required since an FPGA can be used to assign a photon number to a given trace based on a set of pre-defined threshold values [286,300]. We can see from the voltage signals in figure 5.8 that the photon numbers are clearly separated and the histogram of peak heights in figure 5.9(a) shows we obtain a reasonable energy resolution based on this very simple method.

However, this requires careful setting of the output amplitude threshold and is very sensitive to noise. Another method exploits the fact that the total energy removed by ETF is given by the integral of the pulse. Integrating the area under the pulse in this way uses all of the available information recorded in the trace and so can be more tolerant to noise and fluctuations. We found in section 5.3.4 that most of the energy resolution comes from the high frequency components of the signal which in practise means that the fast rising edge contains the most important information. Achieving good photon number discrimination

using this integration method in practise actually requires reducing the integration limits to only a narrow window centred around the rising edge. Thus we do not make full use of the available signal.

Instead, we use a convenient method which computes the overlap (inner product) between each signal trace and a 'template' or 'matched-filter' waveform [301]. This template waveform, $\hat{V}_{\langle n \rangle}$ must be constructed for each data-set of a given mean photon number $\langle n \rangle$ in post processing to achieve the highest sensitivity. One particularly simple template waveform is the mean of all the waveforms in a data-set,

$$\hat{V}_{\langle n \rangle}(t) = \frac{1}{M} \sum_{i=1}^{M} V_i(t), \tag{5.22}$$

where M is the total number of waveforms taken at a given average photon number. Another, more sophisticated, template could be generated using principle component analysis. Computing the dot-product between each waveform and this template yields an overlap which is proportional to the photon energy. A histogram of these overlap values is plotted in figure 5.9(c) which is clearly the better of the three filter methods used. This method is basically the same as the integration technique discussed above but we first weight the data by the average signal-to-noise ratio at each time point. In this way, we weight our signal more at the rising edge as is required without having to sacrifice the information in the rest of the pulse. Further pulse processing can include applying non-linear Weiner filters [302,303], enhanced dot-product routines and K-means clustering algorithms [301]. Although these methods can achieve better energy resolution, especially at higher photon numbers, they are more complicated to implement and are still bound by the same fundamental resolution

set by the intrinsic noise of the device and amplifiers.

Using this method we are able to quite clearly resolve 8 photons as seen in figure 5.10. If we make the reasonable assumption of white noise (see figure 5.5) then we expect a Gaussian distribution of individual photon number outputs. Gaussian fits to the individual photon number peaks are shown in figure 5.10 and are a reasonable fit to the data. We can use these fits to characterise our detector response and cluster our signals into groups belonging to the same detected photon number. However, one must be careful about using this Gaussian fit to assign a photon number uncertainty to traces which occur at the intersection of two peaks. The imperfect assumption of white noise and linear response complicates an accurate calculation of the confidence one can ascribe to a certain trace belonging to a given photon number and is easy to underestimate[301]. We are currently working on a more rigorous detector tomographic routine which is able to better assign such confidence levels.

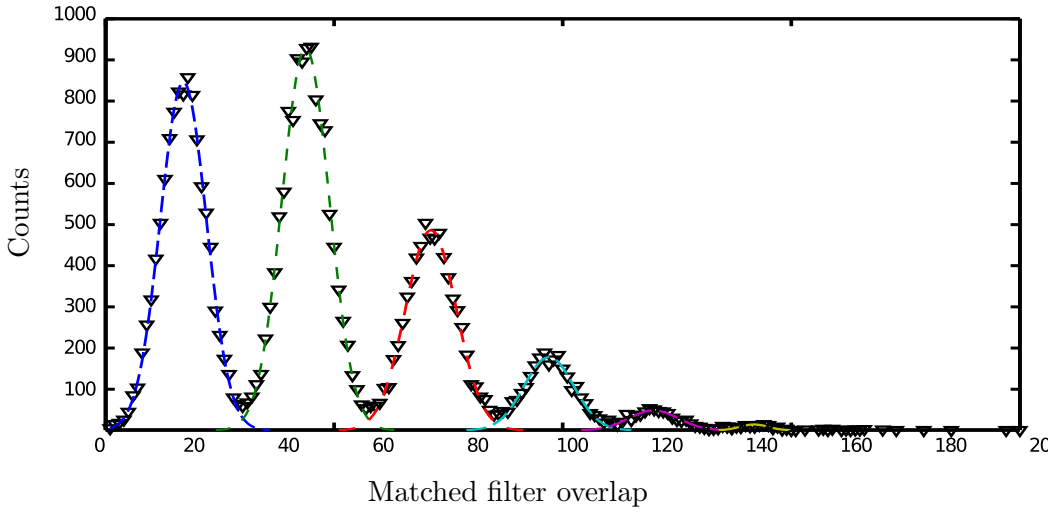
5.4.3 Characterisation results

Energy resolution

Using the results from figure 5.10 we can estimate the energy resolution of our detector from the widths of the fitted Gaussian peaks. We scale our results knowing that the distance between two peaks equals 1 photon of energy (0.8 eV). We plot the results in figure 5.11 for a variety of different input coherent states. The measured energy resolution of the one photon peak is 0.36 eV which is only slightly worse than we expected based on our noise analysis. This means our matched-filter method is already quite close to achieving the 'optimal' resolution.

As expected, the energy resolution does not depend upon the *average* input photon

(a)



(b)

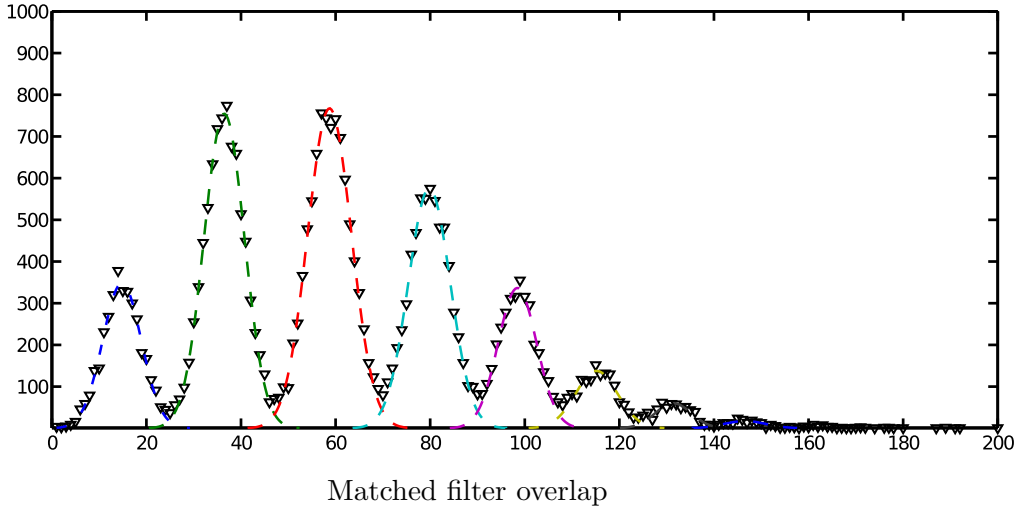


Figure 5.10: Counting photons from coherent states. Using the matched-filter method we plot the resulting histograms for a variety of input coherent states. (a) average input photon number of $\langle n \rangle = 1.1$, (b) $\langle n \rangle = 2.1$. We fit a Gaussian to each of the peaks from which we extract the peak position and width. To look at traces corresponding to only one photon number we bin the signals which occur within 1 standard deviation of a peak mean.

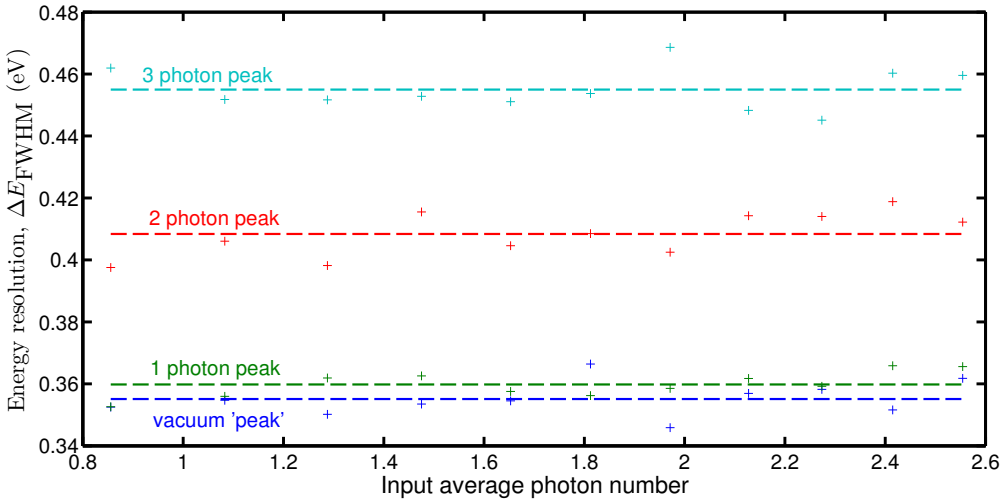


Figure 5.11: Energy resolution. We estimate the measured energy resolution of our detector using the peak widths of the fitted Gaussian curves similar to those in figure 5.10. Each width is scaled by the distance between neighbouring peaks which helps account for the non-linear energy scale. We calculate this resolution for a variety of input coherent states and find that it is independent of the mean photon number.

number but it does get worse with increasing detected photon number. This is due to the non-linear response of the open-loop SQUID. As more photons are absorbed by the TES, we move further away from the linear, and most sensitive, part of the SQUID response curve. This reduces both the TES signal and increases the contribution of external noise - leading to a decrease in observed energy resolution. The non-linear response of the SQUID with photon number is shown clearly in figure 5.12.

Timing jitter

The timing jitter of our detectors was measured by recording the time interval from the laser trigger until the rising edge of the signal crossed some threshold value, V_{th} [293]. For this measurement we sorted the raw voltage traces according to their photon number using the histograms in figure 5.10. To avoid misidentification of traces at the overlap of the Gaussian fits, we sampled traces which were less than one standard deviation away from the peak

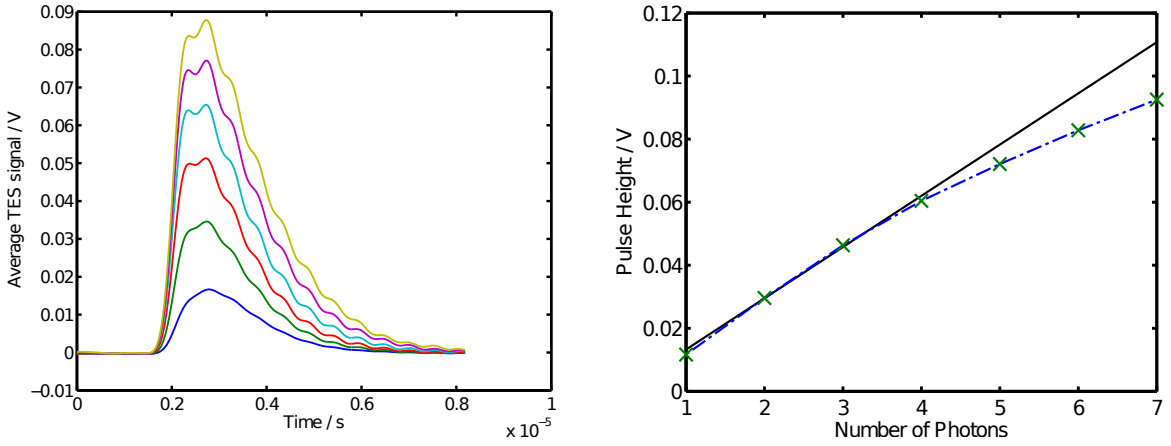


Figure 5.12: Non-linear response of the SQUID array. (a) The sinusoidal response function of the SQUID array manifests itself as a non-linear increase in pulse peak height as a function of absorbed energy. These traces were obtained by taking the mean of all the individual traces which were binned into a specific photon number using the matched-filter method. (b) Plotting the pulse height vs photon number we can clearly see the deviation away from linearity.

height means. In this way, we trade better photon number discrimination for fewer usable data traces.

For any timing pulse signal, the jitter is determined by a combination of the noise and the underlying gradient of the rising edge,

$$\Delta t_{FWHM} \approx 2.35 \frac{\sigma}{V_{\max}} \tau_{\text{rise}} \tag{5.23}$$

where $V_{\max}$ is the maximum amplitude of our photon signal, σ, is the RMS noise and τ_{rise} is the rise-time of our signal[289]. The 10%-90% rise-time of our one-photon signals is $\sim 500\,\text{ns}$. This is limited by both the bandwidth of our amplifier, $\tau_{ext} = 0.35/\Delta f_{amp}$ and the finite inductance of input coil $\tau_{ind} = L/R_0$ which give a combined expected rise time of $\tau_{rise} = \sqrt{\tau_{ind}^2 + \tau_{ext}^2}$. For an amplifier bandwidth of 1 MHz this implies we have around 200 nH of stray inductance.

The results of the jitter measurements are summarised in figure 5.13. For each photon

number, we plot the histogram of crossing times for a defined threshold voltage and fit it to a Gaussian as shown in figures 5.13(a-b). The variation of the FWHM crossing times as a function of the threshold level is shown in figure 5.13(c). As expected, the jitter is lowest when the steepest part of the rising edge is used, which in the case of our detector is at the onset of the pulse. We find there is sometimes a slight dip in these plots as the threshold value first needs to rise above the noise. The FWHM jitter is plotted as a function of photon number in figure 5.13(d) for a variety of input coherent state amplitudes. For the 1 photon signal we find a jitter of 150 ns which falls to around 30 ns for the 5 photon traces since the local steepness of these signals is much sharper.

Using the estimates of our detector noise from section 5.4.1, and the measured rise time of our signal we would expect a jitter of around $\Delta t_{FWHM} \approx (1/2.5) * 500\,\text{ns} = 200\,\text{ns}$. We might expect our actual jitter to be less than this since it only depends on the high-frequency noise components which we know from figure 5.5 are smaller than the white noise average. This jitter will require us to clock our experiments at a maximum rate of $\sim 1\,\text{MHz}$ in order for us to unambiguously distinguish photons from adjacent pulses. This will likely require fast pulse-picking from a standard 80 MHz Ti-Sapphire oscillator. However, we know that this will only contribute a fixed factor of at most 80 in reducing our count rates. As we move to beyond three photon experiments, the extra detection efficiency will easily outweigh any loss from this reduced timing performance.

We can see that this jitter is mostly limited by the amplifier and inductance of our system. Recent results have shown that by using high-bandwidth (20 MHz) electronics and wire-bonding the TES directly to a low-inductance SQUID system the jitter could be

reduced to as low as 2 ns [3] without sacrificing any other detector parameters [289].

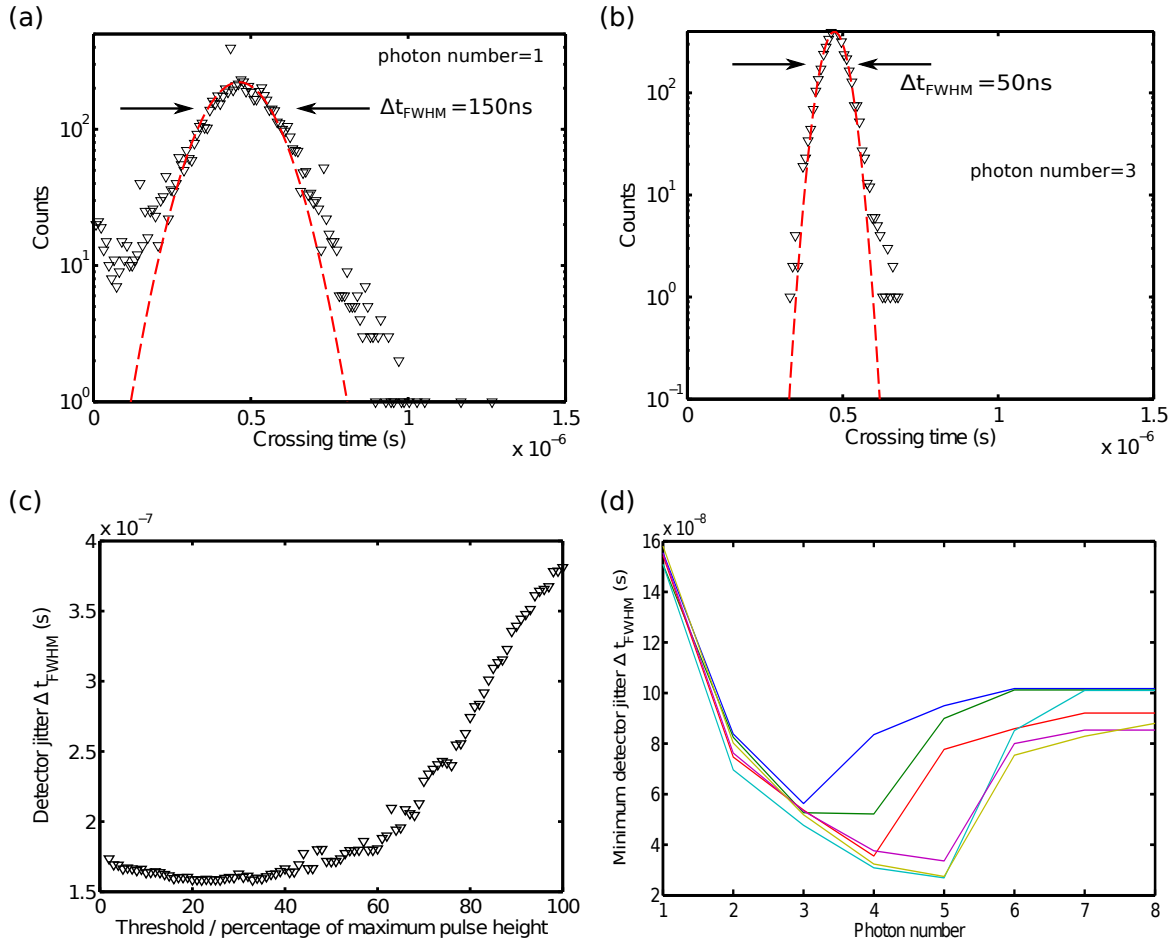


Figure 5.13: Detector jitter. (a) Histogram of pulse crossing times for a 1-photon signal fitted with a Gaussian. (b) The timing uncertainty gets smaller as for the 3-photon signals since the rising edge is steeper. (c) Detector jitter for a 1-photon state as a function of threshold voltage level. The pulse is steepest close to the bottom of the signal which is where the lowest timing jitter is found. (d) Minimum detector jitter as a function of photon number for a selection of input coherent state amplitudes. For the input coherent states with smaller amplitudes (blue, green, red curves) the jitter appears to increase at certain photon numbers because there is insufficient data to get a good fit to the histogram.

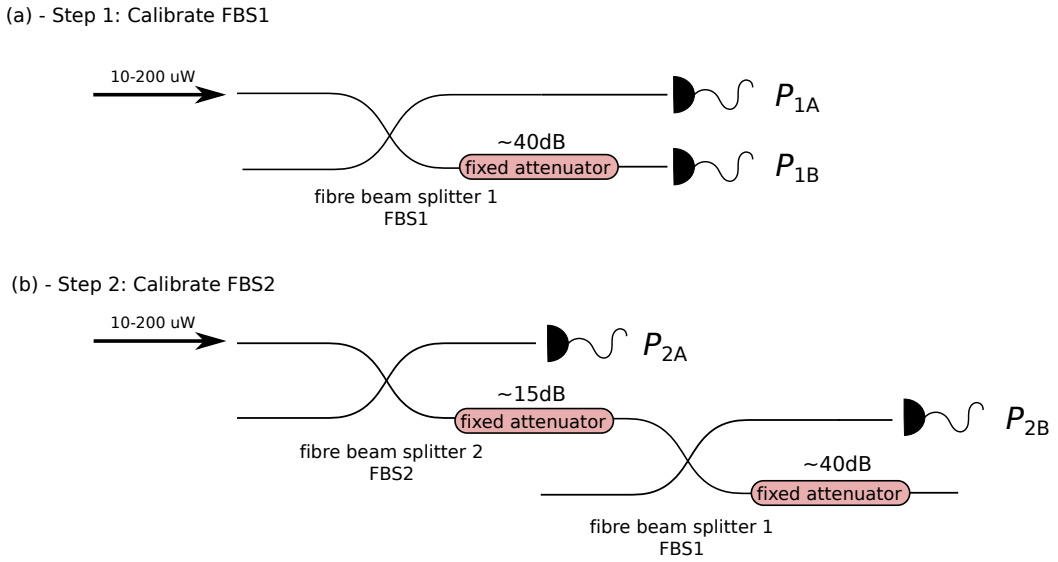


Figure 5.14: Setup to calibrate attenuators. Because the dynamic range of the calibrated power meter is insufficient to span our required attenuation we perform the calibration in two steps at the expense of increased error. (a) We use the laser diode running in CW mode to calibrate the attenuators since the power meter is most accurate in this mode. We first take a series of measurement of the output powers P_{1A} and P_{1B} for a range of input powers to the fibre beam splitter. This lets us calibrate the output power in port $1B$ if we know the power in port $1A$. (b) We connect a second fibre beam splitter to the first one before we calibrate it to make sure the effect of the FC/FC connection is properly accounted for. We then make a series of measurements of P_{2A} and P_{2B} for a range of input powers. Concatenating these results we now know the output power in port $1B$ given the power recorded in $2A$. For the detector characterisation we switch the laser to pulsed operation and attenuate the input light to the nanowatt level.

5.4.4 Estimate of detection efficiency

Detection efficiency is the most important parameter for any photon detector to be used for LOQC. However, it is also one of the most challenging to accurately measure due to the need for some calibrated light source [245,284,285]. Accurate measurements can be made in this way by using a power-meter calibrated to some radiometric standard, which is then used to calibrate some fixed attenuators to reduce the output from a pulsed laser to the single-photon level with a known mean-photon number per pulse. The issue is that calibration errors in the power meter propagate linearly through the process which results in large

[3] A 20 MHz amplifier results in a minimum possible signal rise time of $\tau_{\text{ext}} = 0.35/\Delta f_{amp} = 17.5\,\text{ns}$. The low inductance SQUID system keeps the actual signal rise time to $< 25\,\text{ns}$, which combined with a large signal-to-noise ratio of results in a small signal jitter which has shown to be as low as 2 ns.

overall uncertainties in the measurement. National laboratories such as NIST have made careful measurements of the detection efficiency of TES devices [284] using very carefully calibrated power meters [304] and high-repeatability switchable attenuators.

We do not have access to the precision equipment required to make such a measurement but if we are careful about our errors we can at least bound the detection efficiency of our devices. We use the Newport 918D-IG-OD3R power meter to make our measurements which provides a quoted calibration accuracy of 2% of absolute power and a linearity of better than 0.5 %. We first use this power meter to calibrate a set of fixed in-line fibre attenuators as shown in figure 5.14. The basic idea is to use a fibre beam splitter with a fixed fibre attenuator connected on one of the output ports. So long as this attenuation is well within the linear dynamic range of the power meter we can obtain a calibration curve for this device. However, the attenuation required to reach the single photon level is much greater than the dynamic range of the power meter which forces us to use a second, calibrated fiber beam splitter/attenuator device as described in the caption to figure 5.14. A weighted total least-squares algorithm [305] is used to find the calibrated attenuation properly taking account of the absolute power errors in both variables. The total attenuation is given by the product of the two attenuators, but the errors in the measurements add linearly since they are not independent. Our final calibrated attenuation is accordingly found to be

$$\eta_{att.} = (2.03 \pm 0.15) \times 10^{-6}, \tag{5.24}$$

which relates the power measured at the monitor port $2A$ to the output port $1B$. The variable attenuator is used to set the input power level before the calibrated attenuator so

that we can probe our detector with a variety of coherent state amplitudes. We monitor the input power to the attenuator using port $2A$ and calculate the average photon number per pulse in port $1B$ which is coupled to the TES. The return loss from the unterminated fibre plugged into the monitor power meter will make us underestimate the total power. Fibre specifications give this loss at about 3.3% but there is around a 1% uncertainity in this figure[306]. We estimate the error in the absolute power measured at port $2B$ to be $\sigma_P = 2\%$. Thus our input average photon number per pulse has an uncertainty which is given by the linear sum $\sigma_{\text{tot}} = \sigma_{att.} + \sigma_{RL} + \sigma_P$ which equates to about a 10% uncertainty in the input average photon number.

For each input coherent state we take a total of 32768 triggered TES signals and extract the measured photon number distribution of the data set by fitting a Gaussian to each of the peaks in the matched filter histogram (figure 5.15). The number of counts in each photon number bin are then fit to a Poisson distribution to find the average detected photon number. The fit between the measured counts and a Poisson distribution is typically very good with reduced chi-squared values of around 0.5. We estimate the error in the extracted mean photon number by assuming the error in the counts in each bin is itself Poissonian. This leaves us with an error of around 3% in each detected photon number.

