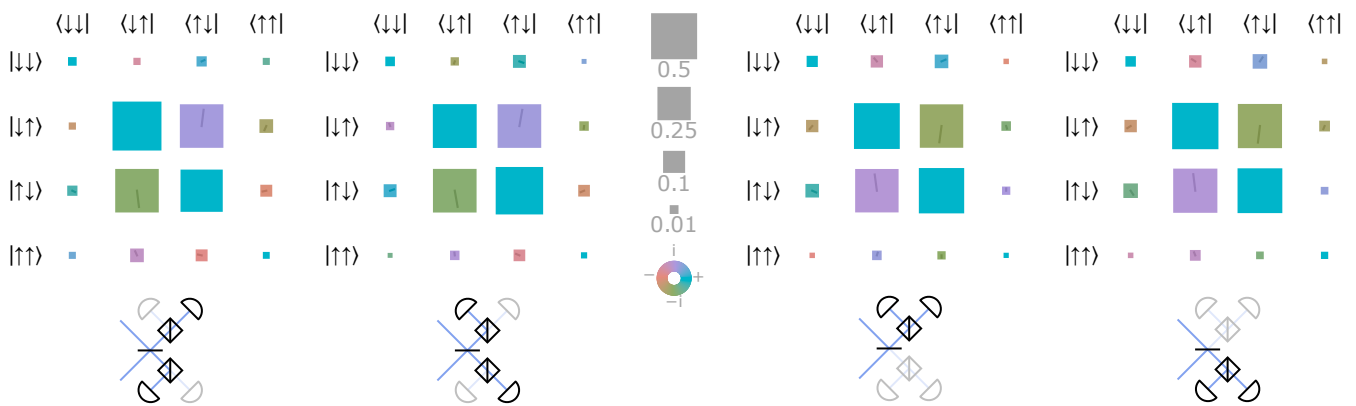


UNIVERSITY OF OXFORD
DEPARTMENT OF PHYSICS

Entanglement between nodes of a quantum network

Laurent J. Stephenson



A thesis submitted for the degree of
Doctor of Philosophy

Wolfson College
Trinity term, 2019

Abstract

Entanglement between nodes of a quantum network

Laurent J. Stephenson

A thesis submitted for the degree of Doctor of Philosophy

Trinity term 2019

Wolfson College, Oxford

Quantum devices are rapidly gaining momentum as a technology that will induce a paradigm shift in computing, communication and cryptography. Trapped ion qubits are one of the leading candidates for implementing a quantum computer, having previously demonstrated all of the required criteria. Local gate fidelities between ions exceed those for all other platforms, but the total number of ions in a trap is limited by unavoidable issues – one possibility for scaling the ion trap quantum processor is to create entanglement between ions in separate traps via single photons. The work in this thesis demonstrates the generation of remote entanglement between ions at high fidelity and rate, paving the way towards protocols using multiple entangled pairs for computations.

We describe the construction from scratch of a twin-trap apparatus designed to entangle separated ions. We demonstrate the entanglement of strontium ions in traps separated by ~ 2 m by swapping entanglement from single photons emitted by the $^{88}\text{Sr}^+$ ions. A novel photon collection geometry is used, maximising the entanglement between the ion and photon without impeding optical access for standard ion trap laser beam geometries. The Bell state fidelity of the ion-photon state, which is itself a valuable entanglement resource for blind computation, is at least 97.70(12)% in either trap, with entangled pairs detected at a rate of at least $3.98 \times 10^3 \text{ s}^{-1}$. The remote ion-ion Bell state fidelity is 94.0(5)%, and is generated at a rate of 182 s^{-1} , representing the highest fidelity remote entanglement reported in ions by a large margin, at a rate more than an order of magnitude faster than previous experiments.

The two identical trap systems are designed with capabilities beyond those demonstrated here. We can co-trap $^{88}\text{Sr}^+$ and $^{43}\text{Ca}^+$ in a microfabricated trap suitable for local ion transport operations, opening up the possibility of performing distillation of remote entanglement in the same apparatus. Entanglement distillation can be used to make remote entanglement with similar fidelity to local operations, and would represent a significant step towards a fully scalable ion trap quantum computer architecture.

Acknowledgements

My greatest thanks must go to my supervisor, Prof. David Lucas; his engaging undergraduate atomic physics tutorials were one of the main reasons I chose to continue my studies, and years of experience always provided sage advice on the occasions he poked his head into the lab.

Dr. Chris Ballance led (and still leads) the twin-trap experiment on which I have spent the past few years, and I count myself incredibly lucky to have worked with him. His enthusiasm for science is truly infectious¹ and somehow he found the patience to teach me most of what I know about ion trapping. The experiment was very much a team effort, and I owe an enormous amount to the other members past and present, namely Dr. Tim Ballance, David Nadlinger, Beth Nichol, Dr. Shuoming An and Peter Drmota, without whom this would not have been possible – I remain astounded by what we managed to achieve together in such a short time. I'd like to thank Dr. Joe Goodwin for always asking questions that left me utterly flummoxed, inevitably leading to both of us scratching our heads for the next hour, and for being able to reach the elusive fibre coupling Zen state of mind. A special mention goes to Rustin Nourshargh, who, aside from providing ample amounts of levity during his time as a Master's student in the group, also proofread this thesis from cover to cover². The wider ion trap group proved to be a wonderful bunch to work (and occasionally drink) alongside, so thank you to Dr. Tom Harty, Dr. Martin Sepiol, Dr. Vera Schäfer, Dr. James Tarlton, Keshav Thirumalai, Jochen Wolf, Amy Hughes, Clemens Löschnauer, Marius Weber, Shaobo Gao, Will Hughes, and Prof. Andrew Steane.

My friends and family have supported me throughout this endeavour, and deserve thanks for enduring frequent dinner table physics explanations without (visible) signs of indifference. My grandmother always used to say, “Do your research!” – I suppose research is never complete, but I can definitively say that I have now done *some* research.

¹Despite a near identical acknowledgment in Vera Schäfer's thesis, this comment was written independently!

²And shall hereafter be blamed for all errata...

Finally, I gratefully acknowledge funding from EPSRC and from the Wolfson Harrison scholarship.

Contents

1	Introduction	1
1.1	Quantum information processing and networking	3
1.2	Trapped ions for quantum computing	5
1.3	Thesis outline	9
2	The choice of ions	11
2.1	Strontium	12
2.1.1	Photoionisation	14
2.1.2	Doppler cooling	15
2.1.3	State preparation	16
2.1.4	Single qubit manipulations and readout	18
2.1.5	Entangled photon emission and collection	20
2.2	Calcium	23
3	Dual ion trap apparatus	24
3.1	The ion trap	26
3.1.1	Trap design	26
3.1.2	Electrical connections and trap mounting	27
3.1.3	DC supply and filtering	33
3.1.4	Trap RF supply	34
3.2	Vacuum system	36
3.2.1	Assembly and baking	38
3.2.2	Pumping	41
3.2.3	Atomic ovens	43
3.2.4	Vacuum mounting structure	46
3.3	Magnetic field coils	46
3.4	Experimental Control	47
3.5	Laser systems	49
3.5.1	Diode lasers	50
3.5.2	Ti:Sapphire lasers	51
3.5.3	AOM switching and frequency tuning	52

3.6	Trap beam paths	54
3.6.1	Fixed final mirrors	59
3.6.2	Alignment onto ions	61
3.7	Rear imaging system	64
4	Photon collection and entangling apparatus	68
4.1	Photon collection	69
4.1.1	Fibre coupling procedure	69
4.2	Photonic Bell state analyser	76
4.2.1	Beam splitters	79
4.2.2	Waveplates	80
4.2.3	Optical fibres	82
4.3	Experimental control	83
4.3.1	The “entangler core”	83
4.3.2	AOM switching latency	86
5	Ion state manipulation and readout	89
5.1	State preparation	89
5.2	Readout in the computational basis	91
5.3	Readout in other bases	94
5.4	Errors in preparation and readout	95
6	Ion-photon entanglement	96
6.1	Collecting photons from ions	98
6.2	Ion-photon entanglement	103
6.3	Error sources	110
6.3.1	Ion dephasing	111
6.3.2	Ion rotation errors	112
6.3.3	Laser leakages	113
6.3.4	Photon polarisation errors	114
6.3.5	Summary of errors	115
7	Ion-ion entanglement	116
7.1	Beamsplitter theory	119
7.2	Entangled state phase	127
7.3	Indistinguishability criteria	133
7.3.1	Temporal matching	134
7.3.2	Field equalisation	135
7.4	Fibre birefringence compensation	138
7.5	Remote entanglement	139
7.6	Error sources	141