We plot the measured average photon number as a function of input average photon number in figure 5.16 from which we can see the effect of the large errors. Using a weighted least-squares fitting routine we can robustly find the variance of the gradient given the errors in both variables. This gives us a measured detection efficiency of,

$$\eta_{det} = 0.99 \pm {}^{0.01}_{0.13}, \tag{5.25}$$

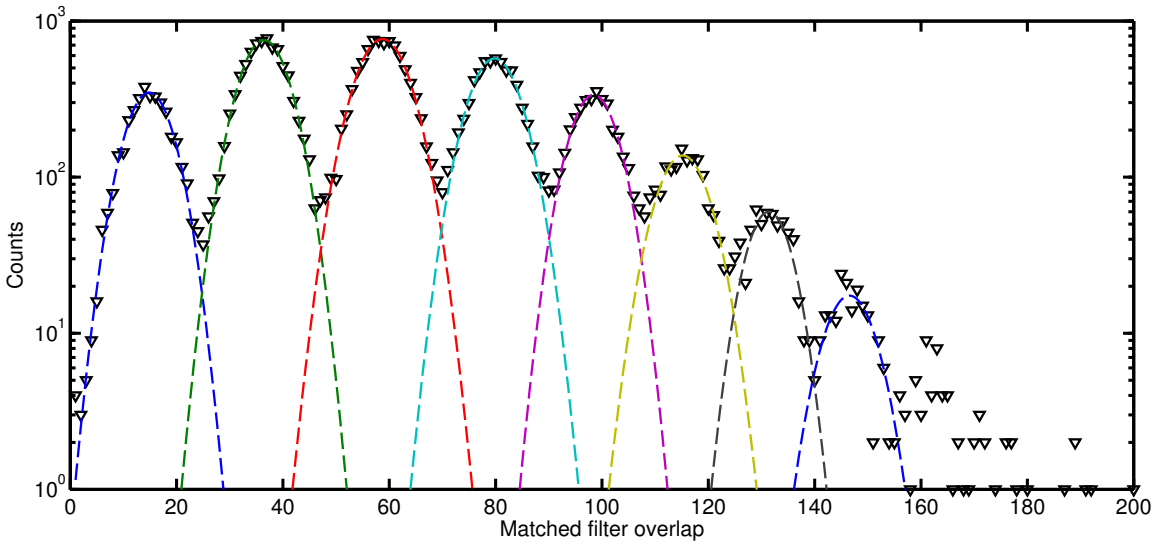


Figure 5.15: Fitting a Poisson distribution. For each input coherent state amplitude we fit Gaussian peaks to the pulse-energy histogram. Counting the number of traces which are within each Gaussian fit we can derive the photon number distribution of our measured signal. This distribution is fit against a Poisson function to extract the detected average photon number.

which includes the efficiency of the FC/FC connector at the fibre breakout box, two fibre fusion splices within the fridge and the efficiency of the TES detector itself. From previous

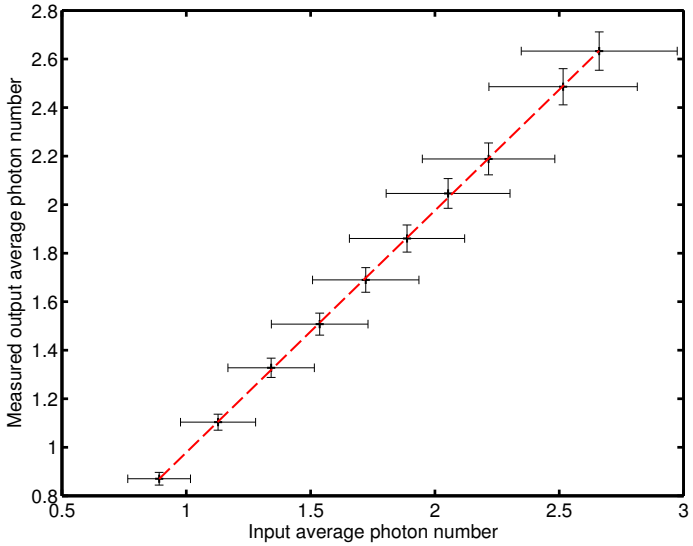


Figure 5.16: Detector efficiency. We plot the measured input average photon number against the measured detected average photon number along with the inferred errors on each measurement. A range of different coherent state amplitudes were used to test the device and we plot a straight line fit to extract an estimate of the detection efficiency.

Parameter	Measured Value
System detection efficiency	$> 86\%$
Energy resolution	0.35 eV
Saturation energy	14 photons (11.2 eV)
Timing jitter, Δt_{FWHM}	150 ns
Recovery time, τ_{etf}	1.1 μs

Table 5.4: Summary of measured detector parameters.

measurements at NIST we expect the TES to show close to 99% efficiency by itself[284] but the two splices and the FC connection should reduce this to around 95%. Our results are obviously consistent with this but the mean is likely too high since we have not accounted for the effects of dark counts. A summary of the characterised detector results is given in table 5.4

Chapter 6

On-chip TESs

THE integration of single photon sources and linear optical circuits onto waveguide chips has seen considerable progress over the past few years. However, until recently, there was little work devoted to integrating single-photon detectors onto a chip. On-chip detectors carry a number of important advantages over their free-space or fibre-coupled normal-incidence counterparts. Since there is no need to come off-chip, it removes one more potentially lossy interface in the photonic circuit. Additionally, on-chip detection allows high efficiency detection over a wide range of wavelengths, polarisation-dependent detection, arbitrary placement of detectors within a planar circuit and transmission of the undetected signal. This last point, in particular, will be shown to be of considerable benefit.

In this chapter, as part of a collaboration with NIST, we demonstrate the first integration of the TES detectors described in chapter 5 with the UV-written silica-on-silicon waveguide architecture used in chapters 3 and 4. After introducing the evanescently-coupled detector concept, we briefly review the first proof-of-principle experiments that were carried out [307]. By carefully optimising the device geometry, accounting for both optical and thermal effects

we then show operation of a device with increased detection efficiency [308]. The results from this experiment confirm the absence of significant scattering losses in the system and opens the possibility for more novel devices.

Review of on-chip single-photon detectors

Although conducted independently, the work in this chapter coincided with a wider global effort by numerous quantum optics and nano-fabrication groups to build on-chip single-photon detectors. Before we describe our specific work integrating transition-edge sensors, we briefly summarise the current state of the art for waveguide integrated single photon detectors in table 6.1. The work described in this chapter resulted in the first two realisations of an on-chip TES detector and are highlighted in the first two rows of table 6.1. The remaining devices shown in the table are all based on superconducting nanowire single photon detectors (SNSPDs) integrated onto various waveguides architectures.

The table includes both the 'on-chip' and 'system' detection efficiencies of each waveguide-based detector on the list. The 'on-chip' detection efficiency represents the probability of a photon already coupled into the correct waveguide mode causing the detector to register a signal. The 'system' efficiency, meanwhile, represents the probability that a photon coupled into the input optical fibre outside of the cryostat will cause the detector to fire and thus also includes the effects of coupling and propagation loss of the waveguide. In most respects, the 'on-chip' efficiency is the most important figure since it is the stated aim of this research to use photon sources and circuits which are already integrated, thus removing the need to lossy fibre coupling. However, the 'system' efficiencies give a useful indication of the current practicalities of using such devices.

Detector type	Measured detector parameters					
	Detector material	Waveguide architecture	$\eta_{\text{on-chip}}$	η_{system}	PNR?	Notes
TESs						
T. Gerrits et al. (2011)[307]	W	SoS	7.2%	3.4%	5	
B. Calkins et al. (2013)[308]	W	SoS	88%	12.7%	5	three detectors multiplexed for increased efficiency
SNSPDs						
W.H.P. Pernice et al. (2011)[142]	NbN	SoI	91%	2.2%	No	16% uncertainty in the on-chip efficiency
J. Sprengers et al. (2011)[309]	NbN	GaAs	19.7%	3.4%	No	
C. Schuck et al. (2013)[310]	NbTiN	SiN	52.5%, 80.1% (1550 nm 768 nm)	4.3%, 3.9% (1550 nm, 768 nm)	No	dual-band operation
C. Schuck et al. (2013)[311]	NbTiN	SiN	67%	-	No	dark counts < 100 mHz
C. Schuck et al. (2013)[274]	NbN	SoI	88%	-	No	fabricated as part of a 20×12 grid of on-chip SNSPDs
G. Reithmaier et al. (2013)[312]	NbN	GaAs	0.1%	†	No	integrated with an on-chip quantum dot source
F. Najafi et al. (2014)[313]	NbN	SoI	45%	11%	No	fabricated using flip-chip technology to increase device yield
D. Sahin et al. (2013)[314]	NbN	GaAs	24%	4%	4	a four-wire SNSPD with single readout

Table 6.1: Comparison of waveguide-integrated single photon detectors A list of the state-of-the-art for various waveguide-integrated single-photon detectors that have been reported. The two on-chip TES detectors listed at the top of this table summarise the results of the work carried out in this chapter. The two efficiency metrics, $\eta_{\text{on-chip}}$ and η_{system}, describe the probability of the detector firing given a photon inside the waveguide underneath the detection region and given a photon in the input fibre respectively. NOTES: † - in this device the single photon source (a quantum dot) was integrated into the waveguide itself meaning there is no equivalent 'system' efficiency to report.

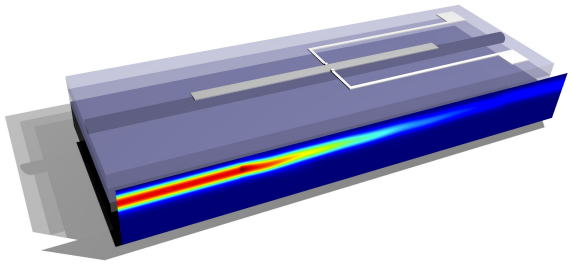


Figure 6.1: Evanescently coupled photon counting detector concept. The tungsten detection layer is deposited on top of an un-clad UV-written silica-on-silicon waveguide structure directly above a guided optical mode. The evanescent field of this mode couples to the tungsten which absorbs some optical power as the field propagates underneath it. An example internal mode profile can be seen at the front of the device schematic showing the gradual depletion of optical intensity as the field is absorbed by the detector.

6.1 Evanescently coupled detectors

Most available photon detectors, whether they are semiconductor-based avalanche photodiodes or superconducting detectors like TESs or SNSPDs, work using normal-incidence absorption. Since the thickness of the detector substrates are intrinsically limited by the detection mechanisms, care must be taken to reduce reflections from the metallic surface and design cavities to effectively confine the light within the detector in order to achieve high detection efficiencies. In the previous chapter, we have seen that this limits the available optical-bandwidth of the detectors, with different dielectric structures required to achieve high efficiencies for both visible and near-IR wavelengths. In addition, normal-incidence detection schemes restrict the placement of detectors to the end-facets of a device, reducing the flexibility to infer information about the quantum state from within the circuit.

In contrast, evanescent coupling, where detectors are placed in the evanescent field of the waveguided mode, avoids these problems completely. This concept is shown in figure 6.1, where the photon absorber, a regular tungsten TES, is deposited above a single waveguided

mode. As the mode propagates through the detector region it is coupled continuously into the detector via absorption. By carefully designing the detector geometry it is possible to achieve a high detection efficiency over the entire optical bandwidth of the waveguide without the need for any cavity structures or anti-reflection elements. The detectors can be patterned anywhere in a complex photonic circuit to target any desired guided mode.

A novel and very useful feature of this evanescent design is the transmissive failure mode of the device. Any optical signal not absorbed by the detector simply remains in the guided mode - available for subsequent detection or processing. So long as the metallic detector layer does not significantly alter the effective index or spatial mode profile of the guided mode, there will minimal scattering out of the guide and multiple detectors can thus be multiplexed one after another in a simple scheme to increase detection efficiency. We demonstrate this idea in the second-generation detector design and show there is no unexplained scattering loss. Ensuring that the optical modes of the rest of the waveguide circuit are impedance matched to the detector regions in this way was a key consideration in the design of the high-efficiency 'double-tadpole' geometry in section 6.2.

6.1.1 Proof-of-principle detector

In the first instance, it was decided to test the evanescently coupled design using a 'standard' TES design such as the ones used in chapter 5. The $25\,\mu\text{m}\times25\,\mu\text{m}\times40\,\text{nm}$ square tungsten TES detectors have been used and developed over many years at NIST and are known to have good noise properties, energy resolution and convenient transition temperatures[284,315]. Thus it was thought that using such a detector provided the best chance for easily demonstrating a successful prototype.

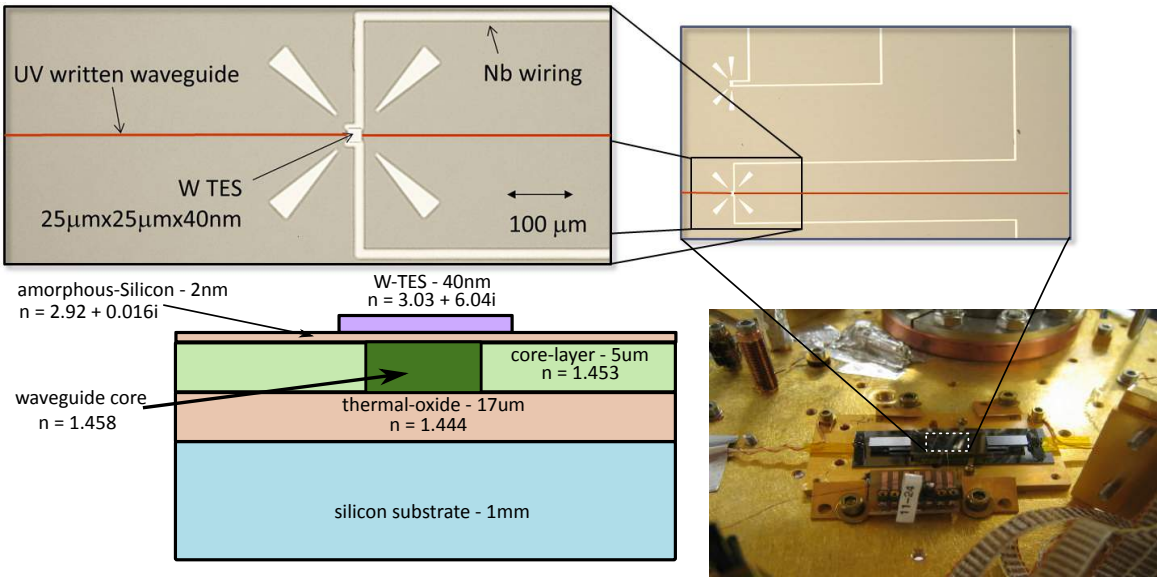


Figure 6.2: The first proof-of-principle test device. Top - a microscope image showing the fabricated device for the first tests. The top right image shows both the main detector over the waveguide (falsely coloured orange) and the reference side-device to test for scattered light absorption. The patterned arrows serve as alignment marks and guides. Bottom - a schematic of the UV-written silica-on-silicon device with TES detector patterned on top. There is no cladding layer deposited above the waveguide core which enhances the evanescent coupling into the tungsten. The photo shows the waveguide device mounted on the cold plate of a dilution refrigerator.

The silica-on-silicon waveguide structure used for these detectors is similar to those used in the previous chapters. Importantly, no overclad layer was deposited for these samples, which is required in order to achieve sufficient overlap of the evanescent field with the tungsten detector layer (see chip schematic in figure 6.2). However, as the waveguide core is now exposed to the air, the chip surface must be kept extremely clean to prevent photons being scattered out of the guide by particles on the surface. In reality, this manifests itself as an extra waveguide propagation loss of about 1dB /cm. In addition, the low index of the air causes the guided mode to take on a more elliptical shape which appears pushed down from the chip surface. Finally, the UV-written silica-on-silicon architecture provides a surface roughness of less than 1nm across the width of the detector region. This is of crucial importance since surface roughness increases the thermal stresses on the deposited TES layer

which can lead to a suppressed transition temperature[316]. In order to achieve the desired transition temperature in the tungsten TES layer, a thin, 2 nm layer of amorphous silicon was deposited between the silica waveguide core and the detector. This helped alleviate residual thermal stresses whilst being thin enough to have essentially no optical effect on the device.

This first test device is shown in figure 6.2 and was fabricated by the team at NIST. The tungsten TES was DC sputtered onto the amorphous silicon layer directly above the waveguide core using alignment markers laid down during the waveguide writing phase. The microscope image in figure 6.2 shows the TES was aligned to better than 10 nm of the centre of the waveguide core. Niobium was used to wire the TES to on-chip bond pads, similar to the method used in the standard TES construction detailed in chapter 5. The chip was then pigtailed to SMF-28 single mode fiber on either end using v-groove arrays and UV-cure epoxy similar to that used earlier in section 4.8.

6.1.2 Numerical modelling

Numerical mode solver - Fimmwave

Before the device was fabricated, the optical absorption of the detector was modelled using Fimmwave - a commercial fully-vectorial numerical mode-solver[317]. Fimmwave uses the Film-Mode-Matching (FMM) method to quickly and accurately model this evanescently coupled detector design[318]. The FMM method works particularly well for rectangular waveguide structures and works by first slicing the entire structure into vertical films of uniform transverse profile. A number of 1D TE and TM modes are then calculated for each of these films. Because these 1D modes can be computed analytically, the size and

index-contrast of the layers in each film have almost no influence on the computation time and do not need to be approximated. Across all of the films, modes with the same z-component of the modal wavevector, k_z are collected and their field distributions matched at each slice interface by adjusting the modal amplitudes in each film. Only for certain wavevectors, k_z, can a set of non-zero film mode amplitudes with matching distributions at the slice interfaces be found. Each set constitutes a sought-after guided eigenmode. The remaining set of continuous, unbound (i.e. radiative) modes are discretised via the finite computational boundary window. By monitoring the change of propagation constant of the lowest-order guided mode as we increased the number of 1D modes we found that 330 of the latter provided an optimal trade off between speed and accuracy. With these settings, the calculation of 50 modes took around 15 minutes per structure on a quad-core 3GHz processor using 12GB of RAM.

Eigenmode-expansion-method — Fimmprop

Whilst enabling the calculation of the solutions of the wave-equation in order to calculate the modes present at a single point in the waveguide, the FMM method on its own does not produce solutions for modes propagating along the guide. Instead, we use the Eigenmode-Expansion-Method (EME)[319,320] as implemented by the add-on program, FimmProp, to calculate the overall scattering matrix of the device to investigate how much power would be coupled into the detector.

The EME technique uses the set of 2D eigenmodes (both guided and radiation modes) found using the FMM method as a complete basis set in which all solutions to Maxwell's equations in that waveguide can be expressed. During straight-waveguide propagation in

a uniform guide, each mode simply acquires a linear (possibly complex) phase. At a join between different waveguide structures an overlap integral is computed between the set of modes on each side of the join to obtain a scattering matrix. The scattering matrix for the entire device is found by multiplying the matrices of each joint and is exact in the limit of an infinite set of 2D basis modes. For all of the calculations in this chapter a set of 50 2D modes were used. This method is also inherently bi-directional, so we obtain all of the effects of reflections occurring between different waveguide sections automatically [1]. To estimate the fraction of power absorbed by the detector we compare the amplitude of the initial single mode to that remaining at the end of the EME propagation. We account for any power scattered into radiation or backward-travelling modes by adding up the power present in them at the end of the calculation.

Model results

The results from this modelling are shown in figure 6.3. One important finding is the confirmation that the addition of the tungsten detector layer causes only small distortion to the optical mode. We also observe a strong preferential absorption for TM-polarised modes, likely explained by the presence of surface plasmon polaritons (SPPs) propagating along the metal-dielectric interface. Such SPPs can only be excited by TM-modes and are characterised by tight spatial mode confinement and high local field intensity. Thus, they interact much more strongly with the metallic layer and hence exhibit higher propagation loss. The absorption peak feature which is seen in the TM-mode is a known effect of SPPs [321,322]. Since the TE-mode is not able to excite any SPPs, this peak-like absorption

[1] For the first test device we considered 3 waveguide sections — propagation region, detector region, propagation region. For the multiplexed high-efficiency devices presented later a total of 13 waveguide sections were used.

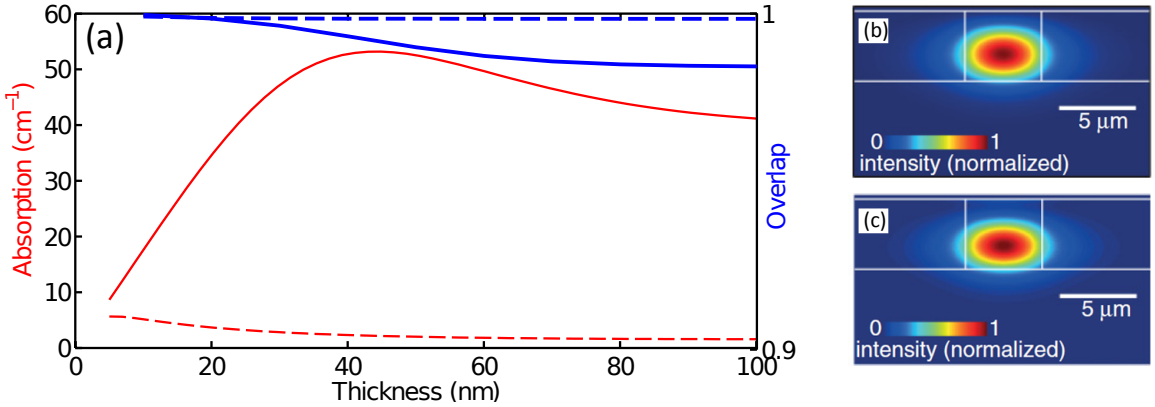


Figure 6.3: Device modelling. (a) Plot showing how the overlap between the mode supported by the region with and without the detector (blue) and the absorption coefficient (α/cm^{-1}, red) of both the TE (dashed) and TM (solid) mode varies as the thickness of the layer is altered for the 25 μm-wide detector. The maximum absorption coefficient for the TM mode occurs around 45 nm with both the shape and preferential absorption over the TE-like mode clear evidence of a surface plasmon polariton effect. (b) The modelled normalised intensity distribution of the TM mode in the waveguide region without the detector. (c) The modelled normalised intensity distribution of the TM mode in the detector region with a 25 μm $\times$ 40 nm tungsten square deposited directly above the core. Only a small amount of distortion in the vertical direction can be seen, along with a slight concentration of intensity near the detector layer. The calculated scattering matrix between these two regions suggests we should expect less than 1% excess loss due to this mode mis-match.

feature is not observed in the simulations.

For this initial test geometry we find that we should expect an absorption coefficient of 54.6cm^{-1} for the TM-mode and 2.3 cm^{-1} for the TE-mode. To test that our model is accurate we compare these figures to a classical loss measurement using the Bragg grating method presented earlier in section 4.8. A sample was fabricated with a 100 μm long, 40 nm thick strip of tungsten across the waveguide. The results gave measured absorption coefficients of 55.6 ± 2.4 cm^{-1} and 3 ± 2cm^{-1} for the TM- and TE-modes respectively. This compares well with our model predictions and thus gives us confidence in its validity. For the specific geometry used for the initial test devices this translates to an expected detector efficiency of 13.2% and 1.2% for the TM- and TE modes respectively.

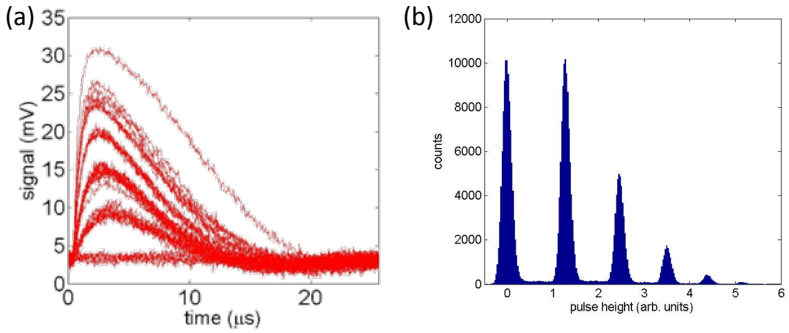


Figure 6.4: Test-device results. (a) Raw electrical output traces recorded from the on-chip TES when an attenuated pulsed laser is coupled into the waveguide. The number-resolving capability is clearly visible here. (b) The corresponding pulse-height histogram derived from the raw electrical output traces. The histogram was created using the matched-template method used earlier in section BLAH. Fitting a Poisson distribution to this histogram allows us to derive a mean detected photon number of $\langle n \rangle = 0.986$ for the optimal TM-polarisation.

6.1.3 First results

The on-chip detector was characterised by the team at NIST using a similar scheme as that described in the previous chapter in section 5.4. The device was cooled to 12 mK in a dilution refrigerator and a pulsed 1550 nm laser diode of 10 ns duration at $\sim$ 35 kHz attenuated to a mean photon number per pulse of $\langle n \rangle = 28.5 \pm 0.8$ was coupled into the fibre pigtailed device. The raw output traces from this device are shown in figure 6.4. The input polarisation was adjusted to launch a TM-like input state by optimising the detector response, whilst a TE-like input state was prepared by minimising the detector response. The matched-template method was again used to convert these raw traces into a histogram of pulse heights which were then fitted with a Poisson distributed convolved with Gaussian peaks to extract the mean photon number recorded by each detector. This resulted in $\langle n \rangle_{TM} = 0.986$ and $\langle n \rangle_{TE} = 0.086$. This leads to an overall system detection efficiency of 3.5% for the TM-mode and 0.3% for the TE-mode. These figures include the losses arising from the fiber coupling and waveguide propagation loss and so are not a good estimate of

the actual quantum efficiency of the detectors which is the probability of a photon already in the waveguide mode being detected.

To recover the raw quantum efficiency of the detector, ξ_{TES} we first make a measurement of the total device transmission,

$$\xi_{\mathrm{total}} = (1-\eta_A)\xi_{\mathrm{TES}}(1-\eta_B) \tag{6.1}$$

where $\eta_{A/B}$ are the fiber coupling efficiencies on each side and include the propagation losses in the waveguide up to the detector region. We also measure the system detection efficiency as described above for light coupled from either side,

$$\xi^{(A)}_{\mathrm{device}} = (1-\eta_A)\xi_{\mathrm{TES}}, \tag{6.2}$$

$$\xi^{(B)}_{\mathrm{device}} = (1-\eta_B)\xi_{\mathrm{TES}}, \tag{6.3}$$

from which we can deduce the actual detector quantum efficiency, $\xi_{\mathrm{TES}} = \xi_{\mathrm{total}}/\xi^{(A)}_{\mathrm{device}}\xi^{(B)}_{\mathrm{device}}$. This method yields absolute detection efficiencies of $7.2 \pm 0.5\%$ and $0.65 \pm 0.05\%$ for the TM and TE modes, respectively. Comparing this with our simulations above we find a discrepancy that shows lower than expected detector performance. One possible explanation is that the tungsten was found to have not adhered properly to the waveguide which could lead to photon absorption events not properly thermalising with the tungsten electronic system and instead being lost to the substrate. However, it is important to confirm that there is not some additional unexplained loss mechanism occurring in these devices since this would prevent the efficient multiplexing of these detectors. The second-generation design discussed below in section 6.2 was designed in order to directly measure the effect of any

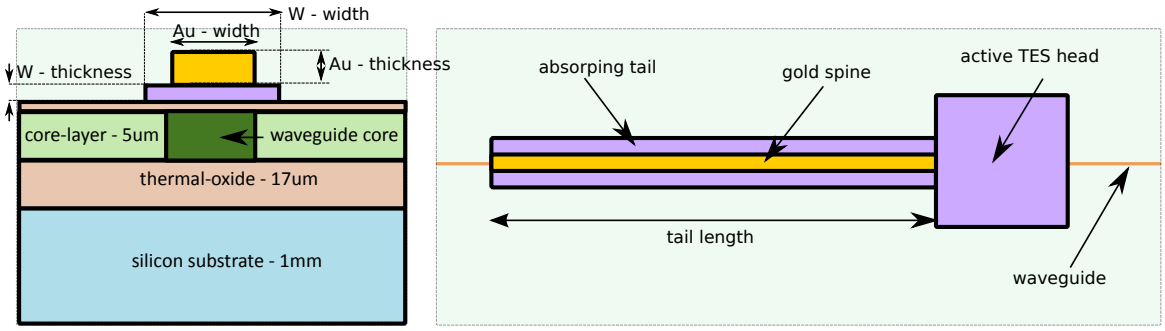


Figure 6.5: The 'tadpole' device design. In this geometry, the primary absorber is now separate to the active TES element. This allows us to use a long and thin absorbing 'tail' which otherwise would not function correctly as a detector. This tail has an added gold spine which acts to drive the tungsten below it into a normal conducting state and to provide greater thermal conductivity along the length of the absorber. The end-view schematic is similar to that used in the previous design shown in figure 6.2. An optimisation scan was performed over the length scales indicated on the figure.

additional loss mechanisms.

6.2 Increasing efficiency

These first detector results were promising and clearly demonstrated the proof-of-principle of evanescent absorption for use in on-chip superconducting detectors. Around the same time as this work, other groups demonstrated similar results using superconducting nanowire single-photon detectors (SNSPDs) integrated with Si[142] and GaAs[309] waveguide structures. Whilst these other works showed promise of high efficiency on-chip detection, the unique number-resolving capability of the TES makes it more difficult to optimise the quantum efficiency without sacrificing energy resolution or timing performance. The next section describes the effort involved in carefully balancing the optical and thermal properties of the device in order to optimise the detector performance.

6.2.1 The 'Tadpole' design

There is a design trade-off between optimising detector efficiency and energy resolution. The first-order optimisation problem occurs between the evanescent interaction length and the detector's thermal heat capacity. For more efficient absorption we desire a longer detector. However, good energy resolution is provided by a small heat capacity and thus a smaller volume of metal. The intensity of the guided mode is mostly concentrated in the centre region of the core a few microns across. It makes sense to trade this width for extra length of the device since the fractional absorption is exponential with length $\propto e^{-\alpha z}$. This is also the case for the tungsten thickness. Although the modal absorption peaks at a thickness of $\sim 40\,\mathrm{nm}$, decreasing the thickness by half only reduces the modal absorption by a factor of ~ 1.8 which means we can again trade some thickness for extra length. By this first-order optimisation we can reduce the tungsten thickness to 20 nm and the width to $4\,\mu\mathrm{m}$ which roughly correspond to the minimum fabrication capabilities. This would enable us to use a length of greater than $300\,\mu\mathrm{m}$ for the detector and potentially achieve a fractional absorption of 55%. In addition, because the modal absorption factor, α, is not increased over the previous device design, the mode mis-match between the detector and propagation regions will remain small.

However, such a device geometry poses a number of technical challenges. Maintaining good thermal equilibrium over such a long, thin length is difficult, and eventually the thermal coupling between the TES and the substrate will become larger than the thermal coupling along the TES itself[323]. This can lead to a phase separation of the metal into normal and superconducting sections[276,324] which can lead to unstable bias conditions and the failure of the device. A large amount of work and trial-and-error testing of TES devices has

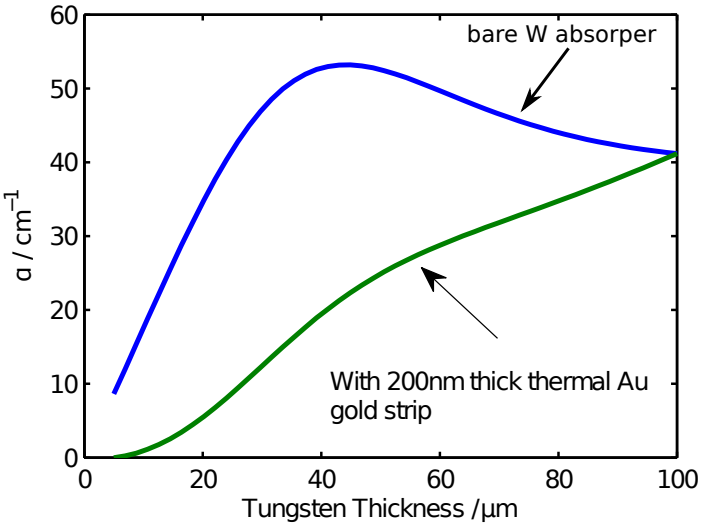


Figure 6.6: Effect of the thermal gold-spine. The low real-part of the refractive index of gold causes a reduction in the absorption coefficient compared with the bare tungsten used in the first proof-of-principle device.

gone into the standard $10\,\mu\text{m} \times 10\,\mu\text{m} \times 40\,\text{nm}$-sized normal-incidence devices to make reliable detectors with good quantum efficiency, energy resolution, convenient superconducting transition temperature and timing properties.

Instead the 'tadpole' design was developed to obtain a good balance of optical absorption and thermal properties. In this design, shown in figure 6.5, we separate the absorber from the TES, instead using a long, normal-metal 'tail' as an absorbing fin connected to a standard-design active TES 'head'. Because of the earlier success of using tungsten as the absorber it was decided to continue using this metal for both the absorbing tail fin and the TES head. However, since tungsten is a superconductor at the operating temperatures we use, it has very poor thermal conductivity. A normal-metal spine must be added to the absorbing tail in order to achieve sufficient thermal conductivity. Gold was chosen both for its excellent thermal conductivity properties and its previous use in TES devices to improve their thermal recovery time[292]. With this new proposed 'tadpole' geometry we now must find the optimal device geometry such that the detector achieves good quantum efficiency whilst retaining

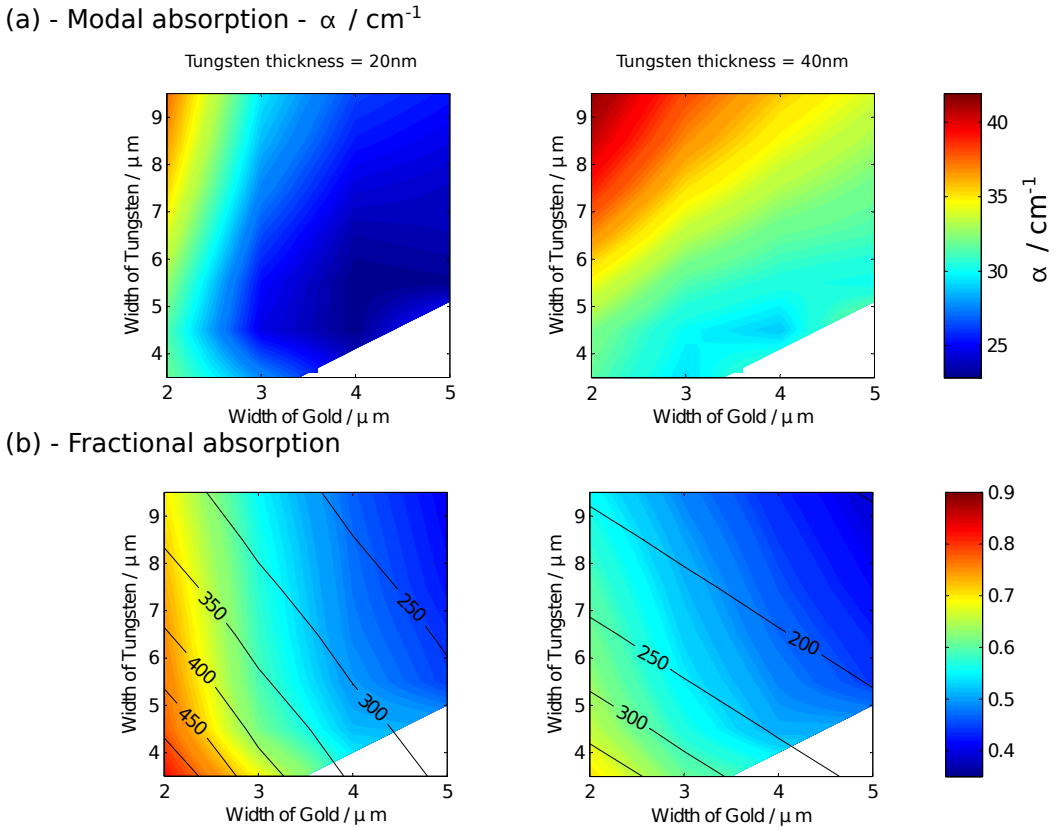


Figure 6.7: Optical modelling of the tadpole design. (a) The modal absorption coefficient, α, is calculated for various detector geometries. For a 20 nm thick tungsten layer, the effect of increasing gold width is a dramatic fall off in achievable absorption. However, once the tungsten thickness reaches 40 nm, the effect of the gold is less pronounced. (b) The resulting fractional absorption is calculated by first determining the maximum device length allowed for a given geometry and using the modal absorption coefficient calculated in (a) to compute $\xi = (1 - e^{-\alpha L})$. Black contour lines show the resulting device length. From these plots, it is clear that the bottom-left corners (i.e. narrow tungsten and gold), are the optimal places to operate for increased absorption.

acceptable energy resolution.

6.2.2 Device optimisation

Geometrical optimisation

Fimmwave and Fimmprop were again used to perform numerical modelling of various potential device geometries. The waveguide schematic in figure 6.5 details the four parameters to be explored. The low real-part of the refractive index of gold ($n = 0.38$) acts to push the

optical mode away from the absorbing surface leading to a decreased modal absorption parameter as seen in figure 6.6. Hence, it is important to find the optimal geometry to minimise the detrimental effect of this gold. The parameter scan proceeds by varying widths and thicknesses of both the tungsten absorbing tail and the thermal gold spine over the ranges: $W_{\text{thick}} = \{20, 30, 40\}\,\text{nm}$, $Au_{\text{thick}} = \{50, 100, 150, 200\}\,\text{nm}$, $W_{\text{width}} = \{3, 4, 5, 6, 7, 8, 9\}\,\mu\text{m}$ and $Au_{\text{width}} = \{2, 3, 4, 5\}\,\mu\text{m}$. In order to preserve the desired energy resolution, the total heat capacity of the device is fixed. Therefore, for a given cross-sectional area of tungsten (A_{W}), and gold (A_{Au}), we can infer the maximum device length,

$$L_{\text{device}} = \frac{C_{\text{tot}}}{A_{\text{W}} + A_{\text{Au}}/1.96}, \tag{6.4}$$

where C_{tot} is the total allowed heat capacity normalised to a volume of tungsten and the factor of 1.96 takes into account the reduced heat capacity of gold with respect to tungsten. In effect, the total length of the device is still mostly constrained by the volume of the tungsten. Figure 6.7 summarises how the modal absorption parameter, α/cm^{-1}, as well the fractional absorption change as the detector geometry is varied. As expected from the results of the previous device, as we increase the width and thickness of the tungsten, the modal absorption value increases. We can also see the detrimental effect the gold spine has on the absorption of mode into the metallic layer. When we use a 20 nm thick tungsten layer, the low refractive index acts to push the mode away from the surface and results in a lower modal absorption as the width of the gold layer increases. Although not shown on the plots, it was found that varying the thickness of the gold spine from 40 nm to 200 nm did not significantly affect the modal absorption values. Since we require at least an 80 nm

layer to achieve significant thermal conductivity we fixed the thickness of the gold layer at this value for the remainder of the simulations.

Increasing the width of the gold spine also contributed to a sharp fall in the mode-matching between the propagation and detector regions of the chip. However, it is seen that once the thickness of the tungsten is increased to 40 nm, the negative effects of the gold are mostly masked. In particular, the drop off of α with increased gold width is dramatically reduced. Translating these modal absorption values into a fractional absorption as shown in the lower half of figure 6.7 shows that once again it is beneficial to trade the width and thickness of tungsten for added length. This would seem to suggest that the detector should be only 20 nm thick and be as thin as possible. If we could build a device 4 μm wide and 500 μm long we could expect close to 90% absorption. However, we must also consider the effect of additional scattering into radiation modes caused by the greater mode-mismatch present with the 20 nm devices. Before proceeding to choose a device geometry based on these results we should also consider the thermal properties of such a device.

Thermal modelling

The addition of the gold spine in the 'tadpole'-design allowed us to consider fabricating longer absorbing devices since we no longer had to keep the entire device superconducting. However, when a photon is absorbed along the tail of this device we need to make sure that most of that energy makes it to the active TES at the head of the tail before it is lost through some other thermal path. Energy lost along the tail, rather than being conducted to the detector, degrades the achievable energy resolution. Since the absorbing fin is a normal-metal the primary heat conduction process is electron diffusion and can

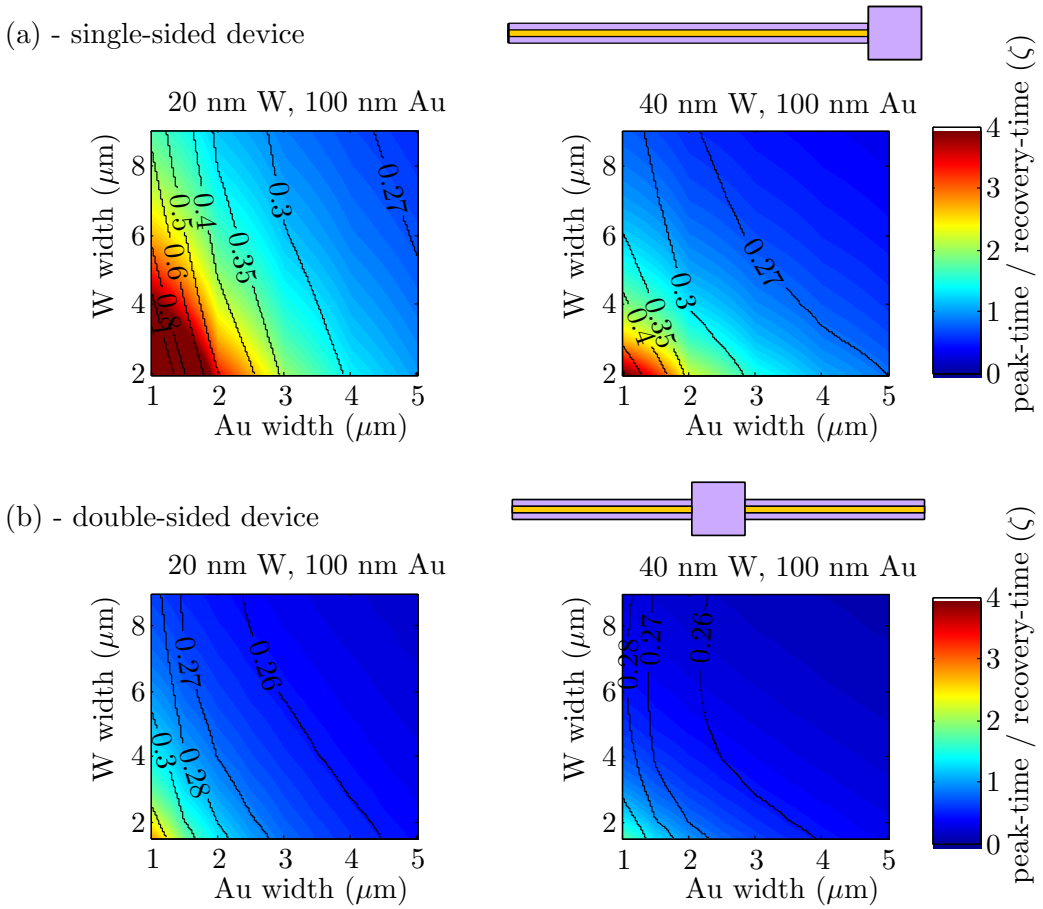


Figure 6.8: Summary of thermal modelling. (a) Single-sided device. The length of the entire absorbing tail is fabricated as one piece. The coloured shading shows the predicted peak-time / recovery-time ratio which should ideally be less than one. The black contours show the predicted energy resolution of the device in electron-volts - one photon at 1550 nm is 0.8 eV. At the bottom left corners, where the optical efficiency is highest (known from figure 6.7), is where thermal performance is worst. (b) Double-side device. The absorber tail length is split into two equal sections before and after the active TES head. The thermal modelling from (a) is performed again which clearly shows for a 40 nm thick tungsten layer we can achieve good optical and thermal performance in the same regions of parameter space. These thermal modelling results were performed by Brice Calkins and are used with permission.

be described by the Wiedemann-Franz law, $\kappa_{wf} \propto \sigma T_e$, where κ_{wf} is the thermal conductivity, σ is the electrical conductivity of the normal-metal tail and T_e is the electron temperature. This conduction along the length of the spine is in competition with the thermal escape-mechanism of electron-phonon coupling which has an effective conductivity described by $\kappa_{e-p} \propto \Sigma(T_e^5 - T_p^5)$, where Σ is a material parameter quantifying the strength of the electron-phonon interaction and T_p is the phonon temperature. We can immediately see that if we want to achieve $\kappa_{wf}/\kappa_{e-p} \gg 1$ it is advantageous to make T_e lower which means that a lower device operating temperature (T_c) is needed. By constructing a full one-dimensional heat equation for this device, and incorporating the temperature evolution and expected noise spectrum of the active TES, it is possible to probe the trade-off between device length and energy resolution for a given transition temperature [2]. Controlling the transition temperature of a TES requires careful optimisation of the deposition process and requires substantial research effort [316]. For this work, no effort was made to reduce the transition temperature below the standard achievable value of 70-80 mK. Thus, all thermal modelling was carried out assuming a transition temperature of 75 mK. In future work it might be possible to achieve better thermal performance by engineering the tungsten film to further suppress its T_c. A second important metric for the thermal performance of the detector is the ratio of the time taken for heat to be conducted along the tail to the TES, τ_c to the time taken for the TES to remove this heat during recovery, τ_{etf}. Ideally, this so-called 'peak-time/recovery-time' ratio, $\zeta = \frac{\tau_c}{\tau_{etf}}$, would be less than one to avoid severe distortions to the electrical pulses produced by the detector in response to photons being absorbed at different locations along the tail.

[2] All of the thermal modelling in this work was carried out by Brice Calkins from the quantum electronics group at NIST, Boulder, USA.

The results of this thermal modelling are shown in figure 6.8. Both the expected energy resolution and ζ are plotted as a function of detector geometry over the same parameter range as the optical modelling above. Unfortunately, there is almost a direct trade-off between optical absorption efficiency and thermal performance. As the detector is made longer and narrower for optical purposes, the increased length of the detector provides more time for heat to escape through electron-phonon coupling, increasing ζ. Therefore, it was decided to split the absorbing tail into two to create a double-tailed tadpole device with the active TES in the centre as shown in figure 6.8(b). With this design, we can maintain the same optical interaction length whilst halving the length of metal heat must diffuse through until it reaches the active TES.

The thermal modelling results for this double-tailed design are shown in figure 6.8(b). It is now possible to access geometries which show simultaneously high optical absorption and good thermal properties. It is clear we wish to operate in the lower-left regions of this parameter space. Therefore we choose a device with the tungsten and gold layers made as narrow as possible - a 3.5 μm-wide tungsten strip and a 2 μm-wide gold strip of 80 nm thickness. We choose to use a 40 nm thick tungsten layer for two reasons: although a 20 nm layer would allow us to make the tail longer, the 40 nm layer minimises the mode mis-match between the propagation and detector regions. Also, the thermal modelling suggests that we may not even be able to make use of the extra device length a 20 nm would give us anyway since it would lead to poor energy resolution.

Given this cross-section we must now decide on a detector length. Although, in principle, the device could be up to 400 μm-long (i.e. with $2 \times 200\,\mu$m tails) and still have acceptable energy resolution, it was found that such devices could not be easily fabricated with a

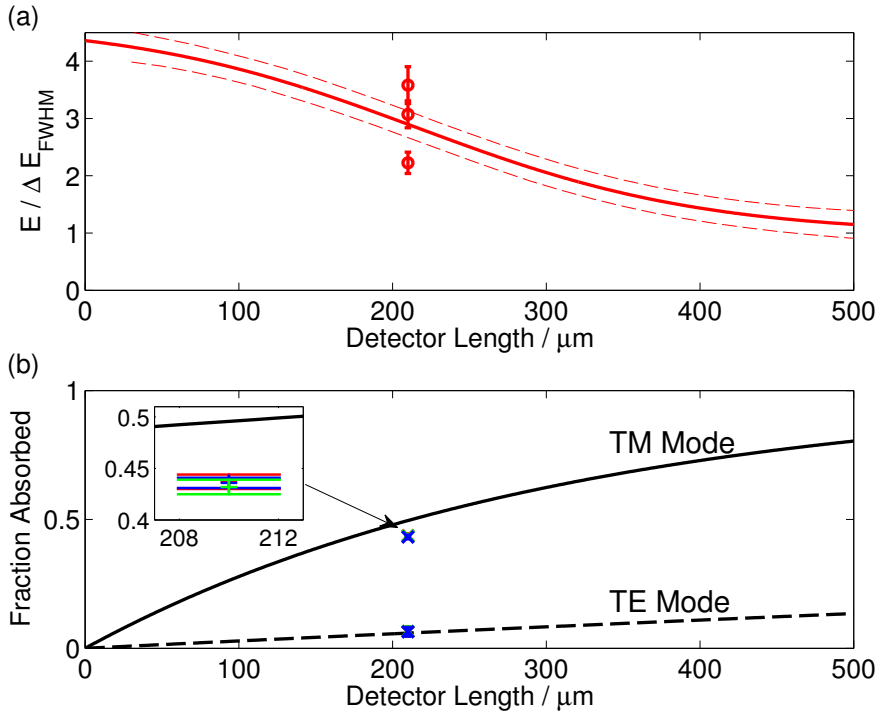


Figure 6.9: Effect of device length on energy resolution and optical absorption. These are modelling results for the final double-tailed tadpole design. (a) Thermal modelling. the red line shows the predicted energy resolution of the device as a function of detector length. The dotted lines indicate the standard error on the model. The circles represent the measured energy resolution of the three fabricated detectors discussed in section 6.3. (b) Optical modelling. The black lines show the predicted optical-absorption efficiency of the TM and TE modes as a function of whole detector length. The crosses represent the measured efficiency of the fabricated detectors as discussed in section 6.3.

low-enough T_c. Fixing all of the other parameters as defined above and considering only a single variable we can plot the dependence of energy resolution and optical absorption as a function of device length as shown in figure 6.9. We choose to fabricate $2 \times 100\,\mu$m-long tails as a good trade-off between efficiency and energy resolution whilst being confident that such a device would have an acceptable T_c. These final design parameters, along with the constants used for the thermal modelling are shown in table 6.2.

This final geometry is simulated using Fimmprop to accurately calculate the predicted detection efficiency of this new device. The full simulation is shown in figure 6.10 along with

	Design parameters	
	Tungsten	Gold spine
'Tail' Geometry		
thickness	40 nm	80 nm
width	3.5 μm	2 μm
length	$2 \times 100\,\mu$m	$2 \times 100\,\mu$m
'Head' Geometry		
thickness	40 nm	0 nm
width	10 μm	0 nm
length	10 μm	0 nm
refractive index	$n = 3.03$	$n = 0.38$
	$\alpha = 4.9 \times 10^5 \mathrm{cm}^{-1}$	$\alpha = 8.7 \times 10^5 \mathrm{cm}^{-1}$
Thermal & Electrical parameters		
electron specific heat	$140\,\mathrm{aJ}\mu\mathrm{m}^{-3}\mathrm{K}^{-1}$	$71.4\,\mathrm{aJ}\mu\mathrm{m}^{-3}\mathrm{K}^{-1}$
electron-photon coupling	$0.4\,\mathrm{nW}\mu\mathrm{m}^{-3}\mathrm{K}^{-5}$	$2.6\,\mathrm{nW}\mu\mathrm{m}^{-3}\mathrm{K}^{-5}$
electrical conductivity	$5 \times 10^6\,\Omega^{-1}\mathrm{m}^{-1}$	$2.23 \times 10^{10}\,\Omega^{-1}\mathrm{m}^{-1}$
electron mean-free path	2.8 nm	18.8 μm
thermal conductivity	$0.0122\,\mathrm{Wm}^{-1}\mathrm{K}^{-1}$	$0.232\,\mathrm{Wm}^{-1}\mathrm{K}^{-1}$

Table 6.2: Final design parameters for the high-efficiency on-chip TES detector. The thermal and electrical parameter data was used in the thermal modelling and was obtained from previous literature [308]

the 2D mode profiles in each of the detector regions: ordinary waveguide, absorbing tail, and active TES. The distortion in each mode is quite small leading to predicted overlaps between each region to be $> 98\%$. The simulations predict an absorption coefficient of 32.6 cm^{-1} (2.9 cm^{-1}) for the TM-like (TE-like) guided modes. For the whole double-fin device (total length 210 μm) including the 10 μm TES square, Fimmprop is used to propagate the modes under the detector to arrive at a predicted device efficiency of 50.3 % (6.1 %) for a TM-like (TE-like) input mode. This full set of predicted device parameters is given in table 6.3. For convenient comparison, we also list the final measured values of the fabricated device which are discussed more fully in the next section.

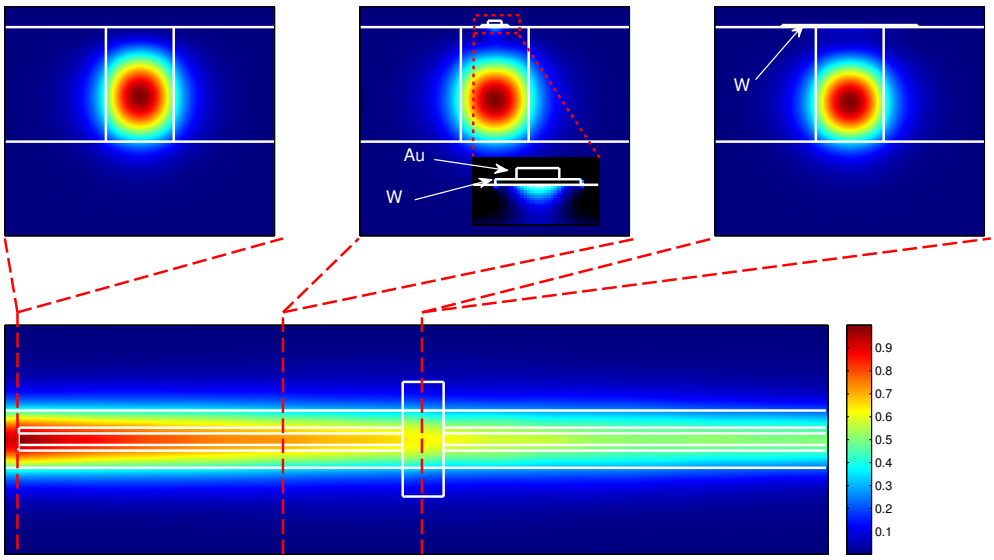


Figure 6.10: Final device design simulation. Using the final device design parameters we have plotted the simulated TM-like field as it propagates under the double tadpole detector (bottom). The colour scale shows around half of the initial optical intensity is absorbed by the detector. Two-dimensional slices at the different positions indicated by the red dotted lines show that the mode suffers almost no distortion under the detector regions leading to very small scattering losses (top). An enlarged view of the metal-silica interface shows a slight concentration of intensity underneath the absorbing tails (top middle).