8	Conclusions	147
8.1	Comparison to other experiments with ions	149
8.1.1	Ion-photon entanglement	149
8.1.2	Ion-ion entanglement	151
8.2	Comparison with alternative platforms	151
8.3	Outlook	152

List of Figures

2.1	Diagram of the $^{88}\text{Sr}^+$ level structure	13
2.2	$^{88}\text{Sr}^+$ two-step photoionisation scheme	14
2.3	$^{88}\text{Sr}^+$ state preparation	17
2.4	$^{88}\text{Sr}^+$ qubit manipulations and readout	18
2.5	$^{88}\text{Sr}^+$ photon emission and collection	20
2.6	$^{43}\text{Ca}^+$ level structure and Raman transitions	22
3.1	Experiment overview block diagram	25
3.2	Photograph of the HOA2 trap	27
3.3	Schematic of the HOA2 electrode layout	28
3.4	DC connection test circuit	32
3.5	Trap RF chain	37
3.6	Vacuum system internals	39
3.7	422 nm AOM setup	55
3.8	674 nm AOM layout	56
3.9	Schematic diagram of $^{88}\text{Sr}^+$ beam paths	57
3.10	Photograph of the fixed final mirrors	60
3.11	Confocal imaging system	63
3.12	Rear side imaging system	66
4.1	Imaging ion light into a fibre	71
4.2	Point spread function images and fits	73
4.3	Fitted Zernike coefficients of point spread function	74
4.4	Photonic Bell state analyser layout	78
5.1	Bright and dark readout histograms	92
5.2	Rabi oscillations of the optical qubit	94
6.1	Three polarisation-encoded photon emission schemes	100
6.2	Fast excitation pulse sequence	104
6.3	Fast preparation of $S_{1/2, m = -1/2}$	106
6.4	Pulsed excitation to $P_{1/2, m = +1/2}$	107

6.5	Ion-photon density matrices	109
7.1	Two types of photonic Bell state analyser	124
7.2	Energy levels of an arbitrary system for entanglement swapping	128
7.3	Photon arrival time distribution and fit	136
7.4	Ion-ion density matrices for each herald pattern	140

List of Tables

3.1	Toroidal resonator parameters	35
3.2	Magnetic field coil specifications	47
4.1	Non-polarising beam splitter properties	79
4.2	Polarising beam splitter extinction ratios	80
4.3	Bell state analyser waveplate properties	82
4.4	Switching latencies of time-critical beams	88
7.1	Ion-ion error budget	146

Chapter 1

Introduction

The history of computing stretches far further back than one might naively expect. As soon as civilisation developed, humans encountered problems that were intractable without some form of assistance; how much grain to levy as taxes from each farmer for example. Tools were invented to aid these calculations, the first being the Sumerian abacus. More and more sophisticated ways of problem solving developed – geared mechanisms for calculating the positions of celestial bodies such as the Antikythera mechanism, dated from c. 125 BCE, and much later astronomical clocks in China, Europe and the Middle East.

These machines had very specific tasks, and were not remotely versatile. The language for describing versatile computing machines came with Alonzo Church and Alan Turing in the 1930s, while simultaneously Claude Shannon proved that electric switching circuits could implement Boolean algebra, paving the way for modern computing. The rate of development accelerated with the invention of the transistor in 1959, and today the classical computer

is embedded in virtually every aspect of our daily lives, from the smartphone in our pocket to the supercomputer calculating the weather forecast.

The early 20th century was a frenzy of scientific discovery in other areas too – the idea that energy might be quantised was put forward by Planck to solve the “ultraviolet catastrophe”¹, where classical statistical mechanics predicts absurdly that an ideal black body would radiate an unbounded amount of energy. This, along with Einstein’s explanation of the photoelectric effect requiring the existence of finite “energy quanta” (which would eventually become known as photons), led to the birth of quantum mechanics as we now know it.

This might appear to have nothing at all to do with computing, and indeed quantum mechanics and computing remained apart for several decades, until two observations were made in the early 1980s: (1) we can simulate any classical Turing machine of finite length with a quantum Hamiltonian [Ben80] and (2) we *cannot* simulate quantum systems efficiently with classical computing [Fey82]. This led to Deutsch describing a quantum Turing machine that he called a ‘universal quantum computer’ [Deu85] that *can* simulate any physical system, including other quantum systems, and so the link between quantum mechanics and computing was made inextricable.

Demonstrating that a quantum computer can solve an abstract class of problems more efficiently than a classical computer is certainly interesting enough for scientists to pursue the goal of creating one, but without more concrete examples it can be difficult to explain the significance of this to those outside the field. Perhaps the two most famous algorithms that explicitly

¹A name I’ve always found amusingly dramatic

show a quantum computer outperforming a classical computer are Grover’s quantum database search [Gro97] and Shor’s algorithm [Sho97] for factorising large numbers. The asymmetry in computational difficulty between the multiplication (easy) and the factorisation (hard) of large numbers underpins the usefulness of encryption techniques like RSA – thus the person wanting to find out the decryption key needs a new abacus.

1.1 Quantum information processing and networking

At its core, our quantum computer consists of *qubits*. A qubit is any quantum system that we can control; we label two states of our choice as $|0\rangle$ and $|1\rangle$ corresponding to the 0 and 1 of a classical bit. The qubits are controlled with quantum gates in exact analogy to Boolean logic gates, and by extension we are motivated to try to implement a “universal gate set” for quantum operations, much like a composition of NAND gates can be used to implement any Boolean gate.

A universal gate set may be composed of single qubit operations and one ‘entangling’ gate [Kit97], but other alternatives are possible. If we can initialise an entangled state of two qubits, then with the addition of single qubit gates and a method of measuring the entangled state (a Bell basis measurement) we can construct a different universal set [GC99]. Thus it appears that creating entangled two qubit states is an important process in building a quantum computer. Furthermore, if these two qubits can be

entangled at a distance from each other, we can begin to explore new ideas in quantum information processing.

Qubits entangled at a distance can be used to implement a quantum network; one application may be to generate an encryption key [Eke91] with a fundamentally different approach to the classical example of RSA given earlier, which is reliant on the properties of quantum mechanics rather than the computability of the factorisation problem. Classical communications over long distance are subject to attenuation of signals, requiring amplifiers at intermediate points – quantum communication channels are also subject to losses and decoherence that degrade the link quality with increasing distance. We might attempt to copy the state of the qubit that we are trying to share, but this is not possible, assuming the quantum no-cloning theorem [WZ82, Par70] holds, so how do we extend our link? What we can do is *teleport* a quantum state [Ben95], consuming an entangled Bell pair in the process. In this way, we can teleport the qubit’s state across many shorter links to get to its eventual destination – the devices in between being called ‘quantum repeaters’. Shared entangled pairs can also be used to implement a ‘blind’ quantum computation, where a client need only make measurements on their qubit while a quantum computing server performs computations on the other qubit [MF13], and in this way the client can execute an algorithm without the server gaining knowledge of what the algorithm, its inputs or its outputs were.

We have established that quantum computers and qubits entangled at a distance may well be very useful, without specifying in detail what a qubit is, beyond calling it a ‘quantum system’. Luckily, nature is quantum, and so

many candidates present themselves as potential qubits – from quantum dots, Josephson junctions and nitrogen vacancy centres in diamond to photons, atoms and trapped ions to name but a few.

1.2 Trapped ions for quantum computing

Trapped ions were first proposed as an implementation for a quantum computer in 1995 [CZ95]. A qubit is encoded in two energy states of the ion; generally one is chosen in the ground level, while the other may be in a metastable excited level, resulting in an optical frequency qubit, or a different state of the ground level, giving a microwave or radio frequency qubit.