6.3 High-efficiency on-chip detector results

6.3.1 Multiplexed detectors

The transmissive failure mode of these evanescently coupled waveguide detectors is a key feature which allows the possibility of device multiplexing. The multiplexed detector concept is shown in figure 6.11(a). Each TES has a finite detection efficiency, ξ, which would ideally leave the remaining fraction, $(1 - \xi)$, travelling through the waveguide available for subsequent detection. In principle, by placing multiple detectors one after another it is possible to increase the total detection efficiency arbitrarily close to 1. However, each detector also has some finite scattering probability, η_{scatt}, which gives the probability that a photon is removed from the waveguide but does not get detected. This could occur by

	Predicted values		Measured values	
	TM-mode	TE-mode	TM-mode	TE-mode
Optical properties				
Modal absorption	32.6 cm^{-1}	2.9 cm^{-1}	27.1 cm^{-1}	3.2 cm^{-1}
Fractional absorption	50.3 %	6.1 %	43.5 %	6.5 %
Scattering probability	1×10^{-2}	1×10^{-3}	$\sim 5 \times 10^{-2}$	$< 1 \times 10^{-3}$†
Thermal Properties				
Energy resolution	0.27 eV		0.32 eV	
TES Recovery time	1.6 μs		1.9 μs	
Peak-time / Recovery-time, ζ	1.2		0.9	
Transition temperature, T_c	75 mK		84 mK	

Table 6.3: Predicted properties for the high-efficiency on-chip TES detector. The results of the optical and thermal modelling performed in the previous section for the final device design as specified in table 6.2 are shown for both the TM-like and TE-like modes supported by the waveguide. For comparison we also list the measured parameters as discussed in the results section 6.3. NOTES: † - the results shown in table 6.5 would indicate that there can be no additional scattering of the TE-like mode since the classically obtained losses at each detector are lower than the measured detection efficiency. However, this difference is within the measurement error and so we simply bound the scattering probability to be less than our measurement uncertainty.

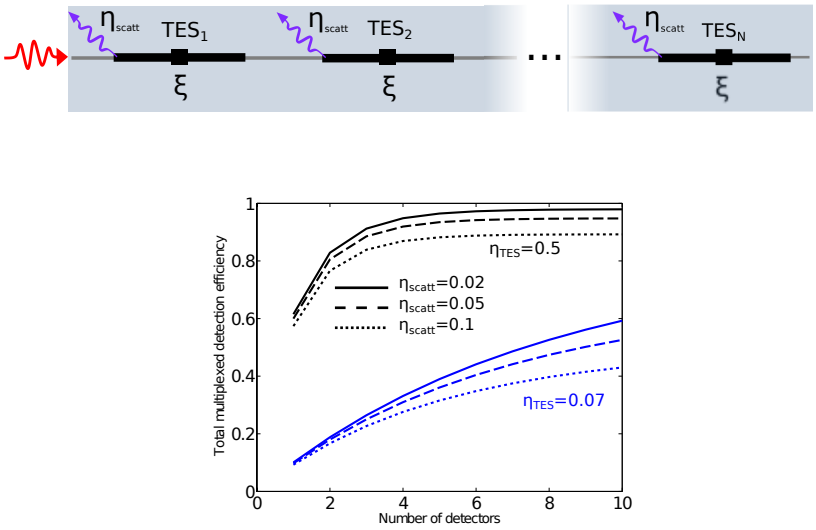


Figure 6.11: Multiplexed detector concept. Each TES has a finite detection efficiency, ξ, and a scattering probability, η_{scatt}. After propagation under each detector region there is an undetected fraction, $(1 - \eta_{\text{scatt}})(1 - \xi)$, transmitted through the guide. By placing another detector along the guide we can increase the overall detection efficiency. (b) Plot showing the dependence on overall detection efficiency as the number of multiplexed detectors increases for a variety of ξ and η_{scatt} values.

scattering from the detector region interface, by waveguide propagation loss or by other non-detector absorption mechanisms. Including this scattering loss we can write the total effective detection efficiency of a device consisting of N multiplexed detectors as,

$$\xi_{\mathrm{TES}}^{(\mathrm{total})} = \sum_{n=1}^{N} (1-\eta_{\mathrm{scatt}})^{n} (1-\xi_{\mathrm{TES}})^{n-1} \xi_{\mathrm{TES}} \tag{6.5}$$

which asymptotes to a total detection efficiency as $N \to \infty$ of,

$$\xi_{\mathrm{TES}}^{(\mathrm{total})} \underset{N\to\infty}{\longrightarrow} \left(1 + \frac{\eta_{\mathrm{scatt}}}{\xi_{\mathrm{TES}}(1-\eta_{\mathrm{scatt}})}\right)^{-1}. \tag{6.6}$$

The result of this finite scattering loss is summarised in figure 6.11(b) for different detector efficiencies. The optical modelling suggests we should achieve $\xi_{\mathrm{TES}} \sim 50\%$ and $\eta_{\mathrm{scatt}} < 0.02$ which means we could achieve 96% total device efficiency. In practice, multiplexing many detectors is quite challenging due to the significant overhead in readout electronics required. However, even with three detectors we should be able to achieve nearly 85% total device efficiency. It is clear from figure 6.11(b) that it is important to achieve low scattering loss as this fundamentally limits the achievable detection efficiency. In this experiment we wanted to measure three things: that the double-sided tadpole device design works with the efficiency expected, that a detector multiplexing strategy can be used to increase the overall device efficiency and to try and measure the scattering loss, η_{scatt}. The schematic of the fabricated device is shown in figure 6.12 along with several photos of the waveguide chip and detector in figure 6.13. The device consisted of three multiplexed double-sided tadpole TES detectors fabricated with the geometry as optimised in section 6.2.2. Each detector was connected to its own bias readout electronics and separate SQUID channel. The chip was fibre pigtailed

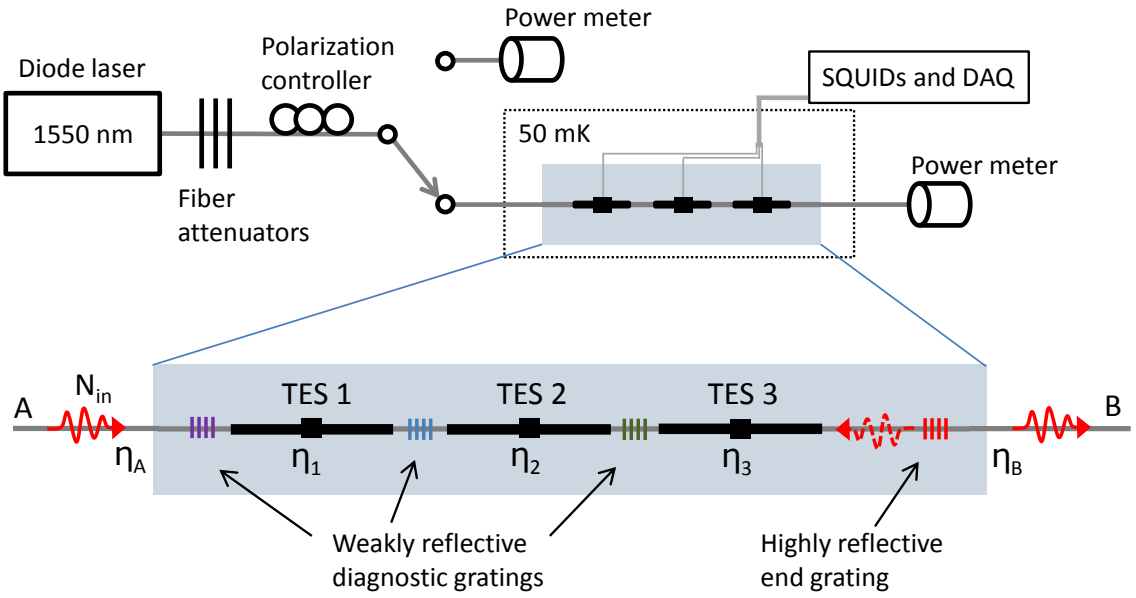


Figure 6.12: Experimental setup. Top - each TES requires its own SQUID and readout electronics. A pulsed 1550 nm diode laser is attenuated to achieve a given average photon number per pulse using a calibrated power meter. Bottom - detailed view of the on-chip TES showing the different launch directions A and B, the three detector efficiencies, η_1, η_2, η_3 and the fiber-coupling efficiencies, η_A and η_B. The weak reflectivity integrated Bragg gratings are used to make a room-temperature characterisation of device performance whilst the high-reflectivity grating at the end of the device provides a double-pass of the detectors for an appropriately tuned source launched into port A.

from both ends to allow a determination of the detector efficiencies separately from the fiber coupling efficiencies as before. In addition, we fabricated a range of weakly reflective out-of-band diagnostic Bragg gratings along the waveguide and between each detector. These gratings enabled us to make an independent measurement of the waveguide propagation loss and the amount of light absorbed by each detector region in a manner similar to the determination of the chip-to-chip gluing loss discussed in section 4.8. Because this grating loss measurement does not distinguish between the photons absorbed by the detector and the photons lost through scattering or other absorption, a comparison between the measured detection efficiency and this loss measurement provides a way of bounding the scattering loss. Finally, an in-band highly reflective Bragg grating was fabricated at the end of the device to effectively double the number of multiplexed detectors to six. From the graph

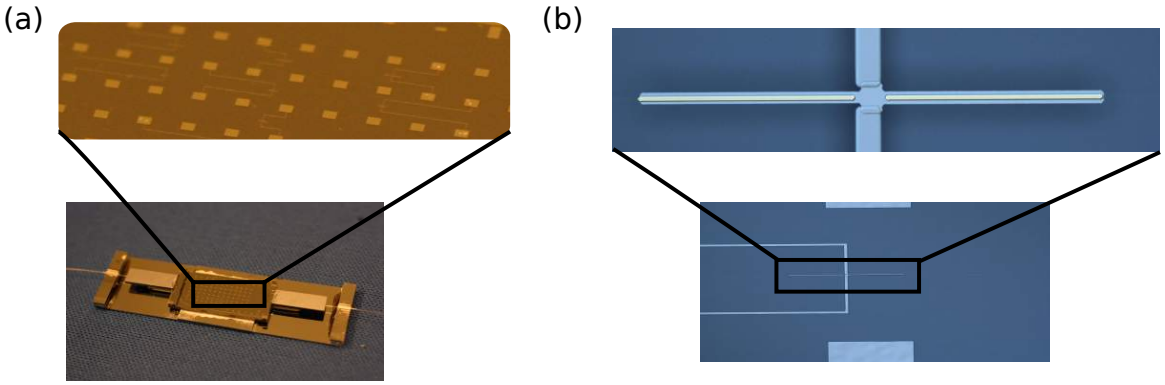


Figure 6.13: Photos of fabricated device. (a) The waveguide chip is pigtailed using two PM v-groove arrays and UV-cure epoxy. The pigtailed chip is securely mounted to a silicon wafer to provide additional support. The zoomed-in region shows the array of bond pads plus the three TESs deposited above the central waveguide. (b) Microscope photo of the double-sided tadpole TES design. The two absorbing tails have a gold spine deposited along their lengths and the central, active, TES is biased using niobium wiring which can be seen coming in from the side.

in figure 6.11(b), it is clear that six detectors essentially saturates the overall detection efficiency.

6.3.2 Results

Measurement of the total device and individual TES detection efficiency was done following a similar scheme as described earlier. For this experiment, we measured the photon response for three in-line detectors simultaneously. A 1550 nm pulsed diode laser produced nanosecond pulses at a repetition rate of 100 kHz, with known mean photon number $\bar{N}_{\text{in}}$, and was coupled to the structure and detector array from the two facets of the device (A and B).

When measuring the individual detector response for both possible propagation directions (A $\rightarrow$ B and B $\rightarrow$ A), a total of six measurements are done. The outcomes of all six

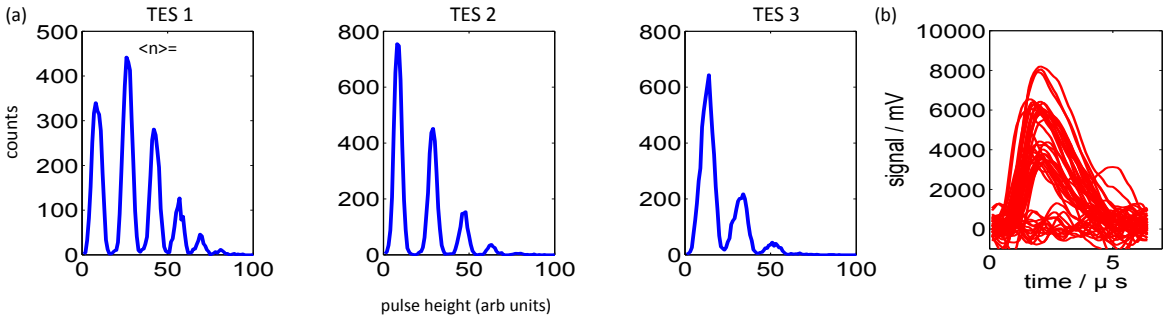


Figure 6.14: Results of high efficiency multiplexed detector. (a) Pulse-height histograms for all three TES, whilst sending triggered photon pulses along the chip from port A and port B (see figure 6.12). The mean detected photon number per pulse decreases as the light travels from TES 1 to TES 3 (left to right) due to absorption of the preceding TES. (b) Example raw pulse signal for TES 1. Individual photon numbers are still visible in the raw traces but the energy resolution and pulse jitter are visibly worse than with the first proof-of-principle device as shown in figure 6.4.

measurements are captured by six equations:

$$\bar{N}_i = \begin{cases} \eta_A \eta_i e^{-\alpha x_i} \prod_{j=1}^{i-1} (1-\eta_j) \bar{N}_{\text{in}} & \text{for} \quad A \to B \\ \eta_B \eta_i e^{-\alpha (L-x_i)} \prod_{j=i+1}^{3} (1-\eta_j) \bar{N}_{\text{in}} & \text{for} \quad B \to A \end{cases} \tag{6.7}$$

where $\bar{N}_i$ $(i = \{1,2,3\})$ are the measured mean photon numbers at TES 1, TES 2 and TES 3 for propagation directions A $\to$ B and B $\to$ A. The fibre-waveguide coupling efficiencies are $\eta_{\text{A/B}}$ and α is the measured waveguide-propagation extinction coefficient (Section 6.3.2), x_i is the position of each TES measured from edge A and L is the length of the waveguide chip.

We measured the detector array response for various input mean photon numbers, ranging from $\sim 0.1 - 3$, and observed consistent efficiency results for all input states. Figure 6.14 shows an example of the response of the TES detector array, while sending photon pulses from A to B. The light is first absorbed by TES 1 (left) and subsequently absorbed by TES 2 (centre) and TES 3 (right). The photon-number distribution of the input coherent state can be clearly observed in the pulse-height histograms, showing that the good photon-number

	TES detector efficiency			Fiber coupling efficiency	
	η_1	η_2	η_3	η_A	η_B
TM	$43.7\% \pm 0.7\%$	$43.6\% \pm 0.5\%$	$43.2\% \pm 0.7\%$	$22.1\% \pm 0.3\%$	$14.8\% \pm 0.2\%$
TE	$6.5\% \pm 1.8\%$	$6.6\% \pm 1.7\%$	$6.4\% \pm 1.7\%$	$8.1\% \pm 1.8\%$	$8.4\% \pm 2.1\%$

Table 6.4: Derived individual detector efficiencies $\eta_{1/2/3}$ and fibre-waveguide coupling efficiencies $\eta_{A/B}$.

resolution of the TES has been preserved in this design. The mean photon-number response decays as the light pulse travels from A to B, due to the absorption at each of the previous detectors.

The derived detection efficiencies for the TM input polarisation are 43.7 % ± 0.7 %, 43.6 % ± 0.5 % and 43.2 % ± 0.7 % for TES 1, TES 2 and TES 3, respectively. The combined efficiency of all three detectors is 79 % ± 2 % [3]. In order to use the high-reflectivity grating at the end of the device to make a double-pass efficiency measurement we used a tuneable CW laser source. Since the central wavelength of the Bragg grating shifts with changing temperature it is difficult to design a given operating wavelength from the outset. The central wavelength of the reflector when cold was found to be 1552.0 ± 0.1 nm. and had a measured reflectivity of ∼ 50%. When this CW laser was on resonance with this grating we observed a measured increase of the combined detector efficiency to 88 % ± 3 %. Since light could only be launched into port 'A' for this measurement, we calculated $\eta_{1/2/3}$ simply using the measured photon counts assuming that η_A and α remained unchanged. This result is consistent with the measured ∼ 50 % reflectivity of the terminal grating.

[3] The uncertainties quoted for the individual detector efficiencies include only the standard deviation of experimentally measured values, whereas the combined results also include the $1-\sigma$ uncertainty due to the calibration of the optical setup.

Fibre-waveguide coupling

These measured efficiency values are the probability that a photon inside the waveguide will be detected by any of the three detectors. This does not account for the coupling efficiency from the optical fiber to the waveguide. In this case the coupling efficiencies are 22.1 % $\pm$ 0.3 % and 14.8 % $\pm$ 0.2 % for the A and B side, respectively. This relatively low coupling efficiency is due to misalignment of the fibre v-groove pigtail with the waveguide when cooling the system to cryogenic temperatures, as well as fibre/fibre coupling losses outside of the cryostat.

The fibre v-groove is pigtailed to the waveguide chip in a similar manner to that described in the chip-to-chip gluing description in section 4.8.2. A six-axis alignment stage, together with an axial force sensor, is used to align the fibre v-groove to the waveguide whilst ensuring the faces of each remained very parallel. This ensured that an even layer of the index-matched, low-shrinkage UV-cure epoxy (Dymax OP-29) was deposited between the faces, reducing the chance of a tilt being introduced when the glue was cured. In addition, a thin strip of silica was glued to the top of each facet of the waveguide chip in order to match the height of the adjoining v-groove [4]. These bars act to strengthen the pigtail and reduce the chance that they break off during a cooldown. Using this method, were able to achieve > 40 % coupling efficiencies at the operating temperature for first proof-of-principle design. However, the high-efficiency device was pigtailed using a different epoxy with a lower viscosity. Although we achieved similarly high room-temperature coupling efficiencies as the first device, the coupling deteriorated more with this new glue at cryogenic temperatures.

Of course, we expect some reduction in coupling due to the different thermal expansion

[4] These silica bars are visible in the lower part of figure 6.13(a)

coefficients of the silicon-based waveguide chip, the silica strengthening bar, the v-groove and glue. It is likely that the difference in thermal expansion coefficient of the new glue was responsible for the poorer coupling efficiencies we measured. We expect to be able improve the low-temperature coupling by optimising the combination of glue and strengthening bar material to find a good combination of thermal expansion coefficients. We have already started this empirical search by fabricating various waveguide test samples fitted with strengthening bars made from silica doped with varying degrees of boron as well germanium-based glasses. We will test these samples using both types of epoxy and cool them down to 77 K using liquid nitrogen to explore how they survive at low temperatures. In this way we are confident of attaining the $\sim$ 40 % coupling efficiencies we have seen before and hopefully find a combination of materials which will allow us to go beyond this.

Optical loss measurements

In addition to the quantum efficiency measurements conducted using the TES detectors, room-temperature characterisation was conducted using the integrated Bragg gratings with high-intensity light. The grating-based loss measurement technique previously detailed in section 4.8 was used to provide a relative power profile along the length of the waveguide. This ratiometric approach allows the loss profile to be determined independent of any coupling losses, and since gratings are positioned between each TES element, the loss associated with each of the three detector regions may be determined. Comparing this to the actual detection efficiency provides a method of determining any excess loss present in the detector regions. Figure 6.15 displays the relative power profile of the device for the TM polarisation. The difference in relative power on either side of each of the tungsten regions,

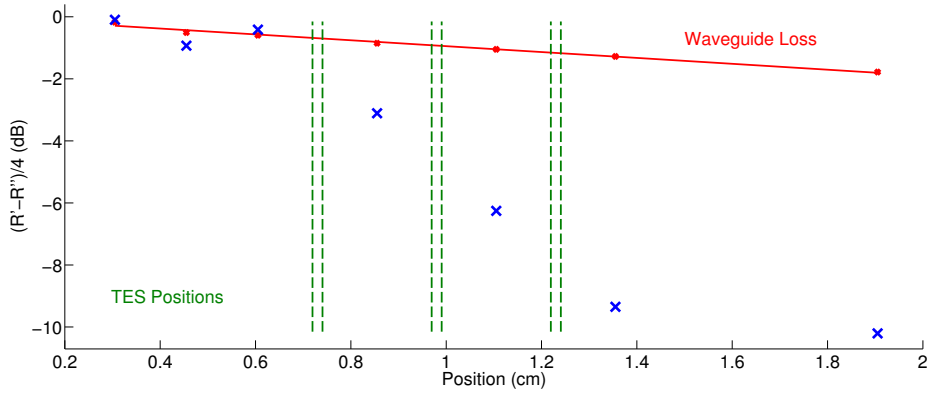


Figure 6.15: Classical optical loss measurements using integrated Bragg gratings for the TM-like mode. Each point corresponds to a different grating along the guide. Red (blue) marks indicate measurements taken before (after) TES deposition. *R'* and *R"* are the powers measured in reflection from the Bragg gratings in the forward and reverse launch directions, respectively. Both waveguide propagation loss and detector absorption can be inferred from these measurements.

after subtracting the measured waveguide loss over this distance (0.920 ± 0.020 dB/cm and 0.947 ± 0.023 dB/cm for TE and TM respectively), gives an estimate of the power coupled into the detector. The full results are listed in table 6.5 and the average absorption at each TES was found to be 46.7% for the TM-mode and 5.3% for the TE-mode.

The values are in good agreement with the detection efficiency measurements above, implying that radiative and reflective losses due to the detector are minimal. The optical modelling suggests that this scattering should be less than 2% and the results of TES 1 seem to support this. However, for TES 2 and 3, the grating-inferred losses are around 5 percentage points larger than the measured detection efficiencies. It is difficult to confirm that this is due to larger scattering losses in these detectors and not a result of different experimental conditions between the two measurements. However, with this data we can at least bound the excess scattering loss to below 5% although more carefully controlled measurements may be able to tighten this. This is already low enough to suggest that we could construct devices with overall efficiencies above 90% by multiplexing more detectors

	$\eta_{\mathrm{abs},1}$	$\eta_{\mathrm{abs},2}$	$\eta_{\mathrm{abs},3}$
TM	$43.2\% \pm 0.6\%$	$48.8\% \pm 0.6\%$	$48.2\% \pm 0.6\%$
TE	$6.3\% \pm 0.6\%$	$4.4\% \pm 0.7\%$	$5.3\% \pm 0.6\%$

Table 6.5: Classically obtained losses at each detector element $\eta_{\mathrm{abs},1/2/3}$, based on a grating based loss measurement.

and thus attain the same performance as their normal-incidence counterparts.

Finally, both the measured detection efficiency and the classically obtained losses are in excellent agreement with the optical modelling for the TE-like modes and show a slight discrepancy for the TM-like modes. That the modelling works better for the TE-like modes perhaps suggests that there are some features of the TM-induced surface plasmons which are not being captured correctly in the model. Surface roughness of the silica guide or fabrication imperfections in the metallic layer can affect the formation and propagation of these SPPs and and might thus explain the small modelling discrepancy.

Chapter 7

Conclusion

In order to realise a regime where optical quantum computers provide a performance advantage over their classical counterparts requires using more photons distributed over more modes. The recent use of integrated photonic circuits is certainly promising for building large, stable, alignment-free optical networks[13,33,35,325,326] but the issues of added photon loss and deviations from device design pose a challenge. These challenges have so far limited most demonstrations to the two-photon level and experiments still lag behind what can be achieved using bulk optics.

Some effort toward this goal of increasing the number of photons and modes which can be used on an integrated photonic chip has been made in this thesis. In the first part we were able to show the first on-chip realisation of a three-qubit protocol, whilst in the second part, the ability to perform experiments in the low-loss telecomm-band was realised with the installation and integration of highly-efficient, number-resolving detectors.

7.1 Summary

7.1.1 Silica-on-silicon waveguide circuits

Chapter 3 introduced the silica-on-silicon waveguide architecture and detailed a 6-port device consisting of 12 beam splitters, 5 phase shifters and 5 interlocking interferometers. A low-loss, high purity source of indistinguishable single photons was optimised, which enabled one of the first experiments to show genuine quantum interference between three separate input photons on an integrated photonic circuit. Moving beyond the previous two-photon, two-mode experiments, highlighted several challenges which will need to be addressed to continue scaling up these systems. In particular, it was found that loss and device characterisation play a crucial role in the operation of these circuits. In an effort toward addressing the latter, an element-wise circuit characterisation scheme was presented together with a method for simulating quantum and classical multi-photon interference visibilities. Applying a statistical test to the data allowed verification of genuine three-photon quantum interference terms in the experiment.

Realising such genuine multi-photon quantum interference within our photonic chip opened up a number of interesting experiments. First, at the end of chapter 3 we reviewed the Boson Sampling problem as a promising avenue toward a near-term demonstration of the computational power of quantum mechanics. Using the circuit and the genuine three-photon interference demonstrated at the start of chapter 3, we showed that the central premise of Boson Sampling — that the output probability distribution is given by matrix permanents — remains true in the face of real-world experimental errors. We identified the main sources of error as arising from the residual photon distinguishability and the higher-order terms

from the down-conversion source.

Second, in chapter 4 we explored an application of three-photon quantum interference to realise the first on-chip implementation of the quantum teleportation protocol - an important primitive in scalable fault-tolerant quantum information processors. The thermo-optic phase shifters were calibrated and used to prepare a variety of different input states and perform on-chip quantum-state tomography. The element-wise characterisation enabled us to model the behaviour of the quantum circuit responsible for teleportation separately from the state preparation and state tomography sections of the chip. In this way, we were able to accurately assess the effects of various sources of error in our experiment, including the effect of beam splitter deviations from their designed values. By carefully adjusting the final unitary transformations according to this characterisation, we were able to successfully teleport three linearly independent quantum states with a fidelity beyond what is classically possible. However, the error analysis conducted in this experiment highlighted a number of important issues which will need to be addressed in future, more complex devices.

First, the quantum teleportation experiment highlighted the role that component deviations from design play in reducing the achievable circuit fidelity. In this experiment we could compensate for some of this error by adjusting our post-processing analysis but it still fundamentally limited the process fidelity we could achieve. As the number of integrated elements in a circuit increases, the overall process fidelity can rapidly deteriorate even if each component has only moderate individual error. The modular chip design proposed at the end of chapter 4 contained integrated Bragg reflectors to perform in-situ device characterisation, and possesses enough reconfigurability to correct for any device deviations or unbalanced loss. The first prototype devices are being fabricated now and the initial chip-

to-chip gluing tests suggest fully reconfigurable devices over a few tens of spatial modes may be possible.

Although this remains far short of being able to realise a universal quantum computer, it will still provide a useful test-bed for quantum photonic experiments in the next few years. Boson sampling experiments with only a few tens of photons may provide the first hard evidence of the computational power of quantum mechanics. Since we will be able to correct for fabrication errors, these devices will also play a crucial role in understanding the high-fidelity operation of linear optical quantum gates, investigating whether it remains possible to achieve fault-tolerant error thresholds as the number of components increases.

Of course, there are many other potential integrated photonic platforms on which to realise complex linear optical networks. A particularly promising one is femtosecond direct-writing in silica-based glasses. Excellent control over fabricated parameters has been shown[133], as well as low-loss propagation, polarisation control[137,138,155] and novel three-dimensional waveguide arrays[131,135]. Since waveguides can be fabricated directly into a bulk glass substrate, devices have a very rapid turn-around time. In our group, we have developed a low-loss high-purity single photon source, based on four-wave mixing in such a chip[140]. However, phase shifters and other reconfigurable elements are more difficult to implement in this architecture. The poor thermal conductivity of silica makes localised heating effects difficult to achieve and there is no electro-optic effect to exploit. However, recent results exploiting the strain-optic effect[234] or evanescent coupling into an electro-optic polymers[327] offer a potential solution. Ultimately it will also be necessary to increase the on-chip component density if we are to realise truly complex optical networks. However, current high-index devices such as silicon-on-insulator, silicon-nitride, or gallium-arsenide

suffer from higher propagation loss and poor coupling into single-mode fibre. They also require expensive and timely lithographic fabrication methods. This currently restricts such devices to only a small number of photons and modes. However, since many of these substrates possess a $\chi^{(2)}$-non-linearity, it possible to fabricate integrated photon sources [145] as well as fast electro-optic phase shifters [328], increasing the appeal of these devices.

Finally, imperfect photon sources were found to be responsible for the remaining fraction of the observed reduction in fidelity in the teleportation experiment. Although residual photon distinguishability accounted for some of this error, the majority was found to be due to the higher-order terms from the down-conversion source. It is here that circuit propagation loss begins to compound our problems. The large circuit loss necessitated pumping the sources hard to give a high pair-production probability ($\sim$ 3%) so that reasonable multi-photon count rates were observed. However, this also increases the fraction of trials contaminated with higher-order photon numbers, the effect of which is further exacerbated by loss [242]. Overcoming these imperfect photon source issues is a ubiquitous problem in quantum optics experiments [329] but progress on developing better single photon sources is encouraging. Source multiplexing [330–333] and high-efficiency number-resolved photon detection will help to alleviate the higher-order contamination. Recent results from our group using silica-waveguide-based photon sources have shown excellent indistinguishability and photon state purity across many different waveguides [334] providing a potential route toward large-scale, low-loss source multiplexing.

One route toward lower loss photonic chips is to work in the telecommunications band at wavelengths around 1550nm. In this band, silica exhibits particularly low optical loss as well as increased mode-coupling efficiency. Working at 1550 nm should make it feasible to

build even larger scale circuits. In addition, the vast array of photonic products developed for the telecommunications industry becomes available to use at these wavelengths, including fast switches and modulators, which raises some hope of being able to implement the critical feed-forward operations required for LOQC. Unfortunately much of the work carried out so far has been restricted to operating at 800nm where commercial silicon avalanche photo-diodes (APDs) are reasonably efficient. Chapters 5 and 6 detailed the installation and on-chip integration of a superconducting detector, not only able to operate in the telecommunications band but with an efficiency far superior to any APD and with the unique capability of having intrinsic photon-number resolution — an essential feature for almost all LOQC schemes.

7.1.2 TES detectors

In chapter 5 we reviewed the available technologies for single photon detection and argued that detection efficiency is the single most important parameter in the context of LOQC. Until recently the silicon APD has been the work horse of single photon detection but is limited to detecting only visible wavelength photons. However, the small energy gap associated with superconducting detectors allows us to achieve high efficiency detection up into the infrared wavelengths including the important low-loss telecommunications band around 1550 nm.

The superconducting transition edge sensor is an advanced detector for LOQC which combines high efficiency detection at 1550 nm with high fidelity photon number resolution. We set up a new cryogenics lab and installed two TES detectors and characterised their performance using an attenuated coherent state. The measured noise performance was

compared to that theoretically achievable, confirming we had limited excess noise that allowed us to achieve clear photon number resolution at 1550 nm. We are currently in the process of improving the noise performance by upgrading the SQUIDs, adding a mu-metal magnetic shield and improving the room-temperature shielding of the signal cables. In addition, it should be possible to reduce the detector jitter to below 10 ns by mounting the TES closer to the SQUID. Such an improvement will allow us to run the detectors at a faster repetition rate.

Even taking account of the enforced reduction of laser repetition rate, these TESs will increase the data-rates of heralded, four-photon experiments by a factor of at least 3 compared to using silicon APDs. This enhancement will enable the next generation of on-chip quantum experiments, moving beyond the three-photons demonstrated here. We have just installed a fibre-link between the cryogenics lab housing the detectors and our main photonics lab. A telecomm-band single photon source is being built and should be available for experiments shortly. Using this source with the TESs and low-loss at 1550 nm it might be possible to achieve the first truly quantum enhanced precision measurement and surpass the standard quantum limit using loss-tolerant (Holland-Burnett) entangled states [335]. Quantum enhanced metrology schemes such as this demand not only very high efficiencies but require projections onto certain photon-number states, making the TES the ideal detector for such an experiment [252]. Most heralded entangling gates, which play a vital role in any LOQC scheme, require at least four photons to operate and often rely on photon-number discrimination making the TES the ideal detector to perform the first such on-chip experiments.

7.1.3 On-chip TES detectors

Integrating detectors on-chip removes one more coupling interface, allows for novel circuit designs and completes the integration of photonic circuit elements. In chapter 6 we introduced the evanescently coupled detector concept and demonstrated its use in the first number-resolving on-chip single photon detector. After successfully demonstrating a first generation test-device, detailed optical and thermal modelling was performed resulting in an optimised high-efficiency detector device design. The fabricated device performed well and was found to agree with simulations. We confirmed that such detectors can be multiplexed and confirmed that any excess scattering loss introduced by the detector regions was small ($< 5\%$).

Several options are being pursued in the search for even greater on-chip detection efficiency. One route being tested now is to simply try and optimise the transition temperature of the current double-sided tadpole design. It is hoped that with by achieving a lower T_C for the device we will be able to extend the absorbing tails to $2 \times 200\,\mu$m. This should allow us reach $\sim 70\%$ fractional absorption per device whilst still retaining acceptable energy resolution and timing performance.

Already, the on-chip detector presented in this chapter opens up some interesting possible experiments. The unique transmissive failure mode of these devices permits us to increase the effective saturation energy of a TES by multiplexing many 'low-efficiency' detectors together. This could provide a way to increase the saturation energy without degrading the energy resolution, opening up potentially interesting high-photon number metrology experiments and investigating the quantum-to-classical transition for high mean-photon number states. The excellent mode-matching of these detectors also allows them to be

used to perform perfectly efficient photon-subtraction. This is particularly appealing since the loss introduced by conventional off-chip detectors reduce the fidelity and purity of the subtracted state that can be achieved. This high-fidelity subtraction, together with the capability for high-efficiency vacuum heralding and photon number-resolution, provides a route toward performing the important entanglement distillation protocol on chip.

The greatest benefits of this work will come with further integration. The first simple passive quantum circuits with integrated detectors are now being fabricated. The first will feature simple beam splitters to perform an on-chip HOM experiment and directly measure the photon bunching at zero delay via the photon number-resolution.

However, to take full advantage of the integrated nature of the device, it will be necessary to develop a method of active phase control at cryogenic temperatures. Clearly, the thermo-optic phase shifters used earlier will not be suitable but we have recently demonstrated the use of the strain-optic effect in silica[234] to induce controllable phase shifts which could work in the dilution fridge. It is also possible to deposit electro-optic polymers above the unclad waveguides and achieve fast phase shifting through evanescent coupling[327]. Such active phase control allows for the construction of a range of heralded on-chip quantum gates such as the non-linear sign shift gate presented in section 2.2.2.

7.2 Outlook

The work in this thesis has demonstrated the suitability of the silica-on-silicon platform to perform complex on-chip logic operations as well as on-chip high-efficiency, number resolved photo-detection. Although outside the scope of this work, our group has also realised a silica-on-silicon based on-chip array of identical, pure, single photon sources, based on spontaneous

four-wave mixing[334]. Thus, the photon sources, circuits and detectors can all be built on one common platform. Although there are still technical challenges to overcome in achieving the complete integration of these components on a single chip, the intrinsic mode-matching, both between the chips and with single mode fibre, provides a low-loss interface between the different components. Although only one small part of the effort to reduce losses, further gains will be made from the move to telecomm wavelengths and use of the high efficiency TES detectors.

Further, the availability of low-loss high-frequency modulators and fast optical switches at this wavelength opens the possibility for performing operations in the time-frequency space. The time-frequency degree of freedom naturally provides access to a high dimensional Hilbert space in which we can realise an ever greater number of addressable optical modes. Exploiting this extra space points to a route toward manipulating many thousands of optical modes as required for full, universal LOQC[336].

Ultimately, the progress of linear optical quantum computing is limited by the lack of a suitable on-demand single photon source and the absence of a deterministic photon-photon non-linearity. These problems are currently addressed by instead using *probabilistic*, but heralded, photon sources and logic gates. This demands that we also posses a suitable *quantum memory* in order to correctly synchronise each of these probabilistic components[337]. Whilst these are subject to much active research, no system has yet achieved a noise low enough to be useful[338]. Instead, increasing effort is being devoted to building single photon sources based on isolated quantum systems and there have been promising recent results demonstrating single-photon-level single-atom-mediated non-linearities. Making such devices compatible with integrated photonic devices is the next challenge although

there have already been some promising recent results[339]. It is not clear yet if linear-optics alone will be practically sufficient to realise a universal quantum computer or if these matter-mediated systems will be required. Either way, programmable, low-loss integrated linear optical networks plus high-efficiency photon-detectors are going to be an essential feature of any future optical quantum machine. Finding a platform on which all components can then be successfully mode-matched will be a critical challenge. The excellent coupling to single-mode fibre makes the silica-on-silicon devices used in this thesis a promising starting point for beginning to realise quantum optics experiments in which all the key parts are built using an integrated photonic architecture.

Bibliography

[1] Landauer, R. The physical nature of information. *Physics Letters A* **217**, 188–193 (1996).

[2] Manin, Y. I. Computable and uncomputable. *Sovetskoye Radio, Moscow* (1980).

[3] Feynman, R. P. Simulating physics with computers. *International Journal of Theoretical Physics* **21**, 467–488 (1982).

[4] Deutsch, D. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **400**, 97–117 (1985).

[5] Deutsch, D. & Jozsa, R. Rapid Solution of Problems by Quantum Computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **439**, 553–558 (1992).

[6] Shor, P. Algorithms for quantum computation: discrete logarithms and factoring. *35th Annual Symposium on Foundations of Computer Science* 124–134 (1994).

[7] Shor, P. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM journal on computing* **26**, 124–134 (1997).

[8] Calderbank, A. & Shor, P. Good quantum error-correcting codes exist. *Physical Review A* **54**, 1098–1105 (1996).

[9] Steane, A. M. Error Correcting Codes in Quantum Theory. *Physical Review Letters* **77**, 793–797 (1996).

[10] Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Physical Review A (R)* **52**, 2493–2496 (1995).

[11] Laflamme, R., Miquel, C., Paz, J. P. & Zurek, W. H. Perfect Quantum Error Correcting Code. *Physical Review Letters* **77**, 198–201 (1996).

[12] Cirac, J. & Zoller, P. Quantum computations with cold trapped ions. *Physical Review Letters* **74**, 4091–4094 (1995).

[13] O'Brien, J. Optical quantum computing. *Science* **318**, 1567 (2007).

[14] Ladd, T., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C. & OBrien, J. Quantum computers. *Nature* **464**, 45–53 (2010).

[15] Nielsen, M. & Chuang, I. *Quantum computation and quantum information*, vol. 1 (Cambridge Univ Press, 2000).

[16] Monz, T. *et al.* 14-Qubit Entanglement: Creation and Coherence. *Physical Review Letters* **106**, 130506 (2011).

[17] Eisaman, M. D., Fan, J., Migdall, a. & Polyakov, S. V. Invited review article: Single-photon sources and detectors. *The Review of scientific instruments* **82**, 071101 (2011).

[18] Kok, P., Nemoto, K., Ralph, T. C., Dowling, J. P. & Milburn, G. J. Linear optical quantum computing with photonic qubits. *Reviews of Modern Physics* **79**, 135–174 (2007).

[19] O'Brien, J. L., Pryde, G. J., White, a. G., Ralph, T. C., Branning, D. & OBrien, J. Demonstration of an all-optical quantum controlled-NOT gate. *Nature* **426**, 264–7 (2003).

[20] Politi, A., Cryan, M. J., Rarity, J. G., Yu, S. & O'Brien, J. Silica-on-silicon waveguide quantum circuits. *Science* **320**, 646 (2008).

[21] Huntington, E. & Ralph, T. Components for optical qubits encoded in sideband modes. *Physical Review A* **69**, 1–5 (2004).

[22] Humphreys, P. C. *et al.* Linear Optical Quantum Computing in a Single Spatial Mode. *Physical Review Letters* **111**, 150501 (2013).

[23] Mair, A., Vaziri, A., Weihs, G. & Zeilinger, A. Entanglement of the orbital angular momentum states of photons. *Nature* **412**, 313–316 (2001).

[24] Leach, J., Padgett, M., Barnett, S., Franke-Arnold, S. & Courtial, J. Measuring the Orbital Angular Momentum of a Single Photon. *Physical Review Letters* **88**, 257901 (2002).

[25] Knill, E., Laflamme, R. & Milburn, G. A scheme for efficient quantum computation with linear optics. *Nature* **409**, 46–52 (2001).

[26] DiVincenzo, D. P. Two-bit gates are universal for quantum computation. *Physical Review A* **51**, 1015–1022 (1995).

[27] Kimble, H. J. The quantum internet. *Nature* **453**, 1023–30 (2008).

[28] Aspect, A., Dalibard, J. & Roger, G. Experimental test of Bell's inequalities using time-varying analyzers. *Physical Review Letters* **49**, 1804–1807 (1982).

[29] Zukowski, M., Zeilinger, A., Horne, M. & Ekert, A. Event-ready-detectors Bell experiment via entanglement swapping. *Physical Review Letters* **71**, 4287–4290 (1993).

[30] Bouwmeester, D., Pan, J.-w., Mattle, K., Eibl, M., Weinfurter, H. & Zeilinger, A. Experimental quantum teleportation. *Nature* **390**, 575–579 (1997).

[31] Weihs, G., Jennewein, T., Simon, C., Weinfurter, H. & Zeilinger, A. Violation of Bell's inequality under strict Einstein locality conditions. *Physical Review Letters* **81**, 5039–5043 (1998).

[32] Bouwmeester, D., Pan, J., Daniell, M., Weinfurter, H. & Zeilinger, A. Observation of three-photon Greenberger-Horne-Zeilinger entanglement. *Physical Review Letters* **82**, 1345–1349 (1999).

[33] Thompson, M., Politi, A., Matthews, J. & O'Brien, J. Integrated waveguide circuits for optical quantum computing. *Circuits, Devices Systems, IET* **5**, 94–102 (2011).

[34] Politi, A. *Integrated quantum photonics.* Ph.D. thesis, University of Bristol (2010).

[35] Politi, A., Matthews, J. C. F., Thompson, M. & O'Brien, J. Integrated quantum photonics. *IEEE Journal of Selected Topics in Quantum Electronics* **15**, 1673–1684 (2009).

[36] Metcalf, B. J. *et al.* Multiphoton quantum interference in a multiport integrated photonic device. *Nature Communications* **4**, 1356 (2013).

[37] Thomas-Peter, N. *et al.* Integrated photonic sensing. *New Journal of Physics* **13**, 055024 (2011).

[38] Wineland, D. J. Quantum information processing and quantum control with trapped atomic ions. *Physica Scripta* **T137**, 014007 (2009).

[39] Monroe, C. & Kim, J. Scaling the ion trap quantum processor. *Science* **339**, 1164–9 (2013).

[40] Clarke, J. & Wilhelm, F. K. Superconducting quantum bits. *Nature* **453**, 1031–42 (2008).

[41] Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: an outlook. *Science* **339**, 1169–74 (2013).

[42] Awschalom, D. D., Bassett, L. C., Dzurak, A. S., Hu, E. L. & Petta, J. R. Quantum spintronics: engineering and manipulating atom-like spins in semiconductors. *Science* **339**, 1174–9 (2013).

[43] Nizovtsev, A. P. A Quantum Computer Based on NV Centers in Diamond: Optically Detected Nutations of Single Electron and Nuclear Spins. *Optics and Spectroscopy* **99**, 233 (2005).

[44] Gisin, N. & Thew, R. Quantum communication. *Nature photonics* **1**, 165–171 (2007).

[45] Braunstein, S. L. & van Loock, P. Quantum information with continuous variables. *Reviews of Modern Physics* **77**, 513 (2005).

[46] Simon, M. K. & Alouini, M.-S. *Digital Communication over Fading Channels* (John Wiley & Sons, 2005).

[47] Menicucci, N. C. Fault-Tolerant Measurement-Based Quantum Computing with Continuous-Variable Cluster States. *Physical Review Letters* **112**, 120504 (2014).

[48] Braunstein, S. L. Error Correction for Continuous Quantum Variables. *Physical Review Letters* **80**, 4084–4087 (1998).

[49] Ralph, T. C. Quantum error correction of continuous-variable states against Gaussian noise. *Physical Review A* **84**, 022339 (2011).

[50] Aaronson, S. A linear-optical proof that the permanent is# P-hard. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **467**, 3393–3405 (2011).

[51] Paunković, N. *The Role of Indistinguishability of Identical Particles in Quantum Information Processing.* Ph.D. thesis, Oxford University (2004).

[52] Gullans, M., Chang, D., Koppens, F., de Abajo, F. & Lukin, M. Single-Photon Nonlinear Optics with Graphene Plasmons. *Physical Review Letters* **111**, 247401 (2013).

[53] Firstenberg, O., Peyronel, T., Liang, Q.-Y., Gorshkov, A. V., Lukin, M. D. & Vuletić, V. Attractive photons in a quantum nonlinear medium. *Nature* **502**, 71–5 (2013).

[54] Zeilinger, A. General properties of lossless beam splitters in interferometry. *American Journal of Physics* **49**, 882 (1981).

[55] Bussieres, F., Soudagar, Y., Berlin, G., Lacroix, S. & Godbout, N. Deterministic unitary operations on time-bin qudits for quantum communication. *arXiv preprint quant-ph/0608183v2* (2006).

[56] Olislager, L., Cussey, J., Nguyen, a., Emplit, P., Massar, S., Merolla, J.-M. & Huy, K. Frequency-bin entangled photons. *Physical Review A* **82**, 1–7 (2010).

[57] Reck, M., Zeilinger, A., Bernstein, H. J. & Bertani, P. Experimental Realisation of Any Discrete Unitary Operator. *Physical Review Letters* **73**, 58 (1994).

[58] Cerf, N., Adami, C. & Kwiat, P. Optical simulation of quantum logic. *Physical Review A* **57**, R1477–R1480 (1998).

[59] Kwiat, P. G., Mitchell, J. R., Schwindt, P. D. D. & White, A. G. Grover's search algorithm: An optical approach. *Journal of Modern Optics* **47**, 257–266 (2000).

[60] Hong, C., Ou, Z. & Mandel, L. Measurement of Subpicosecond Time Intervals between Two Photons by Interference. *Physical Review Letters* **59**, 2044–2046 (1987).

[61] Mandel, L. Quantum effects in one-photon and two-photon interference. *Reviews of Modern Physics* **71**, 274–282 (1999).

[62] Ghosh, R. & Mandel, L. Observation of nonclassical effects in the interference of two photons. *Physical Review Letters* **59**, 1903–1905 (1987).

[63] Ou, Z. & Mandel, L. Further Evidence of Nonclassical Behaviour in Optical Interference. *Physical Review Letters* **62**, 2941–2944 (1989).

[64] Barenco, A. *et al.* Elementary gates for quantum computation. *Physical Review A* **52**, 3457–3467 (1995).

[65] Lloyd, S. Almost Any Quantum Logic Gate is Universal. *Physical Review Letters* **75**, 346–349 (1995).

[66] Scheel, S. & Lütkenhaus, N. Upper bounds on success probabilities in linear optics. *New Journal of Physics* **6**, 51–51 (2004).

[67] Eisert, J. Optimizing Linear Optics Quantum Gates. *Physical Review Letters* **95**, 040502 (2005).

[68] Pittman, T., Jacobs, B. & Franson, J. Probabilistic quantum logic operations using polarizing beam splitters. *Physical Review A* **64**, 062311 (2001).

[69] Browne, D. E. & Rudolph, T. Resource-Efficient Linear Optical Quantum Computation. *Physical Review Letters* **95**, 010501 (2005).

[70] Ralph, T. Scaling of multiple postselected quantum gates in optics. *Physical Review A* **70**, 1–5 (2004).

[71] Bennett, C., Brassard, G., Crépeau, C., Jozsa, R., Peres, A. & Wootters, W. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters* **70**, 1895–1899 (1993).

[72] Lütkenhaus, N., Calsamiglia, J. & Suominen, K.-A. Bell measurements for teleportation. *Physical Review A* **59**, 3295–3300 (1999).

[73] Calsamiglia, J. & Lütkenhaus, N. Maximum efficiency of a linear-optical Bell-state analyzer. *Applied Physics B* **72**, 67–71 (2014).

[74] Raussendorf, R. & Briegel, H. J. A One-Way Quantum Computer. *Physical Review Letters* **86**, 5188–5191 (2001).

[75] Varnava, M., Browne, D. & Rudolph, T. How Good Must Single Photon Sources and Detectors Be for Efficient Linear Optical Quantum Computation? *Physical Review Letters* **100**, 1–4 (2008).

[76] Browne, D. E. & Rudolph, T. Resource-Efficient Linear Optics Quantum Computation via the Cluster State Approach. In *Quantum Information and Quantum Control Conference* (Fields Institute, Toronto, 2004).

[77] Verstraete, F. & Cirac, J. Valence-bond states for quantum computation. *Physical Review A* **70**, 060302(R) (2004).

[78] Wilde, M., Spedalieri, F., Dowling, J. P. & Lee, H. Alternate Scheme for Optical Cluster-State Generation without Number-Resolving Photon Detectors. *International Journal of Quantum Information* **05**, 617–626 (2007).

[79] Aspect, A., Grangier, P. & Roger, G. Experimental Tests of Realistic Local Theories via Bell's Theorem. *Physical Review Letters* **47**, 460–463 (1981).

[80] Clauser, J. F. & Shimony, A. Bell's theorem. Experimental tests and implications. *Reports on Progress in Physics* **41**, 1881–1927 (1978).

[81] Kwiat, P. P., Mattle, K., Weinfurter, H., Zeilinger, A., Sergienko, A. A. & Shih, Y. New high-intensity source of polarization-entangled photon pairs. *Physical Review Letters* **75**, 4337–4341 (1995).

[82] Rarity, J. & Tapster, P. Experimental violation of Bells inequality based on phase and momentum. *Physical Review Letters* **64**, 2495–2498 (1990).

[83] Shih, Y. H. & Alley, C. O. New Type of Einstein-Podolsky-Rosen-Bohm Experiment Using Pairs of Light Quanta Produced by Optical Parametric Down Conversion. *Physical Review Letters* **61**, 2921–2924 (1988).

[84] Ou, Z. Y. & Mandel, L. Violation of Bell's Inequality and Classical Probability in a Two-Photon Correlation Experiment. *Physical Review Letters* **61**, 50–53 (1988).

[85] Kwiat, P. G., Steinberg, A. M. & Chiao, R. Y. High-visibility interference in a Bell-inequality experiment for energy and time. *Physical Review A* **47**, R2472–R2475 (1993).

[86] Boschi, D., Branca, S., Martini, F. D. & Hardy, L. Experimental realization of teleporting an unknown pure quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters* **80**, 1121–1125 (1998).

[87] O'Brien, J., Pryde, G., Gilchrist, A., James, D., Langford, N., Ralph, T. & White, A. Quantum Process Tomography of a Controlled-NOT Gate. *Physical Review Letters* **93**, 1–4 (2004).

[88] Okamoto, R., Hofmann, H., Takeuchi, S. & Sasaki, K. Demonstration of an Optical Quantum Controlled-NOT Gate without Path Interference. *Physical Review Letters* **95**, 1–4 (2005).

[89] Kiesel, N., Schmid, C., Weber, U., Ursin, R. & Weinfurter, H. Linear Optics Controlled-Phase Gate Made Simple. *Physical Review Letters* **95**, 1–4 (2005).

[90] Langford, N. *et al.* Demonstration of a Simple Entangling Optical Gate and Its Use in Bell-State Analysis. *Physical Review Letters* **95**, 3–6 (2005).

[91] Gasparoni, S., Pan, J.-W., Walther, P., Rudolph, T. & Zeilinger, A. Realization of a Photonic Controlled-NOT Gate Sufficient for Quantum Computation. *Physical Review Letters* **93**, 1–4 (2004).

[92] Pittman, T., Fitch, M., Jacobs, B. & Franson, J. Experimental controlled-NOT logic gate for single photons in the coincidence basis. *Physical Review A* **68**, 032316 (2003).

[93] Lanyon, B. *et al.* Simplifying quantum logic using higher-dimensional Hilbert spaces. *Nature Physics* **5**, 134–140 (2008).

[94] Pittman, T., Jacobs, B. & Franson, J. Quantum computing using linear optics. *John Hopkins APL Technical Digest* **25**, 84–90 (2004).

[95] Sanaka, K. & Kawahara, K. Franson-type experiment realizes two-qubit quantum logic gate. *Arxiv preprint quant-ph/0108001* 3–8 (2001).

[96] de Riedmatten, H., Marcikic, I., Scarani, V., Tittel, W., Zbinden, H. & Gisin, N. Tailoring photonic entanglement in high-dimensional Hilbert spaces. *Physical Review A* **69**, 050304 (2004).

[97] Hayat, A., Xing, X., Feizpour, A. & Steinberg, A. M. Multidimensional quantum information based on single-photon temporal wavepackets. *Optics express* **20**, 29174–84 (2012).

[98] Nisbet-Jones, P. B. R., Dilley, J., Holleczek, A., Barter, O. & Kuhn, A. Photonic qubits, qutrits and ququads accurately prepared and delivered on demand. *New Journal of Physics* **15**, 053007 (2013).

[99] Capmany, J. & Fernandez-Pousa, C. R. Realization of Single-Photon Frequency-Domain Qubit Channels Using Phase Modulators. *IEEE Photonics Journal* **4**, 2074–2084 (2012).

[100] Garcia-Escartin, J. C. & Chamorro-Posada, P. Quantum computer networks with the orbital angular momentum of light. *Physical Review A* **86**, 032334 (2012).

[101] García-Escartín, J. C. & Chamorro-Posada, P. Universal quantum computation with the orbital angular momentum of a single photon. *Journal of Optics* **13**, 064022 (2011).

[102] Krenn, M., Huber, M., Fickler, R., Lapkiewicz, R., Ramelow, S. & Zeilinger, A. Generation and confirmation of a (100 x 100)-dimensional entangled quantum system. *Proceedings of the National Academy of Sciences* **111**, 6243–7 (2014).