A list of the criteria that any qubit must satisfy in order to be suitable for quantum information processing was compiled in 2000 by DiVincenzo [DiV00]. The list consists of five criteria pertaining to a localised quantum computer, as well as a further two related to quantum networking. The first five have been demonstrated to exceptional standards with trapped ions:

- *A scalable physical system with well characterized qubits:* Atoms of the same element and isotope are, as far as we know, completely identical, and so any states which we choose our qubit states will always have the same energy. Ionising the atom allows us to confine it in a quasi-electrostatic trap [Pau90]. Electromagnetic fields such as lasers or microwaves can interact with the electronic state of the ion, allowing us to control the state. While scaling to a useful number of qubits remains to be shown, multiple proposals exist for scaling ion

trap systems to large numbers of physical qubits; quantum charge-coupled devices (QCCDs) [KMW02] are one proposal, while photonic interconnects [MRR⁺14] are the main topic of this thesis. See also reviews such as [MK13, BKM16] for further discussions on scalability.

- *The ability to initialize the state of the qubits to a simple fiducial state, such as $|000\dots\rangle$:* The interaction of the electrons with electromagnetic fields allows us to pump population towards a single state [WBID80].
- *Long relevant decoherence times, much longer than the gate operation time:* The ion is trapped in a vacuum; there are few collisions with background gases and the intensity of blackbody radiation at room temperature is far below that needed to drive transitions², and so is well isolated from the environment. Coherence times for ground state ion qubits are limited principally by technical noise sources causing dephasing – coherence times exceeding 10 min have been observed in hyperfine qubits [WUZ⁺17], while magnetically sensitive Zeeman qubits have coherence times up to 300 ms [RSK⁺16], compared to typical gate times of 1–100 μ s.
- *A “universal” set of quantum gates:* Quantum gates with high fidelity have been realised in single and multiple ions in many forms: between ground-state qubits and optical qubits; with lasers [BHL⁺16, GTL⁺16] or microwaves in near [HSA⁺16] and far field; single qubit gates in isolation or individually addressed within an ion string [ANK⁺15]; two

²Although it does perturb the energy levels enough to be a concern in ion clocks [AKR⁺18]

qubit gates between ions of the same or different [TGL⁺15] species.

- *A qubit-specific measurement capability:* Proof of the above results relies on one crucial task: reading out the state of the qubit. In ion traps this is generally accomplished by state-dependent fluorescence, such that on application of the cooling lasers, the atom will scatter many photons from one state and not the other, which we can detect to make a measurement [MSW⁺08].

While these first five criteria appear to have been soundly fulfilled by trapped ions, we note that while excellent two-qubit gates have been realised in systems with a few ions in a single harmonic trap, this method is not easily scalable to arbitrarily large numbers of ions. Adding more ions to the trap increases the number of motional modes of the ion string used to implement gates; more modes means more off-resonant excitation of unwanted modes, leading either to slower gates or larger errors. We note promising work on longer ion chains does exist [LWL⁺19], but this is still experimentally limited to some tens of ions.

The natural development then is to have many traps, each with only a few ions, but then we need to be able to link them. One option is to physically move ions between traps [WMI⁺98], where segmented electrodes can be used to modify the potential such that ions can be split off from a chain and pushed around in a moving potential well. However, in this case the distances traversed are fairly small; if many trap zones are to be used then the required density of laser beams (for addressing each zone) is high. The number of controllable electrodes also presents some difficulty, since each

needs an independent analogue voltage source.

In order to separate the traps further, we can try to use photons to link them [MK13, NFB14] – noting that this of course can apply to any qubit system that can be interfaced with photons. The state of a photon emitted by spontaneous decay is entangled with the ion it originated from; with two such ions and photons, we can perform entanglement swapping [Coe04, SI03] so that by making a Bell measurement on the photons, we can create an entangled pair of ions at a distance. The entangled ion-photon pair can also be useful in its own right – since measurements of photon polarisation qubit states are possible with only linear optics, a blind quantum computer of the type described in section 1.1 can be implemented.

Armed with this knowledge, we examine the final two DiVincenzo criteria:

- *The ability to interconvert stationary and flying qubits*
- *The ability to faithfully transmit flying qubits between specified locations*

Bi-directional interconversion between stationary ion qubits and flying photon qubits is possible with schemes emitting/absorbing a photon entangled with the ion's state [KES⁺16]; in this thesis we create entangled ion-photon pairs via spontaneous emission, and do not map the photon state directly onto another ion. The entangled ion-photon pairs are then used to create entangled pairs of ions at a distance. The second criterion can be fulfilled simply by transmitting photons through optical fibres.

True scalability of the ion trap quantum computer is likely to incorporate a hybrid system, with QCCD-type devices connected with photonic links [BKM16]. Even with noisy links, the entanglement can be *purified* by

consuming multiple low-quality entangled pairs to create a smaller number of higher-quality pairs [BBP⁺96]. Demonstrating this entanglement purification (also known as distillation) is the principle goal of the apparatus built for this thesis, towards which a photonic link is a key intermediate result.

1.3 Thesis outline

In this thesis, we present the implementation of photonic links between two ion traps. The work appearing here was published in *Physical Review Letters* on 16th March 2020 [SNN⁺20]. The thesis is structured as follows.

Chapter 2: The choice of ions gives an introduction to the strontium qubit used throughout, including the ion qubit manipulations and the production of entangled ion-photon pairs.

Chapters 3: Dual ion trap apparatus and 4: Photon collection and entangling apparatus detail the apparatus used in this experiment. Chapter 3 describes apparatus that is common to many ion trap experiments, while chapter 4 is concerned with the photon creation, collection and interference apparatus.

Chapter 5: Ion state manipulation and readout describes the implementation of the ion qubit operations introduced in chapter 2.

Chapter 6: Ion-photon entanglement presents the collection of entangled photons in more detail, expanding on the theory and showing the results of the ion-photon entanglement experiment, with a brief discussion of the errors present.

Chapter 7: Ion-ion entanglement showcases the main experimental

result, that is entanglement between separated ions using photons. We discuss the effects of imperfect photonic Bell measurements on the ion-ion entangled state along with other error sources.

Chapter 8: Conclusions summarises the results of the thesis, their relation to other ion-trap experiments and future prospects.

Chapter 2

The choice of ions

In this chapter we introduce the ion species – strontium – used throughout this work. Many references discuss both the suitability of ions in general¹ and the use of strontium [Ber02] in detail, so we aim only to briefly provide some context and introduce the notation used in the rest of the thesis. An overview of the techniques used to trap, cool and coherently manipulate strontium is given. The mechanism for producing photons entangled with the ion is presented, with detailed discussion deferred to chapter 6. We also give a passing mention to calcium, which will be the second species in this experiment. Calcium will be used for local operations and memory – the infrastructure for its trapping and manipulation has been built as part of this project.

¹See for example [Ste96] and many references therein.

2.1 Strontium

We have extolled the virtues of trapped ions for quantum operations, but thus far have been vague about the nature of the ions themselves. Typically, one of the alkali earth metals is chosen – the full s -orbital of the neutral atom means that the singly ionised form has a simple, hydrogen-like level structure which lends itself to laser cooling and quantum control. Several transition metals (cadmium, mercury, ytterbium, zinc) also share this property and are successfully employed in other groups.