[103] Okamoto, R., O'Brien, J., Hofmann, H. & Takeuchi, S. Realization of a Knill-Laflamme-Milburn controlled-NOT photonic quantum circuit combining effective optical nonlinearities. *Proceedings of the National Academy of Sciences* **108**, 10067 (2011).

[104] Gao, W.-B. *et al.* Teleportation-based realization of an optical quantum two-qubit entangling gate. *Proceedings of the National Academy of Sciences* **107**, 20869–74 (2010).

[105] Prevedel, R., Walther, P., Tiefenbacher, F., Böhi, P., Kaltenbaek, R., Jennewein, T. & Zeilinger, A. High-speed linear optics quantum computing using active feed-forward. *Nature* **445**, 65–69 (2007).

[106] OBrien, J., Pryde, G., White, A. & Ralph, T. High-fidelity Z-measurement error encoding of optical qubits. *Physical Review A* **71**, 060303 (2005).

[107] Gao, W.-B. B. *et al.* Experimental Realization of a Controlled-NOT Gate with Four-Photon Six-Qubit Cluster States. *Physical Review Letters* **104**, 2–5 (2010).

[108] Kalasuwan, P. *et al.* Simple scheme for expanding photonic cluster states for quantum information. *Journal of the Optical Society of America B* **27**, 181–184 (2010).

[109] Vallone, G., Pomarico, E., Mataloni, P., De Martini, F. & Berardi, V. Realization and Characterization of a Two-Photon Four-Qubit Linear Cluster State. *Physical Review Letters* **98**, 1–4 (2007).

[110] Zhang, A., Lu, C., Zhou, X., Chen, Y., Zhao, Z., Yang, T. & Pan, J. Experimental construction of optical multiqubit cluster states from Bell states. *Physical Review A* **73**, 22330 (2006).

[111] Kiesel, N., Schmid, C., Weber, U., Tóth, G., Gühne, O., Ursin, R. & Weinfurter, H. Experimental Analysis of a Four-Qubit Photon Cluster State. *Physical Review Letters* **95**, 1–4 (2005).

[112] Walther, P. *et al.* Experimental one-way quantum computing. *Nature* **434**, 169 (2005).

[113] Lu, C.-Y. *et al.* Experimental entanglement of six photons in graph states. *Nature Physics* **3**, 91–95 (2007).

[114] Yao, X.-C. *et al.* Experimental demonstration of topological error correction. *Nature* **482**, 489–494 (2012).

[115] Yao, X.-C. *et al.* Observation of eight-photon entanglement. *Nature Photonics* **6**, 225–228 (2012).

[116] Barz, S., Kassal, I., Ringbauer, M., Lipp, Y. O., Dakić, B., Aspuru-Guzik, A. & Walther, P. A two-qubit photonic quantum processor and its application to solving systems of linear equations. *Scientific reports* **4**, 6115 (2014).

[117] Cai, X. D. *et al.* Experimental Quantum Computing to Solve Systems of Linear Equations. *arXiv preprint arXiv:13024310* 7 (2013).

[118] Barz, S., Kashefi, E., Broadbent, A., Fitzsimons, J. F., Zeilinger, A. & Walther, P. Demonstration of blind quantum computing. *Science* **335**, 303–8 (2012).

[119] Barz, S., Fitzsimons, J. F., Kashefi, E. & Walther, P. Experimental verification of quantum computation. *Nature Physics* **9**, 727–731 (2013).

[120] Bell, B. A., Herrera-Martí, D. A., Tame, M. S., Markham, D., Wadsworth, W. J. & Rarity, J. G. Experimental demonstration of a graph state quantum error-correction code. *Nature communications* **5**, 3658 (2014).

[121] Smith, B. J., Kundys, D., Thomas-Peter, N. L., Smith, P. G. R. & Walmsley, I. A. Phase-controlled integrated photonic quantum circuits. *Optics Express* **17**, 13516–13525 (2009).

[122] Shadbolt, P. J. P. *et al.* Generating, manipulating and measuring entanglement and mixture with a reconfigurable photonic circuit. *Nature Photonics* **6**, 45–49 (2011).

[123] Matthews, J., Politi, A., Bonneau, D. & OBrien, J. L. Heralding Two-Photon and Four-Photon Path Entanglement on a Chip. *Physical Review Letters* **107**, 1–5 (2011).

[124] Peruzzo, A., Laing, A., Politi, A., Rudolph, T. & O'Brien, J. L. Multimode quantum interference of photons in multiport integrated devices. *Nature Communications* **2**, 224 (2011).

[125] Li, H. W. *et al.* Reconfigurable controlled two-qubit operation on a quantum photonic chip. *New Journal of Physics* **13**, 115009 (2011).

[126] Politi, A., Matthews, J. C. F. & O'Brien, J. Shors quantum factoring algorithm on a photonic chip. *Science* **325**, 1221 (2009).

[127] Spring, J. B. *et al.* Boson Sampling on a Photonic Chip. *Science* **339**, 798–801 (2013).

[128] Marshall, G., Politi, A., Matthews, J., Dekker, P., Ams, M., Withford, M. & O'Brien, J. Laser written waveguide photonic quantum circuits. *Optics Express* **17**, 12546–12554 (2009).

[129] Owens, J. O. *et al.* Two-photon quantum walks in an elliptical direct-write waveguide array. *New Journal of Physics* **13**, 075003 (2011).

[130] Xiong, C. *et al.* Generation of correlated photon pairs in a chalcogenide As2S3 waveguide. *Applied Physics Letters* **98**, 119902 (2011).

[131] Spagnolo, N. *et al.* Three-photon bosonic coalescence in an integrated tritter. *Nature communications* **4** (2013).

[132] Crespi, A. *et al.* Anderson localization of entangled photons in an integrated quantum walk. *Nature Photonics* **7**, 322–328 (2013).

[133] Crespi, A. *et al.* Integrated multimode interferometers with arbitrary designs for photonic boson sampling. *Nature Photonics* **7**, 545–549 (2013).

[134] Corrielli, G. *et al.* Rotated waveplates in integrated waveguide optics. *Nature communications* **5**, 4249 (2014).

[135] Spagnolo, N. *et al.* Quantum interferometry with three-dimensional geometry. *Scientific reports* **2**, 862 (2012).

[136] Sansoni, L., Sciarrino, F., Vallone, G., Mataloni, P., Crespi, A., Ramponi, R. & Osellame, R. Two-Particle Bosonic-Fermionic Quantum Walk via Integrated Photonics. *Physical Review Letters* **108**, 1–5 (2012).

[137] Crespi, A. *et al.* Integrated photonic quantum gates for polarization qubits. *Nature communications* **2**, 566 (2011).

[138] Sansoni, L., Sciarrino, F., Vallone, G., Mataloni, P., Crespi, A., Ramponi, R. & Osellame, R. Polarization Entangled State Measurement on a Chip. *Physical Review Letters* **20**, 200503 (2010).

[139] Spring, J. B. *Single Photon Generation and Quantum Computing with Integrated Photonics*. Ph.D. thesis, University of Oxford (2014).

[140] Spring, J. B. *et al.* On-chip low loss heralded source of pure single photons. *Optics Express* **21**, 13522 (2013).

[141] Bonneau, D. *et al.* Quantum interference and manipulation of entanglement in silicon wire waveguide quantum circuits. *New Journal of Physics* **14**, 045003 (2012).

[142] Pernice, W. H. P., Schuck, C., Minaeva, O., Li, M., Goltsman, G. N., Sergienko, A. V. & Tang, H. X. High-speed and high-efficiency travelling wave single-photon detectors embedded in nanophotonic circuits. *Nature communications* **3**, 1325 (2012).

[143] Peruzzo, A. *et al.* Quantum walks of correlated photons. *Science* **329**, 1500–1503 (2010).

[144] Wang, J. *et al.* Gallium arsenide (GaAs) quantum photonic waveguide circuits. *Optics Communications* **327**, 49–55 (2014).

[145] Martin, A., Alibart, O., De Micheli, M. P., Ostrowsky, D. B. & Tanzilli, S. A quantum relay chip based on telecommunication integrated optics technology. *New Journal of Physics* **14**, 025002 (2012).

[146] Crespi, A. *et al.* Measuring protein concentration with entangled photons. *Applied Physics Letters* **100**, 233704 (2012).

[147] Peruzzo, A., Shadbolt, P., Brunner, N., Popescu, S. & O'Brien, J. L. A quantum delayed-choice experiment. *Science* **338**, 634–637 (2012).

[148] Peruzzo, A. *et al.* A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* **5** (2014).

[149] Spagnolo, N. *et al.* General Rules for Bosonic Bunching in Multimode Interferometers. *Physical Review Letters* **111**, 130503 (2013).

[150] Tillmann, M. *et al.* BosonSampling with Controllable Distinguishability. *arXiv preprint arXiv:1403.3433* (2014).

[151] Carolan, J. *et al.* On the experimental verification of quantum complexity in linear optics. *Nature Photonics* **8**, 621–626 (2014).

[152] Broome, M. A., Fedrizzi, A., Rahimi-Keshari, S., Dove, J., Aaronson, S., Ralph, T. C. & White, A. G. Photonic boson sampling in a tunable circuit. *Science* **339**, 794–8 (2013).

[153] Spagnolo, N. *et al.* Experimental validation of photonic boson sampling. *Nature Photonics* **8**, 615–620 (2014).

[154] Matthews, J., Politi, A., Stefanov, A. & O'Brien, J. Manipulation of multiphoton entanglement in waveguide quantum circuits. *Nature Photonics* **3**, 346–350 (2009).

[155] Sansoni, L. *Integrated devices for quantum information with polarization encoded qubits*. Ph.D. thesis, Sapienza, University of Rome (2012).

[156] Pelton, M., Santori, C., Vucković, J., Zhang, B., Solomon, G., Plant, J. & Yamamoto, Y. Efficient Source of Single Photons: A Single Quantum Dot in a Micropost Microcavity. *Physical Review Letters* **89**, 233602 (2002).

[157] Michler, P. *et al.* A quantum dot single-photon turnstile device. *Science* **290**, 2282–5 (2000).

[158] Mizuochi, N. *et al.* Electrically driven single-photon source at room temperature in diamond. *Nature Photonics* **6**, 299–303 (2012).

[159] Barros, H. G., Stute, A., Northup, T. E., Russo, C., Schmidt, P. O. & Blatt, R. Deterministic single-photon source from a single ion. *New Journal of Physics* **11**, 103004 (2009).

[160] Keller, M., Lange, B., Hayasaka, K., Lange, W. & Walther, H. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature* **431**, 1075–8 (2004).

[161] Nisbet-Jones, P. B. R., Dilley, J., Ljunggren, D. & Kuhn, A. Highly efficient source for indistinguishable single photons of controlled shape. *New Journal of Physics* **13**, 103036 (2011).

[162] McKeever, J., Boca, A., Boozer, A. D., Miller, R., Buck, J. R., Kuzmich, A. & Kimble, H. J. Deterministic generation of single photons from one atom trapped in a cavity. *Science* **303**, 1992–4 (2004).

[163] Kuhn, A., Hennrich, M. & Rempe, G. Deterministic Single-Photon Source for Distributed Quantum Networking. *Physical Review Letters* **89**, 067901 (2002).

[164] Pezzagna, S., Rogalla, D., Wildanger, D., Meijer, J. & Zaitsev, A. Creation and nature of optical centres in diamond for single-photon emissionoverview and critical remarks. *New Journal of Physics* **13**, 035024 (2011).

[165] Tamarat, P. *et al.* Stark Shift Control of Single Optical Centers in Diamond. *Physical Review Letters* **97**, 083002 (2006).

[166] Kurtsiefer, C., Mayer, S., Zarda, P. & Weinfurter, H. Stable solid-state source of single photons. *Physical Review Letters* **85**, 290–3 (2000).

[167] Yuan, Z. *et al.* Electrically driven single-photon source. *Science* **295**, 102–5 (2002).

[168] Grice, W. & Walmsley, I. Spectral information and distinguishability in type-II down-conversion with a broadband pump. *Physical Review A* **56**, 1627–1634 (1997).

[169] Grice, W. P., U'Ren, A. B. & Walmsley, I. A. Eliminating frequency and space-time correlations in multiphoton states. *Physical Review A* **64**, 63815 (2001).

[170] Mosley, P. J., Lundeen, J. S., Smith, B. J., Wasylczyk, P., U'Ren, A. B., Silberhorn, C. & Walmsley, I. A. Heralded Generation of Ultrafast Single Photons in Pure Quantum States. *Phys. Rev. Lett.* **100**, 133601 (2008).

[171] Mosley, P. *Generation of Heralded Single Photons in Pure Quantum States*. Ph.D. thesis, Hertford College, University of Oxford, Oxford University (2007).

[172] Mosley, P. J., Lundeen, J. S., Smith, B. J. & Walmsley, I. a. Conditional preparation of single photons using parametric downconversion: a recipe for purity. *New Journal of Physics* **10**, 093011 (2008).

[173] Gerchberg, B. R. W. & Saxton, W. O. A Practical Algorithm for the Determination of Phase from Image and Diffraction Plane Pictures. *Optik* **35** (1972).

[174] Thomas-Peter, N., Smith, B., Datta, A., Zhang, L., Dorner, U. & Walmsley, I. Real-World Quantum Sensors: Evaluating Resources for Precision Measurement. *Physical Review Letters* **107**, 1–5 (2011).

[175] Kao, C. K. Nobel Prize in Physics for the groundbreaking achievements concerning the transmission of light in fibers for optical communication. Tech. Rep., Royal Swedish Academy of Sciences, Stockholm (2009).

[176] Bernsee, G. Light conducting fibers of quartz glass (1976). URL `http://www.google.com/patents/US3966300`.

[177] Adikan, F. R. M. *Direct UV-written waveguide devices*. Ph.D. thesis, University of Southampton (2007).

[178] Sparrow, I. *Development and applications of UV written waveguides.* Ph.D. thesis, University of Southampton (2005).

[179] Svalgaard, M. *Ultraviolet light induced refractive index structures in germanosilica.* Ph.D. thesis, Technical University of Denmark (1997).

[180] Kawachi, M. Recent progress in silica-based planar lightwave circuits on silicon. *IEEE Optoelectronics, IET* **143**, 257–262 (1996).

[181] Johlen, D., Klose, P. & Brinkmeyer, E. UV-written directional couplers in silica on silicon. In *Proceedings of Optical Fiber Communication Conference*, 279–280 (Opt. Soc. America, 1997).

[182] Kundys, D., Gates, J. & Dasgupta, S. Use of cross-couplers to decrease size of UV written photonic circuits. *IEEE Photonics Technology Letters* **21**, 947–949 (2009).

[183] Emmerson, G. *et al.* All-UV-written integrated planar Bragg gratings and channel waveguides through single-step direct grating writing. *IEEE Optoelectronics, IET* **151**, 119 (2004).

[184] Sima, C., Gates, J. C., Rogers, H. L., Mennea, P. L., Holmes, C., Zervas, M. N. & Smith, P. G. R. Ultra-wide detuning planar Bragg grating fabrication technique based on direct UV grating writing with electro-optic phase modulation. *Optics express* **21**, 15747–54 (2013).

[185] Shabani, A. *et al.* Efficient Measurement of Quantum Dynamics via Compressive Sensing. *Physical Review Letters* **106**, 100401 (2011).

[186] Langford, N. *Encoding, Manipulating and Measuring Quantum Information in Optics.* Ph.D. thesis, University of Queensland (2007).

[187] Rodionov, A. V. *et al.* Compressed sensing quantum process tomography for superconducting quantum gates. *Physical Review A* **144504**, 1–16 (2014).

[188] Knill, E. *et al.* Randomized benchmarking of quantum gates. *Physical Review A* **77**, 012307 (2008).

[189] da Silva, M. P., Landon-Cardinal, O. & Poulin, D. Practical Characterization of Quantum Devices without Tomography. *Physical Review Letters* **107**, 210404 (2011).

[190] Rahimi-Keshari, S., Broome, M. A., Fickler, R., Fedrizzi, A., Ralph, T. C. & White, A. G. Direct characterization of linear-optical networks. *Optics Express* **21**, 13450–8 (2013).

[191] Laing, A. & O'Brien, J. Super-stable tomography of any linear optical device. *arXiv preprint arXiv:1208.2868* 1–6 (2012).

[192] Thomas-peter, N. *Quantum-Enhanced Precision Measurement and Information Processing with Integrated Photonics.* Ph.D. thesis, University of Oxford (2012).

[193] Lanyon, B., Weinhold, T., Langford, N., OBrien, J., Resch, K., Gilchrist, A. & White, A. Manipulating biphotonic qutrits. *Physical Review Letters* **100**, 60504 (2008).

[194] Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. *Proceedings of the 43rd annual ACM symposium on the Theory of Computing* 1–96 (2011).

[195] Valiant, L. The complexity of computing the permanent. *Theoretical Computer Science* **8**, 189–201 (1979).

[196] Motes, K., Dowling, J. & Rohde, P. Spontaneous parametric down-conversion photon sources are scalable in the asymptotic limit for boson sampling. *Physical Review A* (2013).

[197] Shen, C., Zhang, Z. & Duan, L. Scalable Implementation of Boson Sampling with Trapped Ions. *Physical Review Letters* 050504 (2014).

[198] Tillmann, M., Dakić, B., Heilmann, R., Nolte, S., Szameit, A. & Walther, P. Experimental boson sampling. *Nature Photonics* **7**, 1–7 (2013).

[199] Tichy, M., Mayer, K., Buchleitner, A. & Mø lmer, K. Stringent and efficient assessment of Boson-Sampling devices. *Physical Review Letters* **113**, 020502 (2014).

[200] Aaronson, S. & Arkhipov, A. Boson sampling is far from uniform. *arXiv preprint arXiv:1309.7460* (2013).

[201] Shchesnovich, V. S. Conditions for an experimental Boson-sampling computer to disprove the Extended Church-Turing thesis. *Physical Review A* **89**, 5 (2014).

[202] Lund, A. P., Laing, A., Rahimi-Keshari, S., Rudolph, T., O'Brien, J. L. & Ralph, T. C. Boson Sampling from a Gaussian State. *Phys. Rev. Lett.* **113**, 100502 (2014).

[203] Gogolin, C., Kliesch, M., Aolita, L. & Eisert, J. Boson-Sampling in the light of sample complexity. *arXiv preprint arXiv:1306.3995* 22 (2013).

[204] Leverrier, A. & García-Patrón, R. Does Boson Sampling need Fault-Tolerance? *Arxiv preprint arXiv:1309.4687* 20 (2013).

[205] Rohde, P. P., Motes, K. R., Knott, P., Munro, W. & Dowling, J. P. Sampling generalized cat states with linear optics is probably hard. *arXiv preprint arXiv:1310.0297* 8 (2013).

[206] Rohde, P. P., Motes, K. R., Knott, P. A. & Munro, W. J. Will boson-sampling ever disprove the Extended Church-Turing thesis? *arXiv preprint arXiv:1401.2199* (2014).

[207] Seshadreesan, K. P., Olson, J. P., Motes, K. R., Rohde, P. P. & Dowling, J. P. Boson sampling with photon-added coherent states. *arXiv preprint arXiv:1402.0531* (2014).

[208] Gard, B. T., Motes, K. R., Olson, J. P., Rohde, P. P. & Dowling, J. P. An introduction to boson-sampling. *arXiv preprint arXiv:1406.6767* 13 (2014).

[209] Olson, J. P., Seshadreesan, K. P., Motes, K. R., Rohde, P. P. & Dowling, J. P. Sampling arbitrary photon-added or photon-subtracted squeezed states is in the same complexity class as boson sampling. *arXiv preprint arXiv:1406.7821* 5 (2014).

[210] Motes, K., Gilchrist, A., Dowling, J. P. & Rohde, P. P. Scalable boson-sampling with time-bin encoding using a loop-based architecture. *Phys. Rev. Lett.* **113**, 120501 (2014).

[211] Rohde, P. & Ralph, T. Error tolerance of the boson-sampling model for linear optics quantum computing. *Physical Review A* **85**, 1–5 (2012).

[212] Laing, A. *et al.* High-fidelity operation of quantum photonic circuits. *Applied Physics Letters* **97**, 211109 (2010).

[213] Fedorov, A., Steffen, L., Baur, M., da Silva, M. P. & Wallraff, A. Implementation of a Toffoli gate with superconducting circuits. *Nature* **481**, 170–2 (2012).

[214] Gottesman, D. & Chuang, I. L. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature* **402**, 390–393 (1999).

[215] Gisin, N. Nonlocality criteria for quantum teleportation. *Physics Letters A* **210**, 157–159 (1996).

[216] Massar, S. & Popescu, S. Optimal extraction of information from finite quantum ensembles. *Physical review letters* **74**, 1259–1263 (1995).

[217] Li, Y. Averages over the unitary group. In *Selected Topics in Mathematical Physics: Quantum Information Theory (http://www.mpmueller.net/seminar.html)*, vol. 2013, chap. 6, 1–6 (2013).

[218] Kim, Y.-H., Kulik, S. & Shih, Y. Quantum Teleportation of a Polarization State with a Complete Bell State Measurement. *Physical Review Letters* **86**, 1370–1373 (2001).

[219] Furusawa, A., Sorensen, J. L., Braunstein, S. L., Fuchs, C., Kimble, H. & Polzik, E. S. Unconditional Quantum Teleportation. *Science* **282**, 706–709 (1998).

[220] Takeda, S., Mizuta, T., Fuwa, M., van Loock, P. & Furusawa, A. Deterministic quantum teleportation of photonic quantum bits by a hybrid technique. *Nature* **500**, 315–8 (2013).

[221] Marcikic, I., de Riedmatten, H., Tittel, W., Zbinden, H. & Gisin, N. Long-distance teleportation of qubits at telecommunication wavelengths. *Nature* **421**, 509–13 (2003).

[222] Ursin, R., Jennewein, T., Aspelmeyer, M., Kaltenbaek, R., Lindenthal, M., Walther, P. & Zeilinger, A. Communications: quantum teleportation across the Danube. *Nature* **430**, 849 (2004).

[223] Jin, X.-M. *et al.* Experimental free-space quantum teleportation. *Nature Photonics* **4**, 376–381 (2010).

[224] Ma, X.-S. *et al.* Quantum teleportation over 143 kilometres using active feed-forward. *Nature* **489**, 269–73 (2012).

[225] Merali, Z. Data teleportation: The quantum space race. *Nature* **492**, 22–5 (2012).

[226] Nilsson, J. *et al.* Quantum teleportation using a light-emitting diode. *Nature Photonics* **7**, 311–315 (2013).

[227] Riebe, M. *et al.* Deterministic quantum teleportation with atoms. *Nature* **429**, 734–7 (2004).

[228] Barrett, M. D. *et al.* Deterministic quantum teleportation of atomic qubits. *Nature* **429**, 737–9 (2004).

[229] Steffen, L. *et al.* Realization of Deterministic Quantum Teleportation with Solid State Qubits. *Nature* **500**, 319–322 (2013).

[230] Baur, M., Fedorov, A., Steffen, L., Filipp, S., da Silva, M. P. & Wallraff, A. Benchmarking a Quantum Teleportation Protocol in Superconducting Circuits Using Tomography and an Entanglement Witness. *Physical Review Letters* **108**, 040502 (2012).

[231] Pfaff, W. *et al.* Unconditional quantum teleportation between distant solid-state quantum bits. *Science* **345**, 532–535 (2014).

[232] Ralph, T., Langford, N., Bell, T. & White, A. Linear optical controlled-NOT gate in the coincidence basis. *Physical Review A* **65**, 62324 (2002).

[233] Pittman, T. B., Jacobs, B. C. & Franson, J. D. Demonstration of feed-forward control for linear optics quantum computation. *Physical Review A* **66**, 052305 (2002).

[234] Humphreys, P. C. *et al.* Strain-optic active control for quantum integrated photonics. *Optics express* **22**, 21719–21726 (2014).

[235] Lee, H., Chen, T., Li, J., Painter, O. & Vahala, K. J. Ultra-low-loss optical delay line on a silicon chip. *Nature communications* **3**, 867 (2012).

[236] de Burgh, M., Langford, N. K., Doherty, A. & Gilchrist, A. Choice of measurement sets in qubit tomography. *Physical Review A* **78**, 1–15 (2008).

[237] Altepeter, J., Jeffrey, E. & Kwiat, P. Photonic state tomography. *Advances in Atomic, Molecular, and Optical Physics* **52**, 105–159 (2005).

[238] Tichy, M. C., Tiersch, M., Mintert, F. & Buchleitner, A. Many-particle interference beyond many-boson and many-fermion statistics. *New Journal of Physics* **14**, 093015 (2012).

[239] Tichy, M. C. Private communication (2014).

[240] Tichy, M. Interference of Identical Particles from Entanglement to Boson-Sampling. *J. Phys. B: At. Mol. Opt. Phys* **47**, 103001 (2014).

[241] Banaszek, K. Optimal quantum teleportation with an arbitrary pure state. *Physical Review A* (2000).

[242] Barbieri, M., Weinhold, T. J., Lanyon, B. P., Gilchrist, A., Resch, K. J., Almeida, M. P. & White, A. G. Parametric downconversion and optical quantum gates: two's company, four's a crowd. *Journal of Modern Optics* **56**, 209–214 (2009).

[243] Rogers, H. L., Ambran, S., Holmes, C., Smith, P. G. R. & Gates, J. C. In situ loss measurement of direct UV-written waveguides using integrated Bragg gratings. *Optics Letters* **35**, 2849–2851 (2010).

[244] Migdall, A. Introduction to Journal of Modern Optics special issue on single-photon: detectors, applications, and measurement methods. *Journal of Modern Optics* **51**, 1265–1266 (2004).

[245] Hadfield, R. Single-photon detectors for optical quantum information applications. *Nature Photonics* **3**, 696–705 (2009).

[246] Silberhorn, C. Detecting quantum light. *Contemporary Physics* **48**, 143–156 (2007).

[247] Ralph, T., Hayes, A. & Gilchrist, A. Loss-Tolerant Optical Qubits. *Physical Review Letters* **95**, 100501 (2005).

[248] Hayes, A., Gilchrist, A. & Ralph, T. Loss-tolerant operations in parity-code linear optics quantum computing. *Physical Review A* **77**, 012310 (2008).

[249] Varnava, M., Browne, D. E. & Rudolph, T. Loss Tolerance in One-Way Quantum Computation via Counterfactual Error Correction. *Phys. Rev. Lett.* **97**, 120501 (2006).

[250] Larsson, J.-A. k. Bells inequality and detector inefficiency. *Physical Review A* **57**, 3304–3308 (1998).

[251] Pearle, P. Hidden-Variable Example Based upon Data Rejection. *Physical Review D* **2**, 1418–1425 (1970).

[252] Datta, A., Zhang, L., Smith, B. J., Walmsley, I. A., Thomas-Peter, N. & Dorner, U. Quantum metrology with imperfect states and detectors. *Physical Review A* **83**, 063836 (2011).

[253] Gong, Y.-X., Zou, X.-B., Ralph, T. C., Zhu, S.-N. & Guo, G.-C. Linear optical quantum computation with imperfect entangled photon-pair sources and inefficient nonphoton-number-resolving detectors. *Physical Review A* **81**, 052303 (2010).

[254] Silva, M., Rötteler, M. & Zalka, C. Thresholds for linear optics quantum computing with photon loss at the detectors. *Physical Review A* **72**, 032307 (2005).

[255] Jahanmirinejad, S., Frucci, G., Mattioli, F., Sahin, D., Gaggero, A., Leoni, R. & Fiore, A. Photon-number resolving detector based on a series array of superconducting nanowires. *Applied Physics Letters* **101**, 072602 (2012).

[256] Divochiy, A. *et al.* Superconducting nanowire photon-number-resolving detector at telecommunication wavelengths. *Nature Photonics* **2**, 302–306 (2008).

[257] Fitch, M. M., Jacobs, B., Pittman, T. & Franson, J. Photon-number resolution using time-multiplexed single-photon detectors. *Physical Review A* **68**, 043814 (2003).

[258] Achilles, D. *et al.* Photon-number-resolving detection using time-multiplexing. *Journal of Modern Optics* **51**, 1499 (2004).

[259] Achilles, D., Silberhorn, C., Sliwa, C., Banaszek, K. & Walmsley, I. a. Fiber-assisted detection with photon number resolution. *Optics letters* **28**, 2387–9 (2003).

[260] Lacaita, a., Zappa, F., Cova, S. & Lovati, P. Single-photon detection beyond 1 m: performance of commercially available InGaAs/lnP detectors. *Applied optics* **35**, 2986–96 (1996).

[261] Berggren, K. K., Dauler, E. A., Kerman, A. J., Nam, S.-W. & Rosenberg, D. *Single-Photon Generation and Detection - Physics and Applications*, vol. 45 of *Experimental Methods in the Physical Sciences* (Elsevier, 2013).

[262] Enss, C. & McCammon, D. Physical Principles of Low Temperature Detectors: Ultimate Performance Limits and Current Detector Capabilities. *Journal of Low Temperature Physics* **151**, 5–24 (2008).

[263] Burstein, E., Langenberg, D. & Taylor, B. Superconductors as Quantum Detectors for Microwave and Sub-Millimeter-Wave Radiation. *Physical Review Letters* **6**, 92–94 (1961).

[264] Martin, D. D. E. & Verhoeve, P. Superconducting tunnel junctions. *ISSI Scientific Reports Series* **9**, 441–457 (2010).

[265] Zmuidzinas, J. Superconducting Microresonators: Physics and Applications. *Annual Review of Condensed Matter Physics* **3**, 169–214 (2012).

[266] Mazin, B. A. *et al.* ARCONS: A 2024 Pixel Optical through Near-IR Cryogenic Imaging Spectrophotometer. *arXiv preprint arXiv:1306.4674* 12 (2013).

[267] Day, P. K., LeDuc, H. G., Mazin, B. A., Vayonakis, A. & Zmuidzinas, J. A broadband superconducting detector suitable for use in large arrays. *Nature* **425**, 817–21 (2003).

[268] Gao, J. *et al.* A titanium-nitride near-infrared kinetic inductance photon-counting detector and its anomalous electrodynamics. *Applied Physics Letters* **101**, 142602 (2012).

[269] Natarajan, C. M., Tanner, M. G. & Hadfield, R. H. Superconducting nanowire single-photon detectors: physics and applications. *Superconductor Science and Technology* **25**, 063001 (2012).

[270] Dauler, E. A. *et al.* Photon-number-resolution with sub-30-ps timing using multi-element superconducting nanowire single photon detectors. *Journal of Modern Optics* **56**, 364–373 (2009).

[271] Rosenberg, D., Kerman, A. J., Molnar, R. J. & Dauler, E. A. High-speed and high-efficiency superconducting nanowire single photon detector array. *Optics express* **21**, 1440–7 (2013).

[272] Verma, V. B., Marsili, F., Harrington, S., Lita, A. E., Mirin, R. P. & Nam, S. W. A three-dimensional, polarization-insensitive superconducting nanowire avalanche photodetector. *Applied Physics Letters* **101**, 251114 (2012).

[273] Verma, V. B. *et al.* High-efficiency WSi superconducting nanowire single-photon detectors operating at 2.5 K. *arXiv preprint arXiv:1406.1810* (2014).

[274] Schuck, C., Pernice, W., Minaeva, O., Li, M., Gol'tsman, G., Sergienko, A. & Tang, H. Matrix of integrated superconducting single photon detectors with high timing resolution. *IEEE Transactions on Applied Superconductivity* **23**, 2201007 (2013).

[275] Natarajan, C. M. *et al.* Operating quantum waveguide circuits with superconducting single-photon detectors. *Applied Physics Letters* **96**, 211101 (2010).

[276] Irwin, K. D. An application of electrothermal feedback for high resolution cryogenic particle detection. *Applied Physics Letters* **66**, 1998 (1995).

[277] Burney, J. *Transition-edge sensor imaging arrays for astrophysics applications.* Ph.D. thesis, Stanford University (2007).

[278] Irwin, K. & Hilton, G. Transition-edge sensors. *Cryogenic Particle Detection* **149**, 0–67 (2005).

[279] Miller, A. J. *Development of a broadband optical spectrophotometer using superconducting transition-edge sensors.* Ph.D. thesis, Stanford University (2001).

[280] Cabrera, B., Clarke, R., Miller, A., Nam, S., Romani, R., Saab, T. & Young, B. Cryogenic detectors based on superconducting transition-edge sensors for time-energy-resolved single-photon counters and for dark matter searches. *Physica B: Condensed Matter* **280**, 509–514 (2000).

[281] Cabrera, B., Clarke, R. M., Colling, P., Miller, a. J., Nam, S. & Romani, R. W. Detection of single infrared, optical, and ultraviolet photons using superconducting transition edge sensors. *Applied Physics Letters* **73**, 735 (1998).

[282] Irwin, K. *Phonon-mdeiated particle detection using superconducting tungsten transition-edge sensors.* Ph.D. thesis, Stanford University (1995).

[283] Di Giuseppe, G. *et al.* Direct observation of photon pairs at a single output port of a beam-splitter interferometer. *Physical Review A* **68**, 063817 (2003).

[284] Lita, A. E., Miller, A. J. & Nam, S. W. Counting near-infrared single-photons with 95% efficiency. *Optics Express* **16**, 3032 (2008).

[285] Fukuda, D. *et al.* Titanium-based transition-edge photon number resolving detector with 98% detection efficiency with index-matched small-gap fiber coupling. *Optics express* **19**, 870–5 (2011).

[286] Smith, D. H. *et al.* Conclusive quantum steering with superconducting transition-edge sensors. *Nature communications* **3**, 625 (2012).

[287] Christensen, B. G. *et al.* Detection-Loophole-Free Test of Quantum Nonlocality, and Applications. *Physical Review Letters* **111**, 130406 (2013).

[288] Giustina, M. *et al.* Bell violation using entangled photons without the fair-sampling assumption. *Nature* **497**, 227–30 (2013).

[289] Lamas-Linares, A. *et al.* Nanosecond-scale timing jitter for single photon detection in transition edge sensors. *Applied Physics Letters* **102**, 231117 (2013).

[290] Tosi, A., Della Frera, A., Shehata, A. B. & Scarcella, C. Fully programmable single-photon detection module for InGaAs/InP single-photon avalanche diodes with clean and sub-nanosecond gating transitions. *The Review of scientific instruments* **83**, 013104 (2012).

[291] Miller, A. J., Nam, S. W., Martinis, J. M. & Sergienko, A. V. Demonstration of a low-noise near-infrared photon counter with multiphoton discrimination. *Applied Physics Letters* **83**, 791 (2003).

[292] Calkins, B., Lita, A. E., Fox, A. E. & Woo Nam, S. Faster recovery time of a hot-electron transition-edge sensor by use of normal metal heat-sinks. *Applied Physics Letters* **99**, 241114 (2011).

[293] Fukuda, D., Fujii, G. & Numata, T. Titanium Superconducting Photon-Number-Resolving Detector. *IEEE Transactions on Applied Superconductivity* **21**, 241–245 (2011).

[294] Marsili, F. *et al.* Detecting single infrared photons with 93% system efficiency. *Nature Photonics* **7**, 210–214 (2013).

[295] Simon, R., Burns, M. J., Colclough, M. S., Zaharchuk, G. & Cantor, R. *Mr SQUID User's Guide* (STAR Cryoelectronics, Santa Fe, NM, 2004).

[296] Welty, R. & Martinis, J. A series array of dc SQUIDs. *IEEE Transactions on Magnetics* **27**, 2924–2926 (1991).

[297] Nam, S. W. *Development of phonon-mediated cryogenic particle detectors.* Ph.D. thesis, Stanford University (1998).

[298] Drung, D., Assmann, C., Beyer, J., Kirste, A., Peters, M., Ruede, F. & Schurig, T. Highly sensitive and easy-to-use SQUID sensors. *IEEE Transactions on Appiled Superconductivity* **17**, 699–704 (2007).

[299] Alberto, D. *et al.* Optical Transition-Edge Sensors Single Photon Pulse Analysis. *Applied Superconductivity, IEEE Transactions on* **21** (2011).

[300] Namekata, N., Takahashi, Y., Fujii, G., Fukuda, D., Kurimura, S. & Inoue, S. Non-Gaussian operation based on photon subtraction using a photon-number-resolving detector at a telecommunications wavelength. *Nature Photonics* **4**, 655–660 (2010).

[301] Levine, Z. H., Gerrits, T., Migdall, A. L., Samarov, D. V., Calkins, B., Lita, A. E. & Nam, S. W. Algorithm for finding clusters with a known distribution and its application to photon-number resolution using a superconducting transition-edge sensor. *Journal of the Optical Society of America B* **29**, 2066 (2012).

[302] Fixsen, D., Moseley, S., Cabrera, B. & Figueroa-Feliciano, E. Pulse estimation in nonlinear detectors with nonstationary noise. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **520**, 555–558 (2004).

[303] Figueroa-Feliciano, E. & Cabrera, B. Optimal filter analysis of energy-dependent pulse shapes and its application to TES detectors. *Nuclear Instruments and Methods in Physics Research A* **444**, 453–456 (2000).

[304] Vayshenker, I. *et al.* Trilateral optical powermeter comparison between NIST, NMIJ/AIST, and METAS. *Applied Optics* **46**, 643 (2007).

[305] Krystek, M. & Anton, M. A weighted total least-squares algorithm for fitting a straight line. *Measurement Science and Technology* **18**, 3438–3442 (2007).

[306] DeCusatis, C. *Handbook of Fiber Optic Data Communication* (Academic Press, 2013), 4th edn.

[307] Gerrits, T. *et al.* On-chip, photon-number-resolving, telecommunication-band detectors for scalable photonic information processing. *Physical Review A* **84**, 1–4 (2011).

[308] Calkins, B. *et al.* High quantum-efficiency photon-number-resolving detector for photonic on-chip information processing. *Optics Express* **21**, 22657–22670 (2013).

[309] Sprengers, J. P. *et al.* Waveguide superconducting single-photon detectors for integrated quantum photonic circuits. *Applied Physics Letters* **99**, 181110 (2011).

[310] Schuck, C., Pernice, W. H. P. & Tang, H. X. NbTiN superconducting nanowire detectors for visible and telecom wavelengths single photon counting on Si3N4 photonic circuits. *Applied Physics Letters* **102**, 051101 (2013).

[311] Schuck, C., Pernice, W. H. P. & Tang, H. X. Waveguide integrated low noise NbTiN nanowire single-photon detectors with milli-Hz dark count rate. *Scientific reports* **3**, 1893 (2013).

[312] Reithmaier, G. *et al.* On-chip time resolved detection of quantum dot emission using integrated superconducting single photon detectors. *Scientific reports* **3**, 1901 (2013).

[313] Najafi, F. *et al.* On-Chip Detection of Entangled Photons by Scalable Integration of Single-Photon Detectors. *Arxiv preprint arXiv:1405.4244* (2014).

[314] Sahin, D. *et al.* Waveguide photon-number-resolving detectors for quantum photonic integrated circuits. *Applied Physics Letters* **103**, 111116 (2013).

[315] Lita, A. E., Calkins, B., Pellouchoud, L. A., Miller, A. J. & Nam, S. W. Superconducting transition-edge sensors optimized for high-efficiency photon-number resolving detectors. In *SPIE Symposium on SPIE Defense, Security, and Sensing* (2010).

[316] Lita, A., Rosenberg, D., Nam, S., Miller, A., Balzar, D., Kaatz, L. & Schwall, R. Tuning of Tungsten Thin Film Superconducting Transition Temperature for Fabrication of Photon Number Resolving Detectors. *IEEE Transactions on Appiled Superconductivity* **15**, 3528–3531 (2005).

[317] Photon Design Ltd. Fimmwave. URL `http://www.photond.com/products/fimmwave.htm`.

[318] Sudbo, A. S. Film mode matching: a versatile numerical method for vector mode field calculations in dielectric waveguides. *Pure and Applied Optics: Journal of the European Optical Society Part A* **2**, 211–233 (1993).

[319] Gallagher, D. F. G. & Felici, T. P. Eigenmode expansion methods for simulation of optical propagation in photonics: pros and cons. In Sidorin, Y. S. & Tervonen, A. (eds.) *Integrated Optics: Devices, Materials, and Technologies VII*, vol. 4987, 69–82 (2003).

[320] Gallagher, D. Photonic CAD matures. *IEEE LEOS NewsLetter* (2008).

[321] Heeres, R. W., Dorenbos, S. N., Koene, B., Solomon, G. S., Kouwenhoven, L. P. & Zwiller, V. On-chip single plasmon detection. *Nano letters* **10**, 661–4 (2010).

[322] Barnes, W. L., Dereux, A. & Ebbesen, T. W. Surface plasmon subwavelength optics. *Nature* **424**, 824–830 (2003).

[323] Cabrera, B. Introduction to TES Physics. *Journal of Low Temperature Physics* **151**, 82–93 (2008).

[324] Cabrera, B. Design considerations for TES and QET sensors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **444**, 304–307 (2000).

[325] Tanzilli, S., Martin, A., Kaiser, F., De Micheli, M., Alibart, O. & Ostrowsky, D. On the genesis and evolution of Integrated Quantum Optics. *Laser & Photonics Reviews* **6** (2011).

[326] Walmsley, I. & Raymer, M. Toward Quantum-Information Processing with Photons. *Science* **307**, 1733 (2005).

[327] Qiu, F., Spring, A. M., Yu, F., Aoki, I., Otomo, A. & Yokoyama, S. Thin TiO2 core and electro-optic polymer cladding waveguide modulators. *Applied Physics Letters* **102**, 233504 (2013).

[328] Bonneau, D. *et al.* Fast Path and Polarization Manipulation of Telecom Wavelength Single Photons in Lithium Niobate Waveguide Devices. *Physical Review Letters* **108**, 053601 (2012).

[329] Jennewein, T., Barbieri, M. & White, A. G. Single-photon device requirements for operating linear optics quantum computing outside the post-selection basis. *Journal of Modern Optics* **58**, 276–287 (2011).

[330] Christ, A. & Silberhorn, C. Limits on the deterministic creation of pure single-photon states using parametric down-conversion. *Physical Review A* **85**, 023829 (2012).

[331] Migdall, A., Branning, D. & Castelletto, S. Tailoring single-photon and multiphoton probabilities of a single-photon on-demand source. *Physical Review A* **66**, 053805 (2002).

[332] Shapiro, J. H. & Wong, F. N. On-demand single-photon generation using a modular array of parametric downconverters with electro-optic polarization controls. *Optics Letters* **32**, 2698 (2007).

[333] Collins, M. J. *et al.* Integrated spatial multiplexing of heralded single-photon sources. *Nature communications* **4**, 2582 (2013).

[334] Spring, J. B. *et al.* Quantum interference of multiple on-chip heralded sources of pure single photons. In *Research in Optical Sciences: Quantum Information and Measurement (QIM)*, QW1B.6 (OSA, Berlin, 2014).

[335] Holland, M. & Burnett, K. Interferometric Detection of Optical Phase Shifts at the Heisenberg Limit. *Physical review letters* **71**, 1355–1358 (1993).

[336] Humphreys, P. C., Kolthammer, W. S., Nunn, J., Barbieri, M., Datta, A. & Walmsley, I. A. Continuous-Variable Quantum Computing in Optical Time-Frequency Modes using Quantum Memories. *Phys. Rev. Lett.* **113**, 130502 (2014).

[337] Nunn, J. *et al.* Enhancing Multiphoton Rates with Quantum Memories. *Physical Review Letters* **110**, 133601 (2013).

[338] Michelberger, P. S. *et al.* Interfacing GHz-bandwidth heralded single photons with a room-temperature Raman quantum memory. *arXiv preprint arXiv:1405.1470* (2014).

[339] Sprague, M. R. *et al.* Broadband single-photon-level memory in a hollow-core photonic crystal fibre. *Nature Photonics* **8**, 287–291 (2014).

[340] Miller, A., Lita, A., Calkins, B. & Vayshenker, I. Compact cryogenic self-aligning fiber-to-detector coupling with losses below one percent. *Optics express* **19**, 9102–9110 (2011).

[341] Miller, A., Lita, A., Rosenberg, D., Gruber, S. & Nam, S. Superconducting photon number resolving detectors: performance and promise. *Proceedings of the 8th International Conference on Quantum Measurement and Computing* (2007).

Appendix A

Extended quantum teleportation results

The full results from the on-chip quantum teleportation experiment are given in section 4.7.2 as the plot in figure 4.14. The exact numerical values used in that plot are given in table A.1 below. In addition, the full set of reconstructed densitiy matrices for all four BSM outcomes are plotted in figure A.1 and figure A.2.

Bell state outcome	Measured fidelities: Input state $\|V'\rangle$	Input state $\|D'\rangle$	Input state $\|L'\rangle$	Average	Predicted Fidelity
$\|00\rangle$	0.73 ± 0.05	0.78 ± 0.05	0.94 ± 0.03	0.82 ± 0.03	0.83 ± 0.09
$\|01\rangle$	0.91 ± 0.03	0.84 ± 0.04	0.92 ± 0.10	0.89 ± 0.04	0.84 ± 0.09
$\|10\rangle$	0.78 ± 0.05	0.74 ± 0.05	0.88 ± 0.07	0.80 ± 0.03	0.78 ± 0.10
$\|11\rangle$	0.83 ± 0.06	0.85 ± 0.05	0.62 ± 0.10	0.77 ± 0.04	0.84 ± 0.09

Table A.1: Results of on-chip quantum teleportation. The measured fidelities of each teleported input state are listed four each of the four Bell state measurement outcomes. The average over the three linearly independent states is also given to compare to the expected fidelities based on the calculations given in section 4.6. Errors in the experimentally measured fidelities are calculated using a Monte-Carlo method taking account of Poissonian counting statistics and uncertainties in the characterised measurement operators and input states. The predicted fidelities take account of imperfect circuit fabrication, higher order photon emission and residual photon distinguishability. The errors in these simulated fidelities are estimated by a Monte-Carlo method, taking account of the uncertainity in the characterised circuit elements and imperfect photon source parameters.

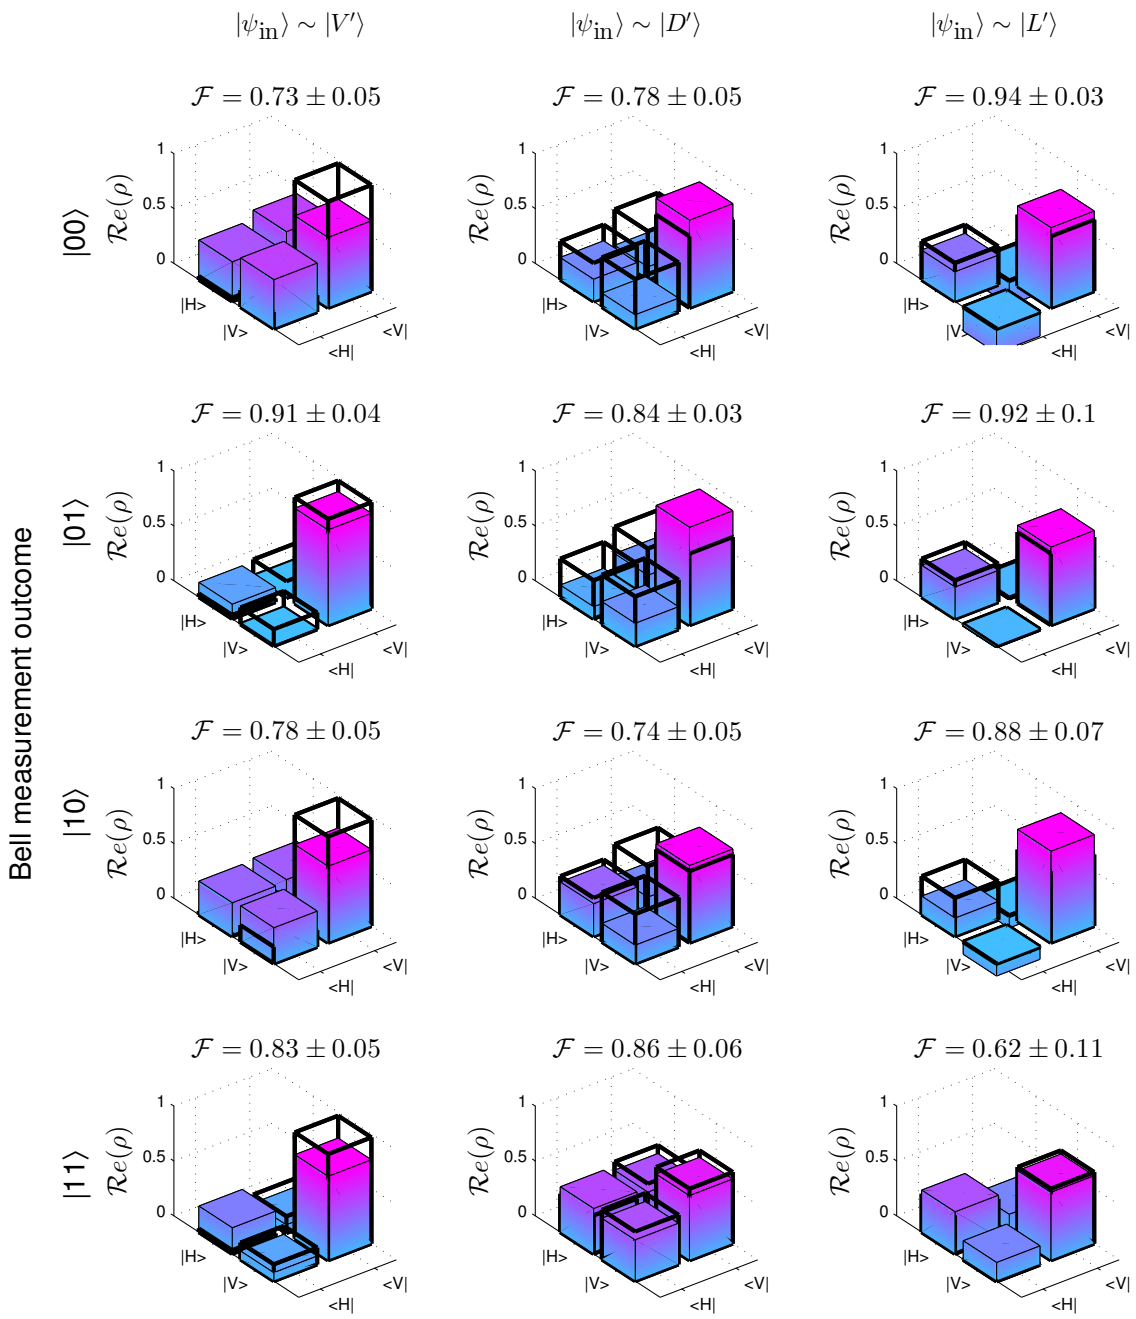


Figure A.1: Real-parts of the reconstructed density matrices. The three columns represent the three linearly independent states which were teleported and each of the four rows corresponds to one of the possible BSM outcomes. The initial states on Q1 are shown as black wire bars with the reconstructed state on Q3 (after applying the correct unitary rotation) are shown as the coloured bars. Errors on the fidelity are calculated using a Monte-Carlo method taking account of Poissonian counting statistics and uncertainties on the measurement operators and input states prepared.