Current proposals for entanglement distillation of remote Bell pairs require two ion species to be used [NBBB16]; one is used for the remote entangling via a photonic interconnect, and the other for high-fidelity local operations. Using two species allows the use of globally addressing beams, since the transitions of each species are far enough detuned from each other that a laser addressing one species does not disturb the quantum state of the other species. Calcium-43 has proven itself as an excellent performer for gate operations [BHL⁺16], and as a memory qubit with coherence times approaching a minute without the use of dynamic decoupling sequences [HAB⁺14], making it the logical choice for our local application qubit and leaving only the selection of the photonic interconnect species. It is desirable to choose two species that are close in mass, since the motional coupling is stronger, making sympathetic cooling and mixed-species quantum gates simpler. The distillation protocol also means that the qubit formed by the interconnect species is not intended for long term storage, so that this qubit only needs a coherence time long compared to the single gate needed to transfer its state

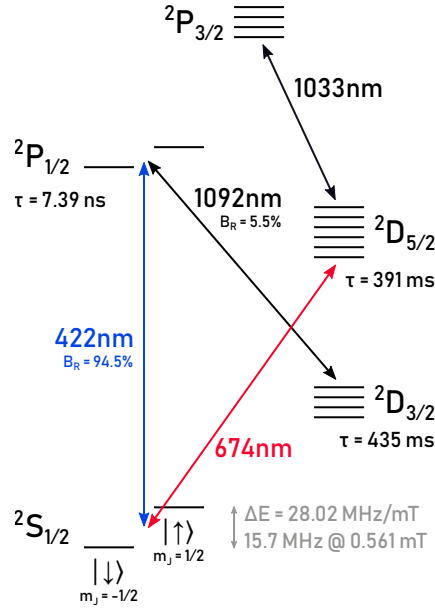


Figure 2.1: Diagram of the $^{88}\text{Sr}^+$ level structure. Relevant branching ratios, lifetimes and wavelengths of transitions are all shown (data collated from [LTGG16, LWGS05, MLN⁺99, San12]). Note that in this thesis a magnetic field of ~ 0.5 mT was used, so the transitions from $P_{1/2} \leftrightarrow S_{1/2}$ are not resolved, unlike at higher fields *e.g.* in [Sch18].

to one of the memory qubits, rather than compared to the whole computation. This allows to choose a species with no hyperfine structure and use a Zeeman qubit.

We choose Strontium-88 as the photonic interconnect species, which is used exclusively throughout this thesis. Strontium-88 is the most common isotope at 82.58 % [BW11] abundance, with highly enriched sources also available²

This even-even isotope has no nuclear spin, and thus no hyperfine struc-

²The radioactive isotope $^{89}\text{Sr}^+$ is used in bone cancer treatment, for delivering beta radiation to bony metastases by replacing calcium due to their similar chemical properties. $^{89}\text{Sr}^+$ is produced by neutron irradiation of $^{88}\text{Sr}^+$, hence highly enriched $^{88}\text{Sr}^+$ is readily available.




Figure 2.2: $^{88}\text{Sr}^+$ two-step photoionisation scheme. The 461 nm laser excites the atom to $5s5p 1P_1$, which is then ionised by the 378 nm laser.

ture; see figure 2.1. The ground level is split in the presence of a magnetic field, forming a Zeeman qubit. We denote the Zeeman qubit states $|\downarrow\rangle \equiv |S_{1/2, m = -1/2}\rangle$ and $|\uparrow\rangle \equiv |S_{1/2, m = +1/2}\rangle$. These are split by $\Delta E = g_J \mu_B B = 28.0 \text{ MHz mT}^{-1}$.

Low-lying metastable D levels provide both a shelf level for readout of the Zeeman qubit, as well as the possibility of an optical qubit. All of the transition wavelengths can be accessed with diode lasers, although in this experiment a mixture of diode and Ti:Sapphire lasers are used. The shortest wavelengths needed are in the visible spectrum³, so conventional optical fibres can be used without having to resort to photonic crystal fibres, and standard optics are available with good transmission properties.

2.1.1 Photoionisation

To load strontium ions into the trap, we first evaporate neutral atoms from a resistively heated oven. The neutral atom beam passes through the trap zone. A two-step photoionisation process, demonstrated first analogously for

³The one exception being the 378 nm laser used in photoionisation.

calcium [GRB⁺01], and subsequently in strontium [VCLB06, BLW⁺07], creates singly ionised strontium in the trap zone, where it is then Doppler cooled. The process uses diode lasers at 461 nm (to excite the $5s^2\ ^1S_0 \leftrightarrow 5s5p\ ^1P_1$ transition of neutral ^{88}Sr) and 378 nm (to excite to the continuum and ionise), as shown in figure 2.2. These wavelengths are easier to realise than the 293 nm light needed for the single wavelength ionisation process [BLW⁺07]. The longer wavelengths also cause less undesirable charging of the trap. The 378 nm laser has a shorter wavelength than is strictly necessary for strontium; this wavelength allows it to fulfil the same purpose for calcium photoionisation [LRH⁺04].

Photoionisation of a neutral atom beam from a heated oven is not the only way to load an ion trap; however, it has largely superseded electron bombardment as the latter requires a much higher atom flux [KHT⁺00] and charges the trap electrodes, increasing the variability of the micromotion compensation voltages [BLW⁺07]. Laser ablation [HGH⁺07] is another alternative, which may well be better suited to more compact experiments in the future. However, for the present, the chosen method offers tried-and-tested reliability and simplicity.

2.1.2 Doppler cooling

Doppler cooling in Strontium-88 is simple due to the lack of hyperfine structure; the $S_{1/2}$, $P_{1/2}$ and $D_{3/2}$ form a closed Λ -system (see figure 2.1). The 422 nm cooling beam is π -polarised, while the 1092 nm beam for clearing the $D_{3/2}$ level has a mixed σ and π polarisation. The Zeeman splitting of the $S_{1/2}$

level (15.7 MHz) is comparable to its linewidth (20.1 MHz) at the 0.561 mT magnetic fields used, so a single laser frequency at the 422 nm wavelength of the $S_{1/2} \leftrightarrow P_{1/2}$ transition suffices to address both ground states.

No resolved-sideband cooling is used in this thesis; the lowest temperature achievable with Doppler cooling is $\sim \Gamma/2$ [EMSKB03], where Γ here is the linewidth of the 422 nm transition, equal to 20.1 MHz⁴. For cooling close to this Doppler limit, the detuning of the 422 nm laser is set around 16 MHz red-detuned from the line centre⁵ and the intensity to a single saturation intensity. The detuning is found by maximising the fidelity of the optical qubit manipulations described in section 2.1.4. At the Doppler limit, and with typical trap secular frequencies of 1–5 MHz, the mean motional occupation number $\bar{n} \sim 2$ –8, substantially above the vibrational ground state.

During readout and experiment idle time, the power is increased and frequency red detuned to maximise fluorescence. The 1092 nm laser is tuned 10–20 MHz above the resonant frequency of the transition to avoid dark states caused by coherent population trapping, with a power equivalent to several hundred saturation intensities.

2.1.3 State preparation

The strontium qubit is prepared into $|\downarrow\rangle \equiv |S_{1/2, m = -1/2}\rangle$ by optically pumping with a σ^- -polarised 422 nm beam as shown in figure 2.3. At or below saturation of the σ^- transition, the fidelity of this operation is limited only

⁴Calculated from the value for the Einstein A coefficient given in Ref. [San12].

⁵The line centre is found by counting the fluorescence over 30 acquisition windows of 500 μ s each, short enough that the ion is not heated and showing the full Lorentzian line shape rather than a half-Lorentzian.

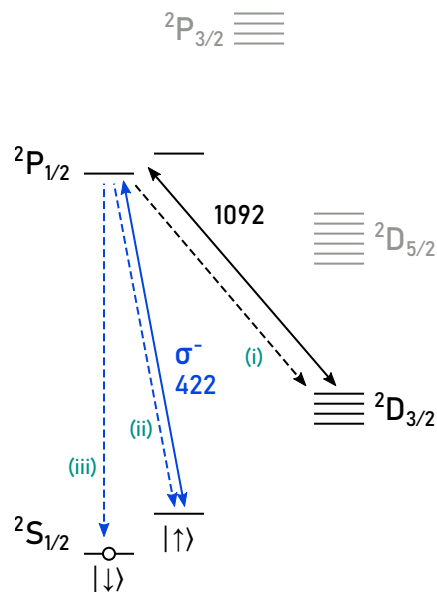


Figure 2.3: $^{88}\text{Sr}^+$ state preparation. The σ^- -polarised 422 nm beam excites population from $|\uparrow\rangle$ to $P_{1/2}$, where it can decay (i) to $D_{3/2}$, from which it is cleared by a 1092 nm beam; (ii) back to $|\uparrow\rangle$, where it is re-excited; or (iii) to $|\downarrow\rangle$. The population in $|\downarrow\rangle$ does not interact with the σ^- -polarised beam because there are no $\Delta m = -1$ transitions to $P_{1/2}$ from this state, and so the population is pumped toward this state.




Figure 2.4: $^{88}\text{Sr}^+$ (a) qubit manipulations and (b) readout. Population is mapped from $|\uparrow\rangle \rightarrow |D\rangle$ with a π -pulse on the “map” transition. Optical qubit rotations occur on the “main” transition between $|\downarrow\rangle \leftrightarrow |D\rangle$. Application of the Doppler beams causes population in $|\downarrow\rangle$ to fluoresce; $|D\rangle$ stays dark.

by the polarisation purity of the 422 nm beam, but if the power in this beam is increased above the saturation intensity, then the relative scattering rate of the impurities increases and the fidelity decreases correspondingly.

2.1.4 Single qubit manipulations and readout

To be able to use both the optical and Zeeman qubits, we must be able to map one of the Zeeman qubit states into a state in the $D_{5/2}$ manifold, as well as performing our main optical qubit operations between that D state and the other Zeeman qubit state.

Our choices of a suitable D state are dictated by the available beam propagation directions, with respect to both the trap axis and the applied magnetic field. The 674 nm beam must have some projection onto the axial direction to perform multi-qubit motional gates on the axial modes; the magnetic field direction is constrained because (a) it must be perpendicular to the imaging direction (see chapter 6) and (b) it is desirable to be able to have other beams both parallel and perpendicular to the field. The electric quadrupole transition $S_{1/2} \leftrightarrow D_{5/2}$ selection rules allow transitions with $|\Delta m| \leq 2$. We choose $|D\rangle \equiv |D_{5/2, m = -3/2}\rangle$ ⁶, since the relative transition strength from the two Zeeman qubit states can be tuned by adjusting polarisation rather than beam direction [Nad18]. The $|\Delta m| = 1$ transition between $|\downarrow\rangle \leftrightarrow |D\rangle$ is the less magnetically sensitive of the two, so we use these two levels as the optical qubit and call this the “main” optical qubit transition. The $|\Delta m| = 2$ transition between $|\uparrow\rangle \leftrightarrow |D\rangle$ is therefore used to transfer between the Zeeman and optical qubits, and is called the “map” transition.

The $D_{5/2}$ manifold is outside of the Doppler cooling cycle – if we have mapped population into the optical qubit, then all we need do to read out the qubit is apply the Doppler beams (see [MSW⁺08] for the exactly analogous readout of $^{40}\text{Ca}^+$). If the qubit was in state $|\downarrow\rangle$ (rather than $|D\rangle$), the ion will fluoresce. To read out the Zeeman qubit, we must first map population from $|\uparrow\rangle$ into $|D\rangle$, and then proceed as before. The readout fidelity of the Zeeman qubit will be directly linked to the fidelity of a π -pulse on the map transition to the optical qubit, since any population remaining in

⁶The notation $|D\rangle$ is potentially ambiguous in this thesis since we also discuss photon polarisations, and $|D\rangle$ traditionally denotes diagonal polarisation – the intended meaning should always be obvious from the context.

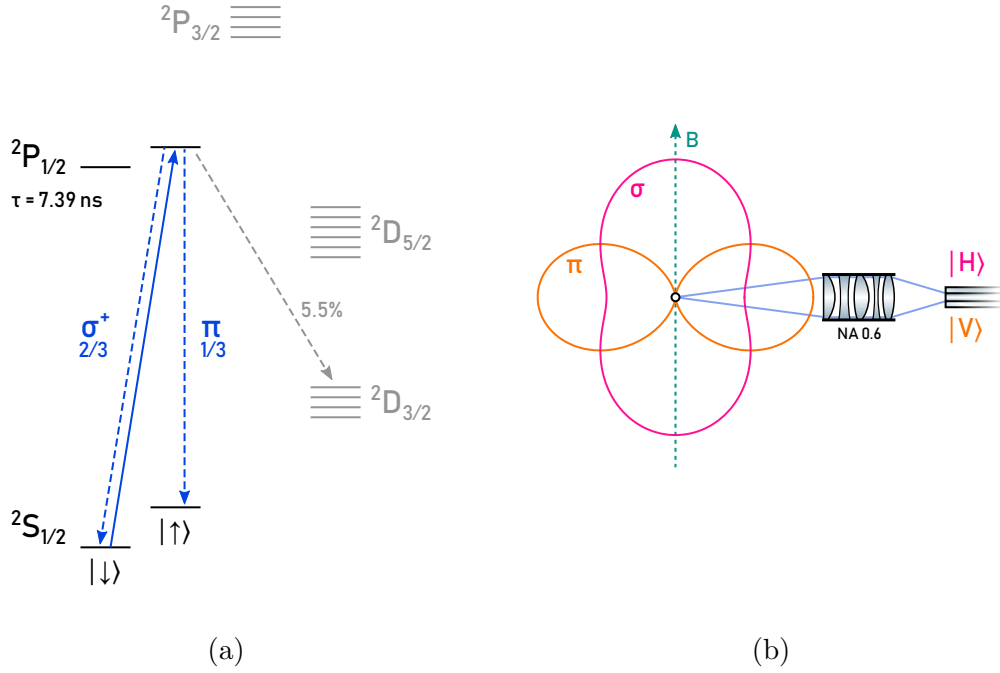


Figure 2.5: $^{88}\text{Sr}^+$ (a) photon emission and (b) photon collection perpendicular to the applied field, with the intensity of the emission patterns for each polarisation represented.

$|\uparrow\rangle$ will fluoresce. The optical qubit readout fidelity is limited fundamentally by the finite photon collection time and the 391(2) ms lifetime of the $D_{5/2}$ level [MSW⁺08, KGA⁺11], however, since this is much longer than the time typically taken to read out the qubit, it is not a practical limitation.

2.1.5 Entangled photon emission and collection

To produce photons entangled with the state of the ion, we use an ultra-fast, coherent pulse at 422 nm to excite the ion from $|\downarrow\rangle \rightarrow |P_{1/2,1/2}\rangle$. On spontaneous decay of the ion, the photon produced is entangled with the ion.

The level structure of Strontium-88 is particularly suitable for the pro-

duction of polarisation entangled photons – disregarding the 5.5% lost to the $D_{3/2}$, there are only two decay paths, directly into our Zeeman qubit, as shown in figure 2.5a. If we condition on the presence of a 422 nm photon, then with the Clebsch-Gordan coefficients of the decay the combined ion-photon state in free space is

$$|\psi\rangle = \sqrt{\frac{2}{3}} |\downarrow\rangle |\sigma^-\rangle + \sqrt{\frac{1}{3}} |\uparrow\rangle |\pi\rangle . \quad (2.1)$$

This state is not maximally entangled due to the imbalance in decay channels; its fidelity to the maximally entangled state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}\{|\downarrow\rangle |\sigma^-\rangle + |\uparrow\rangle |\pi\rangle\}$ is 97%.

The photons are collected perpendicular to the magnetic field into an optical fibre, shown schematically in figure 2.5b. This geometry was first used in free space with $^{111}\text{Cd}^+$ [MMBM04]. Ref. [Nad18] shows that this choice leads to collecting only the maximally entangled fraction of the ion-photon state, such that the state transmitted by the fibre is

$$|\psi\rangle = \frac{1}{\sqrt{2}} \{|\downarrow\rangle |\sigma^-\rangle + |\uparrow\rangle |\pi\rangle\} . \quad (2.2)$$

That we can counter the imbalance in overall emission with a suitable choice of observation axis is extremely useful; this is discussed in greater detail in chapter 6.




Figure 2.6: (a) The $^{43}\text{Ca}^+$ level structure, and (b) a Raman transition for qubit manipulation.

2.2 Calcium

Though not used in this work, the infrastructure to trap and manipulate Calcium-43 was built concurrently with the strontium infrastructure. Its level structure is shown in figure 2.6a. Doppler cooling is performed on the 397 nm and 866 nm transitions between $S_{1/2} \leftrightarrow P_{1/2} \leftrightarrow D_{3/2}$. State preparation is conducted by optical pumping with a circularly polarised 397 nm beam to $S_{1/2} |F = 4, m_F = 4\rangle$, followed by optional microwave pulses to prepare other $|F, m_F\rangle$ states.

The qubit is encoded in the hyperfine states of the $F = 3$ and $F = 4$ manifolds of the ground level – in intermediate strength applied magnetic fields the Zeeman and hyperfine splittings have comparable magnitude, and so some pairs of states have a vanishing first-order field dependence, forming ‘atomic clock’ qubits. Qubit manipulations can be achieved with either microwaves or Raman transitions with lasers (see figure 2.6b).