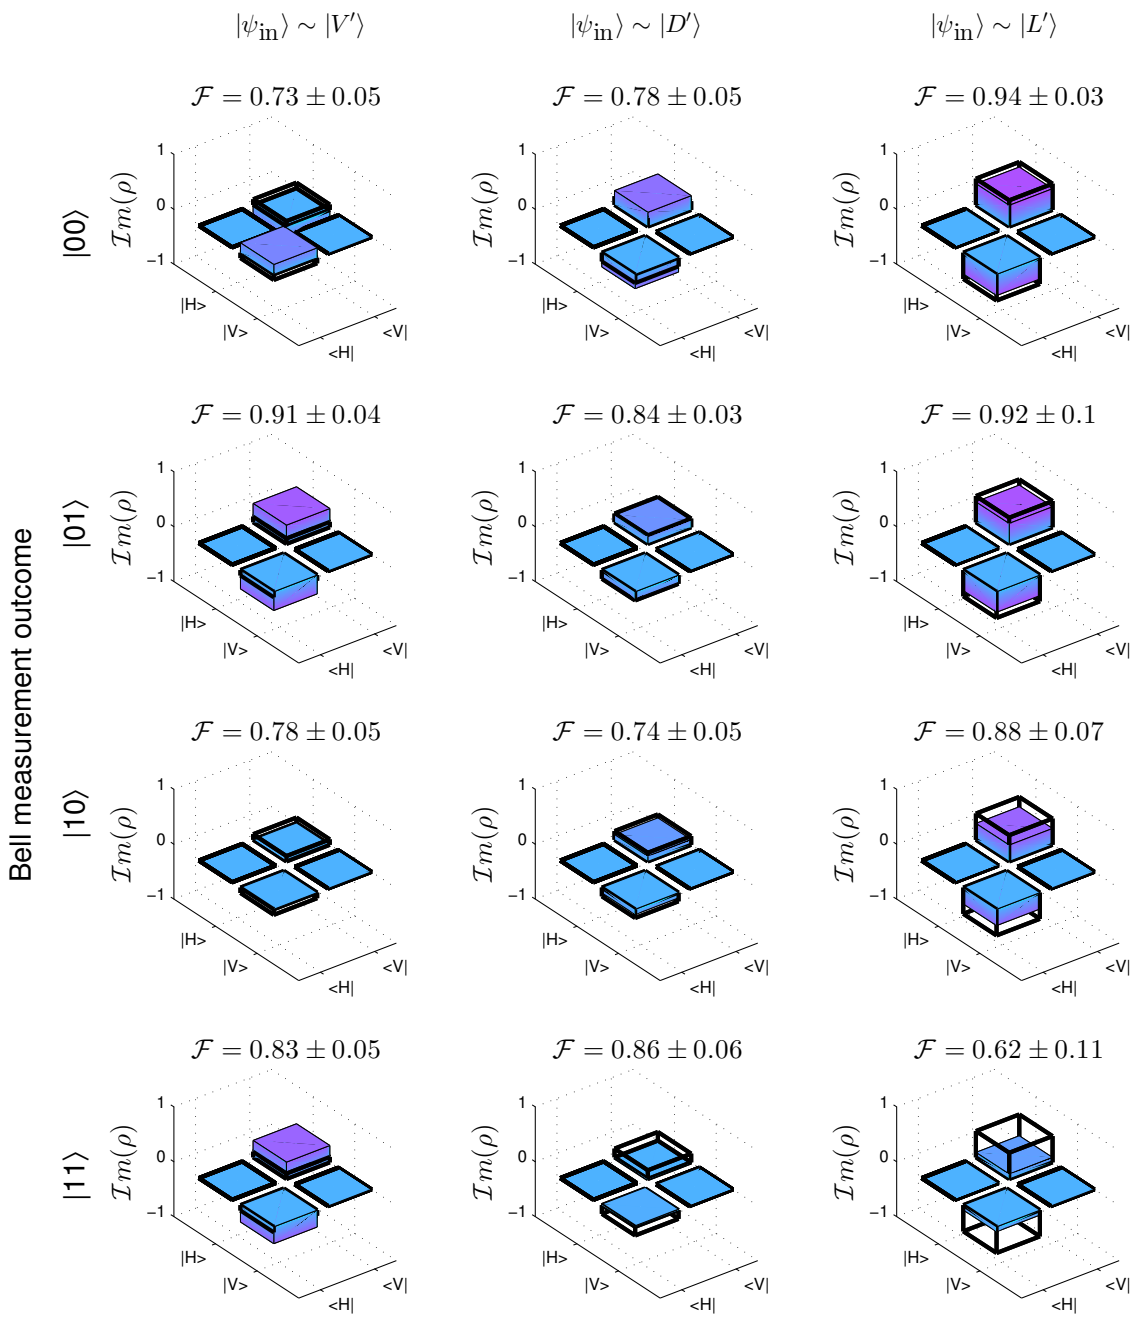


Figure A.2: Imaginary-parts of the reconstructed density matrices. Plots are similar to those described in the caption to figure A.1.

Appendix B

List of symbols used in description of the TES

The list below provides a convenient reference for the various symbols used during chapter 5.

Symbol	Definition
C_e	electronic heat capacity of the TES
g	thermal conductance of the weak thermal link between the TES and the substrate
κ	thermal coupling constant between the TES and the substrate
α	logarithmic 'steepness' of the superconducting transition region
$\xi_{\mathbf{in}}$	input current sensitivity of the SQUID - $\xi_{\mathrm{in}} = dI/d\Phi$
k_B	Boltzmann's contant
n	exponent characterising the thermal conductivity between the TES and the substrate
ϵ	fraction of photon energy which is transferred to the TES electron system
ω	angular frequency of the signal

τ	the natural thermal recovery time of the TES - $\tau = C_e/g$
τ_{etf}	recovery time of the TES in the presence of electro-thermal feedback
T	temperature of the TES electronic system
T_0	quiescent steady-state temperature of the TES
T_S	temperature of the substrate
T_c	superconducting transition temperature of the TES
R	resistance of the TES
R_0	quiescent steady-state resistance of the TES
E_γ	incident photon energy
$E_{\mathbf{sat}}$	saturation energy of the TES
ΔE	energy resolution of the TES
V	voltage across the TES
$v_{\mathbf{djn}}$	TES detector Johnson voltage noise
$v_{\mathbf{cjn}}$	bias circuit Johnson voltage noise
$i_{\mathbf{sqn}}$	current noise due to the SQUID
$i_{\mathbf{bias}}$	current noise due to the bias circuit
p_γ	input photon noise
$p_{\mathbf{TFN}}$	thermodynamic fluctuation noise from power fluctuations across the weak thermal link

Appendix C

Installation of TES detectors

C.1 Cryogenic system

In order to achieve good energy resolution and low noise operation, TES detectors operate with transition temperatures of around 100 mK. Reaching these extreme temperatures is possible by commercial apparatus, but is neither easy nor cheap! This project began with the acquisition of a used dilution refrigerator from Sussex University. Through an existing collaboration with the group of Sae Woo Nam at NIST, we were able to acquire TES detectors. Fast-forward nearly two years and a lab had been found, a pit had been dug, a pump shed built, a gantry erected, vacuum pumping lines installed, a 35 litre dump of Helium-3 bolted to the wall and our very own second-hand dilution fridge sitting in the middle of the room waiting to go!

The dilution fridge itself is an Oxford Instruments Kelvinox-MX400 which is designed for interchangeable experiments using a removable 'experimental insert'. It is specified to achieve a base temperature of 12 mK and has a cooling power of 400 μW at 100 mK. Dilution refrigerators are continuous-cycle systems which use a pumped ^{4}He system to help cool a

mixture of ^{3}He and ^{4}He to below 1K. Once this mixture cools below $\sim$ 700 mK the two helium isotopes separate into a ^{3}He-dilute phase and a ^{3}He-rich phase. The reason dilution fridges are able to get so cold is because there remains a finite fraction of ^{3}He in the dilute phase ($\sim$ 6%) even at zero temperature. Within the dilute phase, the ^{3}He acts like a vapour in an 'atmosphere' of ^{4}He above a reservoir of liquid ^{3}He. Pumping on the dilute ^{3}He phase allows ^{3}He atoms from the rich phase to 'evaporate' into the dilute phase, taking energy from the surroundings as they do so and thus cooling the mixture. However, unlike all other known gases, the effective vapour pressure of the system remains finite at all temperatures allowing the system to cool indefinitely in the abscence of other heat loads. Once the cooling power provided by the dilution fridge equals the heat load coupled down to the cold stage, the system reaches a steady-state base temperature. When installing an experiment on the fridge, care must be taken not to increase this heat load too much so that the base temperature remains low enough to operate the detectors.

Our fridge was first commissioned in July 2013 with help from an ICE Oxford engineer. Having spent two years finding the lab and getting it ready, we still had no guarantee whatsoever that the system would actually work. Therefore, there was cause for much celebration when, after a week of leak-checking, leak-fixing and a crash-course in cryogenic system operation, we saw a base temperature of a little under 20 mK! Unfortunately, the mood was soon punctured by a major leak in the system crashing the fridge and contaminating the mixture. Several months were wasted (or was it just a valuable learning experience?) repairing the system, adding numerous valves in some choice locations and perfecting our cool down routine such that we can now be confident of going from room temperature to 50 mK in a little under 3 days. Once cold, so long as it is regularly topped up with liquid helium,

it will stay cold for weeks or months on end.

C.2 Wiring up the experimental insert

C.2.1 Vacuum feedthroughs

A brief look at figure 5.3 shows that each detector requires three low-temperature bias lines. We use cables consisting of 12 twisted pairs woven into a single loom, using constantan (Co) for the SQUID and Feedback bias lines and beryllium-copper (BeCu) for the TES bias lines to reduce electrical resistance. Each loom was soldered onto a hermetically sealed Fischer connector to act as a vacuum feedthrough and carefully heat sunk at each stage by wrapping the loom around a copper bobbin.

Traditional fibre feedthroughs for cryostat systems rely on hermetically sealed SMA or FC-type connectors mounted to a bulkhead. However, these connectors introduce unnecessary loss into the optical system and take up significant room. Instead, we made our own fibre bundle feedthrough by drilling a small hole ($\sim$ ⌀3mm) through a blank KF16 flange. We assembled a bundle of 16 fibres (8$\times$ Corning SMF28 + 8$\times$ Nufern 780HP both with 250 μm polyimide buffer) and held them in place using some 5 minute epoxy (Hardman Double/Bubble). A tight vacuum seal was made by wicking an ultra-low viscosity epoxy (Stycast 1266) along the fibres into the bored flange (see figure C.2(a)). Around 2 metres of fibre was left on the room-temperature side for subsequent splicing onto connectors or directly to an experiment, stored in a coil in a lightproof box attached to the fridge (figure C.2(b)). The fibre bundle was installed through the line-of-sight port to the mixing chamber and heat-sunk at each stage using copper tubes (inner diameter $\sim$ 1.5 mm) clamped to each plate with a Swagelok compression fitting (see figure C.2(d)). At each stage, the

copper tube was offset slightly to reduce the amount of black body radiation coupled from the room temperature surfaces. At the mixing chamber the fibres were fusion spliced onto AR-coated fiber ferrules which were inserted into self-aligning TES modules C.2(f).

C.2.2 SQUIDs and breakout board

To minimise the heat load on the mixing chamber it is common to place the SQUID electronics board at the 1K stage where the available cooling power is substantially higher. However, by mounting the SQUIDs at the mixing chamber, we can achieve lower noise and make the NbTi wiring to the TES shorter, minimising the timing impact of stray inductance. This comes at the cost of increased heat load so a new SQUID electronics board was designed which could feasibly fit into the tight space available at the 1K stage if the heat load turned out to be too much (see figure C.2(c))). The board was designed to hold 8 SQUID chips with four Micro-D connectors breaking out to the three incoming bias lines and one for the niobium titanium loom (NbTi) connected to the TES (see figure C.1). The board also provides the parallel bias circuit for the TES using cold bias resistors (Yageo 311-001 NTCT-NO) with a nominal resistance of 10 mΩ at 50 mK.

The SQUID chips were fabricated at NIST using a gradiometric geometry designed to reduce the need for any magnetic shielding[298]. The SQUID chips actually contain 100 SQUIDs in series which provides roughly 100x amplification of the voltage signal and improves impedance matching to the room-temperature electronics[296]. Two SQUID chips were glued onto the electronics board and wire-bonded to the contact pads. The board was mounted on an OFHC copper plate and hung from the mixing chamber. With the fibres, bias wiring and electronics board installed at the mixing chamber we achieved a base

temperature of 50 mK. Based on the quoted cooling power of the fridge and knowing each SQUID only dissipates 20 nW of power, this suggests the wiring is the dominant heat load. Since this wiring is already sufficient to run 12 detectors we can be somewhat confident that installing at least this number of detector channels should be possible on our fridge.

C.2.3 Detector box and blackbody radiation

We obtained two tungsten TES detectors from NIST which were fabricated using their standard design - a $20\,\mu\text{m} \times 20\,\mu\text{m} \times 20\,\text{nm}$ square of tungsten placed in an optical cavity with a gold backing mirror to enhance optical absorption into the detector (recall figure 5.2) [284]. One detector was optimised for dual-band operation at 1550 nm and 1330 nm whilst the other had a cavity designed for high efficiency operation at 850 nm. The response of these optical cavities is quite broad with high efficiency operation expected for around 100 nm bandwidth around the central wavelength. The TES is deposited on a 'keyhole'-shaped silicon substrate which is secured within a ceramic ferule as part of a 'self-aligning' mounting system [340]. A fibre ferule with anti-reflection coating is inserted into the sleeve until the tip touches the detector. The $5 - 10\,\mu\text{m}$ sized fibre core is aligned to within better than $3\,\mu\text{m}$ of the centre of the detector which ensures all of the light exiting the fibre will strike the tungsten and remains well-aligned through multiple cool-down cycles (see figure C.2(f)).

Because the TES is a thermal device it is very important to consider any source of heat which could adversely affect device operation. A simple calculation shows that black-body radiation from room-temperature surfaces would contribute a few hundred nanowatts of power to a $20\,\mu\text{m} \times 20\,\mu\text{m}$ device [279] - something we obviously need to avoid! Therefore, we must be very careful to seal up any holes with direct line-of-sight from the top of the

fridge to the mixing chamber. Aluminium tape is very useful here in covering up any gaps between the stages. To make sure the detectors are well protected from this radiation, they are mounted in a sealed box made from gold-plated (with a non-magnetic underclad) OFHC copper. The box was designed to store 6 self-aligned TES modules while allowing the user to remove the experimental insert without having to break fibre splices. By hanging a rod from the mixing chamber we have room to mount 3 of these boxes on our system giving the option of installing up to 18 detectors. Once the detectors are clamped in this box we seal the edges using aluminium tape to make a tight black body radiation shield. The detector boxes are mounted at the lowest portion of the system such that heat dissipated from other components, such the fibres, wiring or SQUIDs is conducted upward toward the mixing chamber and does not couple into the TESs (see figure C.2(e) - item 20).

Having gone to some effort to reduce the impact of black-body radiation from the fridge, one may wonder if connecting a fibre directly to the TES with the other end exposed to room-temperature black-body is such a good idea! Estimates for a 200 μm core -diameter multi-mode fibre suggest that when the reduced transmission of the fibre and reduced absorption of the W are taken into account for the long black-body peak wavelengths, we would expect almost 60 k cts/s due to this radiation [279]. Fortunately, the situation is much improved for single-mode fibre since the peak of the room-temperature black body (about 10 μm) is much longer than the fibre cutoff wavelength. There is negligible transmission through the fibre at this wavelength, and a calculation for the effect of the residual black body tail (below 2 μm) shows we should expect only around 40 cts/s from a 300 K fibre if no other filters are present [341].

C.3 Electronics and readout

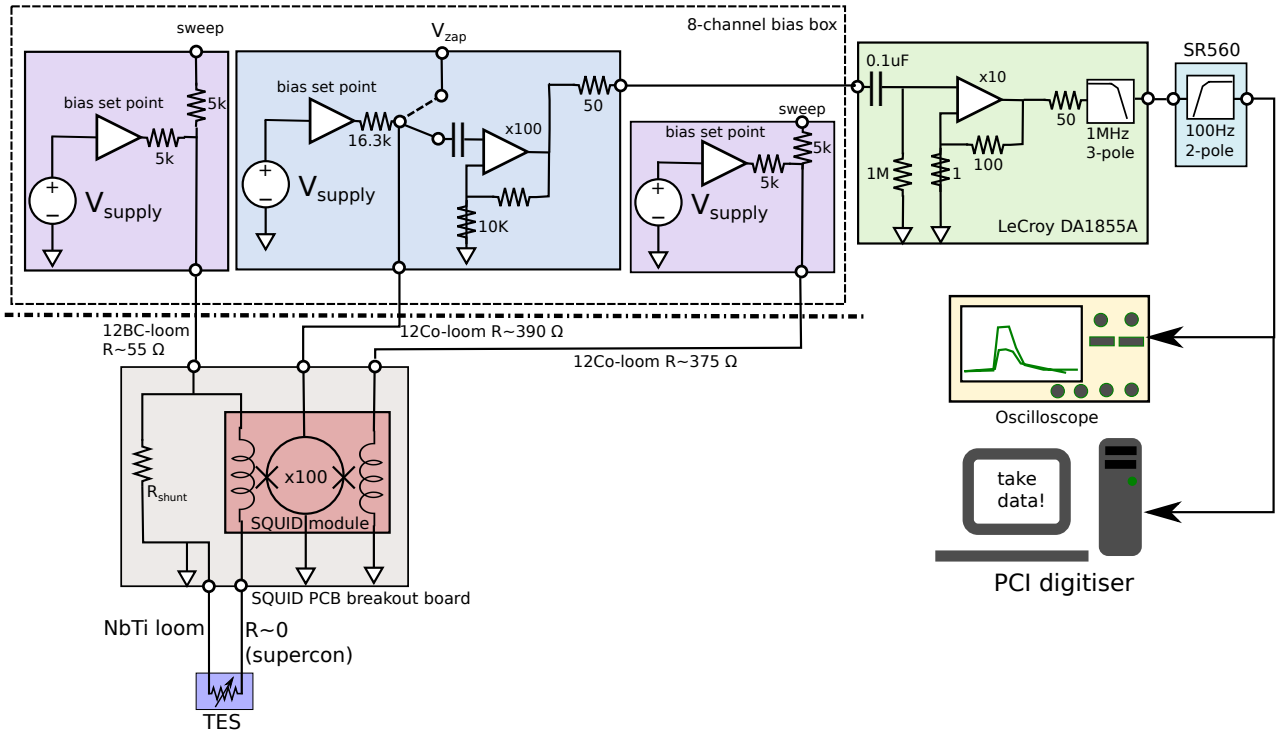


Figure C.1: Detail of data acquisition electronics. An 8-channel bias box provides the three control circuits needed to operate each TES. The TES and feedback bias lines can be connected to an external driving source to sweep the bias current which is used to tune the SQUID response. The SQUID is read-out through a bias-tee circuit before passing through a series of amplifiers and filters. Approximate resistances of all the bias lines are shown including the measured resistances of all of the twisted pair looms inside the fridge.

The room-temperature data acquisition electronics used to read-out the TES detectors are summarised in figure C.1. An 8-channel bias box is connected to a low-noise voltage source (Agilent HP3620A) and used to set the current bias for the three circuits through an internal series resistance (5 kΩ). For the TES and feedback bias lines, there is an option to sweep the applied bias currents using a function generator.

The SQUID dc-bias control is connected through a bias tee and here is an option to connect the bias line to a high current source to 'zap' the SQUID by momentarily driving it normal so that any trapped magnetic flux may escape the ring. 'Zapping' the SQUIDs in

this way helps to improve the modulation depth we can achieve. Immediately after the bias tee is a low-noise FET pre-amplifier with 20 dB gain. The SQUID output ($\sim 10\,\mathrm{mW}$) tends to sit upon a large DC offset ($\sim 1\,\mathrm{V}$) which is either removed via AC-coupling or subtracted using a differential amplifier and a precision voltage source. The signal is further amplified (x10) by another low-noise ($\sim 1\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$) amplifier followed by a low-pass frequency filter with a 3-pole roll-off at 1 MHz. Looking back at figure 5.5, we expect the signal-to-noise ratio to fall below 1 at around this frequency and this filter helps eliminate extra Johnson noise and any excess RF pickup which may be aliased down to our frequencies. Similarly a separate high-pass filter at 300 Hz is used to eliminate most of the $1/f$ noise contribution as well as the dominant 50 Hz signal and its first few harmonics. The signal is then split between an oscilloscope and a data acquisition card for saving signal pulses to disk. We use an Alazar ATS-9440 PCI express digitiser to perform fast analogue-to-digital conversion and stream the voltage signal to memory. This DAQ card is triggered by the laser diode controller running the experiment, leading to a series of stored voltage traces which must be analysed later in post-processing. Unlike conventional detectors, the output signal from a TES consists of a very large amount of data. Each photon trace is around $8\,\mu$s long and is sampled at 125 MS/s with 14-bit resolution leading to each photon signal generating several kilobytes worth of data.

C.3.1 Noise reduction

It is worth mentioning some of the effort which must be spent hunting down and eliminating excess noise sources which can ruin the detection signal. The installation of ceramic breaks on the vaccum pump lines during lab construction proved to be a wise decision. The effect

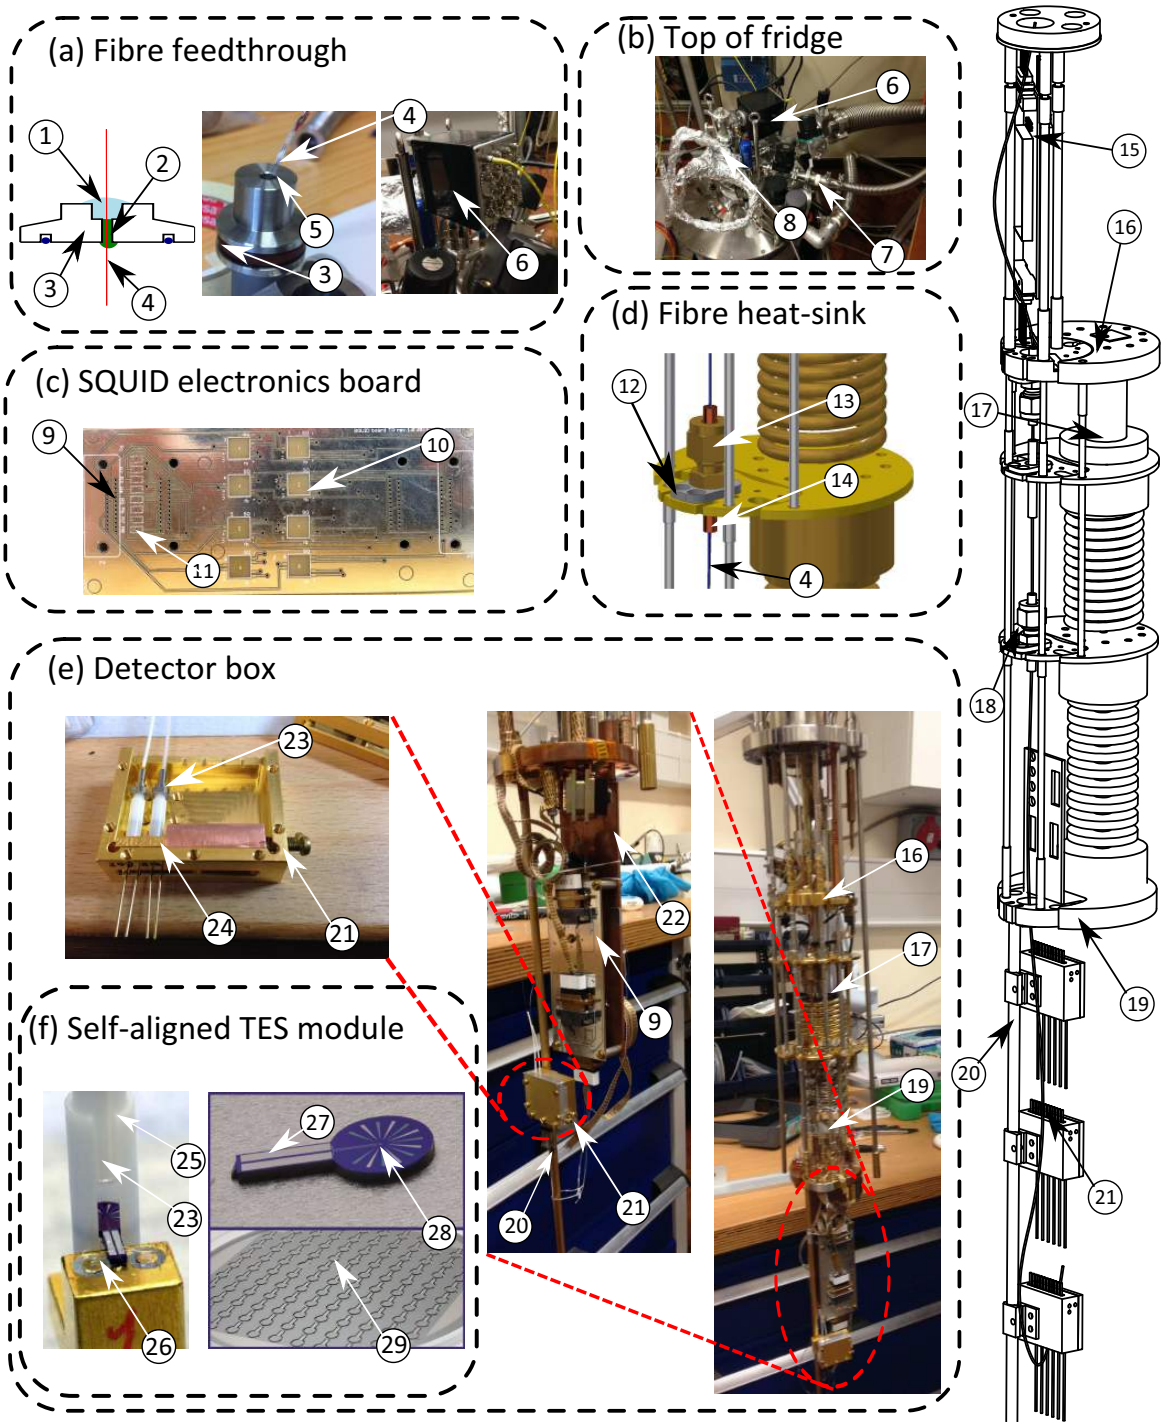


Figure C.2: Installing detectors in the fridge. (1) Five-minute epoxy (2) Stycast 1266 (3) KF16 unbored stub flange (4) Fibre bundle (5) 3 mm bored hole (6) Fibre breakout box (7) Ceramic vacuum break (8) Bias lines for TES, SQUID and Feedback coil (9) SQUID electronics board (10) SQUID chip mounting pad (11) Bias resistor mounting pad (12) Swagelok bracket (13) Swagelok copper clamp (14) Copper fibre heat sink (15) Alternative SQUID board location at 1K stage (16) 1K stage (17) Still (18) Fibre heat sink (19) Mixing Chamber (20) Detector box mounting rod (21) Detector Box (22) SQUID board heat sink / mount (23) AR-coated fibre ferule (24) Self-aligned TES module (25) Ceramic alignment sleeve (26) Wire bonds (27) TES Bonding pads (28) Tungsten TES (29) Etched silicon wafer of many 'keyhole' design TES chips.

of shorting across these breaks is startling - the TES signal almost completely vanishes due to the extra noise introduced.

Making sure each of the components, including the fridge, share a common ground potential is also important for minimising signal interference. This needs to be done carefully to avoid introducing unnecessary ground loops, and usually involves spending several hours connecting various components together with ground straps until the noise is minimised. All of our electronics are housed within a 19" rack which also acts as a basic Faraday shield for external RF noise. Judicious use of aluminium foil was also found to be quite effective - especially at the top of the fridge by the Fischer connectors. It seems that sometimes the rubber o-ring seals on the vacuum feedthroughs break the continuous ground - tightly wrapping foil to bridge these gaps seemed to help. The cables running from the fridge to the bias box were fitted with EMI ferrite filters to reduce RF interference and we were careful to switch off mobile phones whilst using the detectors.

Of course more can always be done, but at some point if the signal looks OK then it is better to just get on and take some data. However, over the course of using the detectors, extra noise-reducing measures were identified. New SQUID chips have been ordered with higher input sensitivity and with in-built chokes on the input coil to reduce high frequency noise. A mu-metal shield has been ordered to fully enclose the SQUID electronics and TES at the mixing chamber to reduce the effect of stray magnetic fields. We are constructing new electrical feedthroughs which will allow better continuous shielding from the fridge to the bias box and also moving the bias box much closer to the fridge itself. Triple-shielded coaxial cable is being installed and additional ceramic breaks added to isolate the fridge from the gas handling system. In the future we may consider adding vibration isolation to

reduce microphonic noise but we are not sure how big an effect this will be. Although this appears like a lot of effort (it is!), if it is done correctly, it only needs to be done once and we are aiming to build up a facility which will be able to be left running and provide multiple high-efficiency detection channels to the next generation of quantum optics experiments performed in our lab.