Chapter 3

Dual ion trap apparatus

The apparatus designed for the experiments in this thesis can be broadly divided into a “standard” ion trap apparatus, described in this chapter, and the photon collection and interference apparatus detailed in chapter 4. While we have called this a “standard” experiment, the level of complexity required to build an experiment consisting of two ion trap vacuum chambers, each able to trap two ion species and perform laser-based gates on them, as well as collecting single photons for remote entanglement experiments, is enormous. The two separate trap systems and all of their ancilliary subsystems are denoted (in time-honoured quantum information tradition) as Alice and Bob, by order of construction. An overview of the subsystems appears in figure 3.1; this also includes the photon apparatus to properly illustrate control flow.

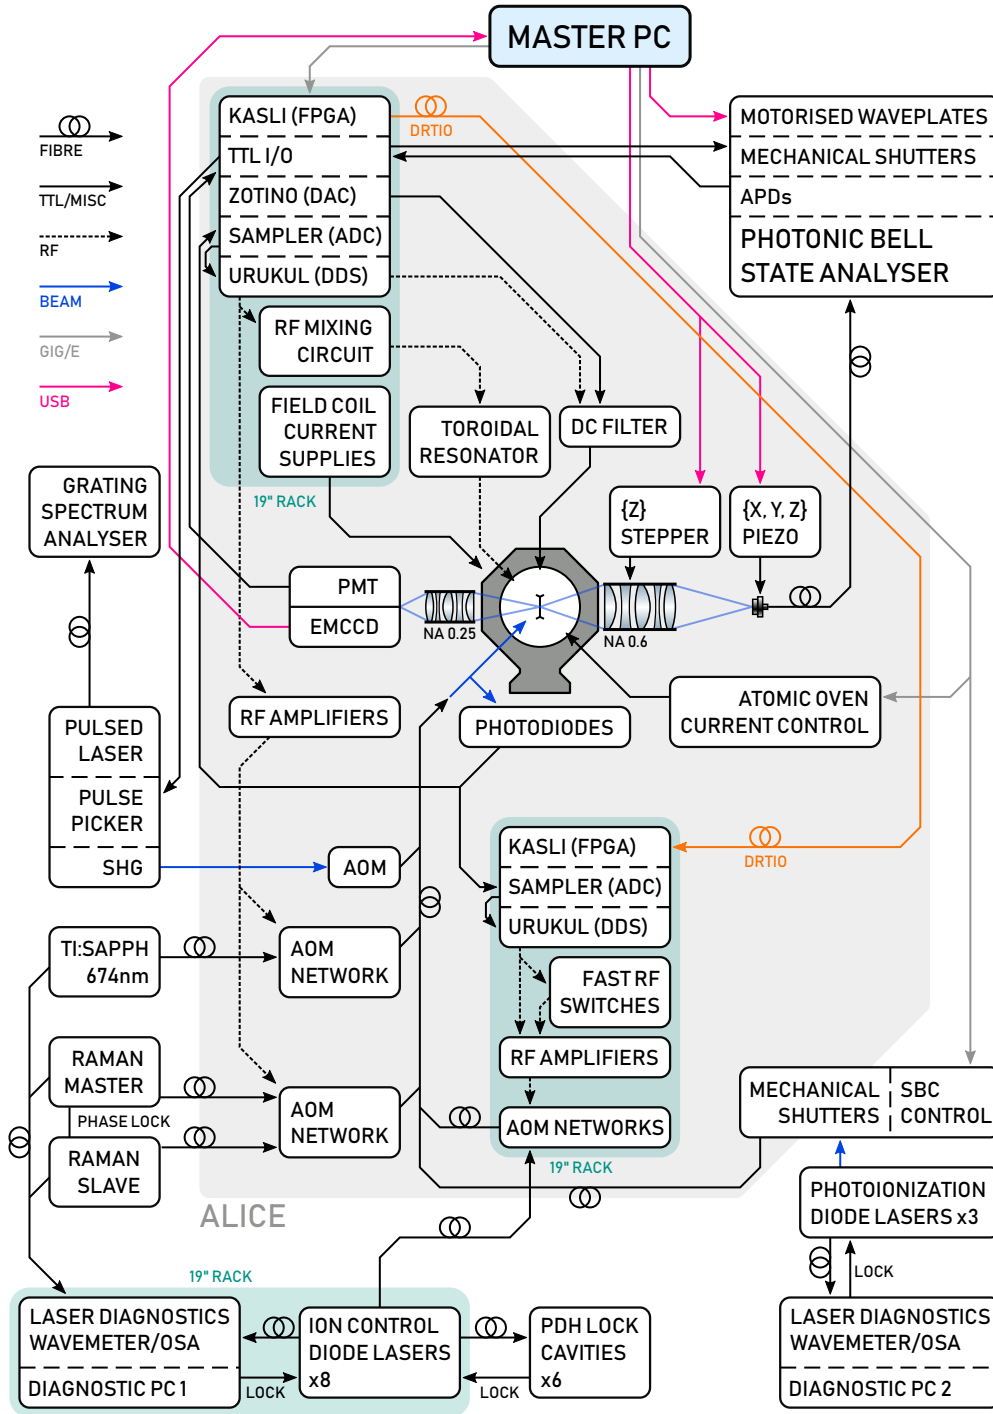


Figure 3.1: Block diagram showing an overview of the subsystems used to control the two trap experiment. The grey shaded background shows the section of the experiment duplicated for each of the traps (arbitrarily labelled Alice here); we omit the duplication in this diagram. Subsystems inside teal boxes are housed in standard 19" server racks. Most subsystems are discussed in detail in this chapter, with the exception of the photon collection and interference apparatus.

3.1 The ion trap

To reduce development time, we used a trap fabricated by Sandia National Laboratories¹, the “High Optical Access 2” trap, henceforth referred to as HOA2 [Mau16]. The trap allows us to image the ion from two directions, one for single photon collection and one for fluorescence detection, and segmented electrodes mean that multiple trap zones are possible for future work on entanglement distillation where shuttling operations are required.

3.1.1 Trap design

The HOA2 trap is a microfabricated planar linear Paul trap, with a total of 94 DC control electrodes, seen schematically in figure 3.3 and imaged in figure 3.2. The number of DC electrodes on the trap itself is much larger than this; however, in order to keep the number of control voltages to within the number available on the trap mounting package, many voltages are repeated. The central linear “quantum” region consists of 19 individually controlled electrode pairs, along with 2 electrodes that can be used to produce a compensation field. This is separated from a Y-junction at each end by “shuttling” and “transition” regions. The Y-junctions allow deterministic reordering of ion crystals.

The trap die is bow-tie shaped with an isthmus that is 1.2 mm across, allowing laser access in the surface plane with NA 0.11 perpendicular to the isthmus and NA 0.08 at 45° and 135°. A 60 μm slot in the trap substrate allows imaging of the ion with NA 0.25 from underneath the trap.

¹Sandia National Laboratories, Albuquerque, NM, USA




Figure 3.2: A photograph of the HOA2 trap – the bow-tie shape of the trap die and the slot through the trap are clearly visible, allowing excellent laser access and imaging from the rear side of the trap respectively. (Image from [Mau16].)

The trap also incorporates holes for loading ions from the rear side into a region on the ends of the Y-junction. This could be useful for loading separate species into different zones and storing them in case of ion loss in the quantum region. However, we chose to load parallel to the trap plane so that the atomic ovens would not obstruct the light cone through the slot for the rear imaging system.

3.1.2 Electrical connections and trap mounting

The trap is mounted in a zero insertion force (ZIF) socket made of PEEK², allowing easy replacement. The ZIF socket is rigidly mounted with screws to a stainless steel plate, which itself is screwed to a stainless steel pedestal on the vacuum base flange. This mechanically registers the trap to the base flange, placing the trap at the centre of the vacuum chamber. The ZIF socket is soldered to a printed circuit board (PCB). This PCB has no components

²Tactic Electronics PN: 100-4680-001A




Figure 3.3: Schematic of the HOA2 electrode layout. Only the central region is used for trapping in this thesis. The RF trapping voltage is on the electrode leading in from the left of the schematic; the mirror electrode to this one is grounded. (Image from [Mau16].)

and serves only to carry electrical connections to the trap. A custom PEEK bodied connector³ takes the DC connections from the 100-pin Micro-D pin feedthrough to the PCB.

The first iteration of the PCB was gold on an alumina substrate⁴, with palladium solder pads. This proved difficult to work with, as the gold traces would wick into the solder if any solder overflowed the pad. This problem could easily be rectified in a future iteration with specification of larger solder pads. However, a more significant issue was that the high density of holes (0.05 inch spacing) for the Micro-D connector made the board fragile and vulnerable to thermal cycling, and the first board failed with a crack across the whole connector region.

The second and final version of the PCB was made with a laminate board⁵. The laminate board had its own drawbacks – being more flexible led to deformation once mounted, with some traces coming very close to the (grounded) support structure, which had been designed with the rigid ceramic substrate in mind. This was not initially noticed when building the first system, Alice, and led to shorting issues, exacerbated by the use of an aged ion pump which shed titanium flakes into the system – this theory was confirmed on Alice’s disassembly. To rectify this, for the second system, Bob, (and when rebuilding Alice) we wedged in a small strip of kapton as an insulator. We used a lead-free no-clean solder⁶, however, as a precaution, we cleaned first with NaOH, then the standard cleaning procedure outlined in

³Winchester Interconnect, formerly SRI Hermetics, PN: SRIMD220-100-NP

⁴CMS Circuit Solutions

⁵Board material: Rogers RT/duroid 6002, board printed by Cirexx International Inc.

⁶Kester Sn 96.5 Ag 3 Cu 0.5

section 3.2.1.

The radio-frequency (RF) trapping voltage reaches the trap mounting PCB via a copper wire connected to a single-sided SMA feedthrough⁷ on the base flange of the vacuum system. A copper wire for grounding is directly connected to the base flange metal. Bare male Sub-D pins are soldered to the PCB to which the RF and grounding wires mate with the appropriate female pin. This arrangement minimises the parasitic capacitance from the trap to ground.

An unintentional consequence of the PCB design was very poor thermal conductivity from the trap to the environment, which is problematic because the power dissipated by the RF trapping voltages will cause a rise of the trap temperature. This is undesirable since trap temperature correlates with the anomalous heating rate [BSC15], as well as causing a local increase in pressure due to outgassing which will reduce the lifetime of ion crystals. The main conduction path for the heat is through the pins of the socket and then into the PCB traces, the cross-sections of which are small and poorly thermally contacted, so conduction is very poor. The HOA2 trap die incorporates a resistive aluminium wire which can be used as either a heater or for temperature measurement. The temperature rise in Bob at steady state induced by the applied RF trapping voltage was measured to be ~ 105 °C. Alice's trap was an earlier version not equipped with this wire so no measurement could be taken, although a similar change can be assumed. To mitigate this effect, future designs using ZIF sockets should incorporate a thermal grounding strap to contact the ceramic chip carrier.

⁷Allectra PN: 241-SMA50-NC

A ground shield 7.3 mm above the trap protects the trap zone from stray electrostatic fields from dielectric materials, *i.e.* the reentrant viewport, which is only 11 mm from the ion. The ground shield is a gold plated alumina substrate PCB⁸. A 18 mm diameter circular hole through through the PCB allows imaging with at least NA 0.6 over the entire trap region. The PCB also incorporates an open-loop microwave antenna for addressing the Calcium-43 ground levels. An SMPM connector⁹ on the PCB enables connection to a semi-rigid cable¹⁰ to a double-sided SMA feedthrough¹¹ on the base flange.

DC connection tester

The HOA2 trap requires 100 connections to be made from the DC feedthrough to the trap via the PCB, the manual debugging of which would be rather time consuming. A circuit was designed by Tim Ballance to test all the connections at once from the air side, and is shown schematically in figure 3.4. The trap incorporates 1 nF trench capacitors, and so by charging a capacitor and measuring the discharge time through a reference resistor, we can infer the capacitance. Any leakage resistance across the capacitor can be measured because the capacitor will not charge up to the supply voltage. Analogue switches allow each of the 100 connections to be tested in turn, as well as a reference 1 nF capacitor on the test board. Before connecting the PCB and

⁸CMS Circuit Solutions

⁹Radiall, PN: R201423110

¹⁰Assembled by Microtek Components Ltd., using Radiall PN:R201052000 (SMPM plug), Radiall PN:R127052001 (SMA 2.9 plug, silicon gasket removed) and Coax Co. Ltd., Japan, PN:SC-219/50-SC-C-PEEK (semi-rigid coaxial cable with PEEK insulation)

¹¹Allectra PN: 242-SMAD18G




Figure 3.4: Diagram of the DC connection test circuit and trap model; A_{OUT} is set to 1 V to charge the trap trench capacitor C_{TR} , the feedthrough capacitance C_{FT} and the cable capacitance C_C . The voltage measured at V_{R-} is dictated by the voltage divider formed by R_{REF} and the leakage resistance R_L . When A_{OUT} is set to zero, the time constant is set by the capacitance and the parallel resistance to ground of the two resistances. A network of analogue switches allows each trap connection to be tested, as well as a reference capacitor included for verification of the test process.

trap, the total capacitance of the cable and feedthrough are measured to be in the range 400–500 pF.

The analogue output A_{OUT} is set to 1 V and the voltage allowed to stabilise for some time. At $t = 0$ the voltage is

$$V_{R-} = 1 \cdot \frac{R_L}{R_{REF} + R_L},$$

where ideally $R_L \rightarrow \infty$ and so $V_{R-} \rightarrow 1$. The output is then set to zero, and the voltage decays exponentially

$$V_{R-}(t) = V_{R-}(t=0) \cdot \exp^{-t/R_{\parallel}C_{\parallel}},$$

where the capacitance is $C_{\parallel} = C_C + C_{FT} + C_{TR}$ and the resistance is $R_{\parallel}^{-1} =$

$$R_{REF}^{-1} + R_L^{-1}.$$

Errors can be easily identified by looking at the value of the capacitance or resistance to ground with the trap and PCB connected. A capacitance of 1500 pF would be considered normal (trench capacitor plus cabling and feedthrough), with infinite (unmeasurable) resistance to ground. Open circuits result in just the capacitance of the cable being measured. Shorts to ground are easily identified because the capacitor does not charge, and so zero resistance to ground is measured. Shorts to neighbours are also easily identified – the trench capacitors will be connected in parallel, so their capacitances add. For two neighbours the capacitance measured is 2500 pF and so on. In this way we can quickly test and debug the DC PCB.

To avoid contamination of the trap during electrical testing, we test the PCB using a sample trap with known electrical failures. By inserting the trap at different orientations all connections can be tested.

The laminate PCB was slightly thinner than the original ceramic PCB specified during the design stage, meaning that for full engagement of the connector to the feedthrough the PCB was flexed downwards. The resulting lateral force on the connector caused intermittent open circuits in the connector to feedthrough interface, which were rectified by retracting the connector slightly from full engagement.

3.1.3 DC supply and filtering

The DC voltages for the trap electrodes are produced by up to three 32-channel 16-bit DACs¹². For this work only one DAC is used, controlling the

¹²ARTIQ Sinara hardware: Zotino

central 15 pairs of electrodes and the two compensation electrodes, with all other DC electrodes grounded. The DAC outputs are filtered, however, due to the large length of cabling between the DAC and the trap a dedicated filter board (designed by Chris Ballance) was placed closer to the trap. One SCSI cable¹³ per DAC connects the DACs to the filter board

The filter board is mounted as close as possible to the feedthrough to minimise the remaining cable length that can pick up electrical noise. The cable to the feedthrough is a 30 cm shielded 2×50-pin Sub-D to 100-pin Micro-D cable¹⁴. The filter circuit is a second order RC low-pass filter¹⁵. On one of the compensation electrodes (Q40) and one pair of centre region electrodes (Q16 and Q18) a 20 pF capacitor bypasses the filter to allow an RF voltage to be coupled via a transformer¹⁶ to the electrode. This allows excitation of the secular motion and hence determination of the secular trap frequencies.

3.1.4 Trap RF supply

In order to achieve stable confinement in a Paul trap, an appropriate choice of frequency and voltage of the driving RF voltage must be made [Pau90]. Ideally, we would like to achieve stable confinement for both Strontium-88 and Calcium-43, with a secular frequency of around 4 MHz for Strontium-88 – with this in mind, an appropriate target is 230 V zero-to-peak amplitude at 60 MHz.

¹³Stonewall Cable, PN: 800-523-3303

¹⁴Custom assembly from Axon' Cable S.A.S.

¹⁵490 Ω , 1 nF

¹⁶Coilcraft WBC1-1LB

	$V_{\text{out}}/V_{\text{in}}$	$V_{\text{out}}/V_{\text{pickoff}}$	f_0/MHz	FWHM/kHz	Q
Alice	10.4	73	28.8	497	57.9
Bob	8.8	57	50.3	1355	37.1

Table 3.1: Toroidal resonator parameters for Alice and Bob. $V_{\text{out}}/V_{\text{in}}$ is the step-up ratio and $V_{\text{out}}/V_{\text{pickoff}}$ the ratio of the pick-off voltage to the voltage at the trap. The centre frequency f_0 , full-width half-maximum (FWHM) and Q values of the resonance are also listed.

We achieve trapping RF voltages with a resonant transformer impedance matching the (approximately capacitive) trap to the standard $50\ \Omega$ RF sources. The transformer is wound on a toroidal carbonyl iron powder core¹⁷. Each resonator includes a capacitive pickoff (shown in figure 3.5) to enable monitoring the RF output voltage in situ with an oscilloscope. The S_{11} parameter of the transformer connected to the trap is measured over the resonance to determine its properties; these results are summarised in table 3.1. The capacitive pickoff was $50\ \Omega$ terminated for these measurements. The pickoff ratio (which is independent of the resonance properties) was measured at multiple frequencies near the resonance with both the transformer’s output and the pickoff measured on an oscilloscope, and found to be independent of frequency over the small range of frequencies near resonance.

We did not achieve the target parameters in either trap, resulting in Alice being unable to trap Strontium-88 and Calcium-43 simultaneously. Trap secular frequencies for Strontium-88 were 4.4–4.8 MHz in Alice, and 2.5–3.1 MHz in Bob.

The trap RF originates from a direct digital synthesiser¹⁸ (DDS), passes

¹⁷Micrometals T50-17

¹⁸ARTIQ Sinara hardware: Urukul

through a mixing circuit with an auxiliary input, and then is amplified¹⁹ before connecting to the toroidal resonator.

The mixing circuit, shown in detail in figure 3.5, is used for micromotion compensation by ‘parametric excitation’ [KPBM15]. It consists of a power splitter²⁰, mixer²¹ and directional coupler²², and allows sidebands to be added to the RF supply, amplitude modulating the RF voltage at the sideband difference frequency. This amplitude modulation translates into a modulation of the secular frequency; if the ion is at the centre of the RF pseudo-potential well, then the net force on it is zero, but if it is displaced from the centre, then it experiences a force proportional to the displacement, oscillating at the modulation frequency [ITU11]. Either the change in fluorescence due to heating [TMA⁺12] or the correlations of photon arrival times [Elt13] can be used as a signal for reducing the excess micromotion.

3.2 Vacuum system

Ion trapping in general requires very high quality vacuum conditions, since collisions with background gas are undesirable regardless of the application. Collisions can cause decoherence, ion loss and, in multi-ion chains, reordering of the crystal [WMI⁺98, ASS⁺10]. We built two vacuum systems, necessitating a design that could be assembled to achieve a high standard of vacuum quickly and efficiently.

¹⁹MiniCircuits ZHL-03-5WH+

²⁰ZFSC-2-4-S+ (Alice, 0.2–1000 MHz) or MiniCircuits ZFSC-2-1S+ (Bob, 5–500 MHz)

²¹MiniCircuits ZP-3MH-S+

²²MiniCircuits ZFDC-20-3-S+

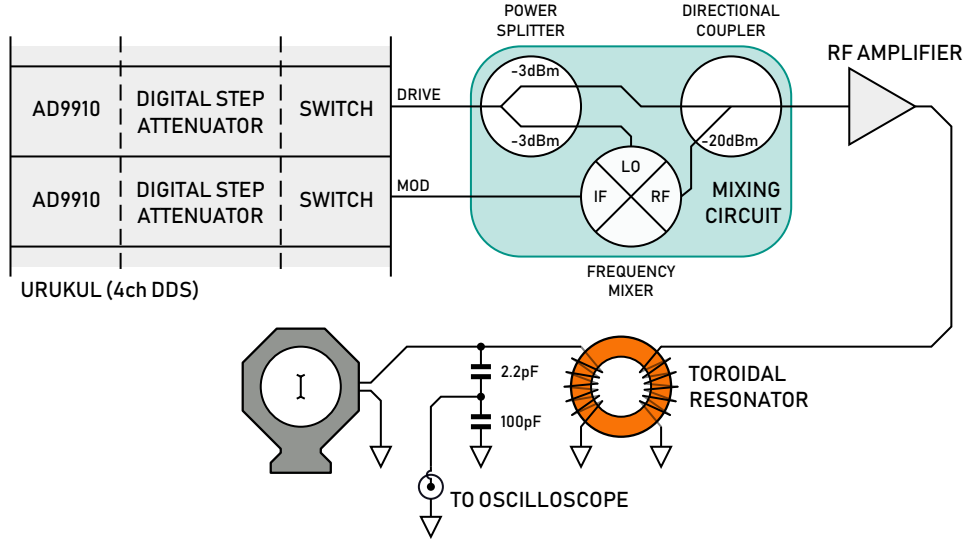


Figure 3.5: The trap RF chain. The mixing circuit (components listed in main text) allows a small amplitude modulation of the RF voltage for parametric excitation. A capacitive pickoff is used to monitor the output voltage of the toroid with an oscilloscope.

The trap is housed at the centre of a 6 inch “spherical octagon” chamber²³, with two CF100 ports and eight CF40 ports. Six CF40 anti-reflective (AR) coated viewports²⁴ allow laser access at $\pm 45^\circ$, $\pm 90^\circ$ and $\pm 135^\circ$ to the trap’s longitudinal axis. A re-entrant CF100 viewport²⁵ allows optical access to the front side of the trap with the short working distance (~ 17 mm) required by the high NA lens.

The CF100 flange²⁶ at the rear side of the trap is custom designed, with two DC feedthroughs (100-pin Micro-D for DC trap voltages, 15-pin Micro-

²³Kimball Physics PN: MCF600-SphOct-F2C8, 316L stainless steel

²⁴Allectra PN: 190-65931-3. Per surface reflectance: $R < 0.25\%$ at 397 nm and 422 nm, $R < 0.35\%$ at 674 nm and 729 nm, $R < 1\%$ at 375 nm, 393 nm, 408 nm, 423 nm and 461 nm, $R < 5\%$ at 850 nm, 854 nm, 866 nm, 1004 nm, 1033 nm and 1092 nm and $R < 10\%$ over 375–1092 nm

²⁵Custom made by Torr Scientific

²⁶Allectra PN: 190-62056-5

D for oven current and thermocouple monitoring), two SMA feedthroughs (for trap RF and microwave access) and with an AR coated viewport (same clear aperture and coating as the CF40 viewports above). This flange also incorporates threaded holes for mounting both the trap pedestal and atomic oven subassemblies; mounting all the internal elements in this way allows the assembly to be conducted on the flange rather than inside the vacuum system, with the result shown in figure 3.6. After the assembly on the base flange is complete, it can be simply lifted (using a lab jack) into the vacuum system. Three lengths of studding threaded into the tapped CF holes of the octagon locate in the through holes of the base flange to ensure that it is well aligned to the octagon during the lifting procedure. This eases the process of initial assembly, as well as maintenance such as trap replacement.

Throughout the assembly procedure and baking, the vacuum system is mounted on a welded steel cradle via M6 tapped holes in the spherical octagon. The cradle forms a stable base to work on the vacuum system, with handles to lift the system without stressing any of the CF joints or inadvertently knocking delicate windows or feedthroughs. The cradle also elevates the system to allow the base flange assembly to be inserted from below with a lab jack.

3.2.1 Assembly and baking

Achieving ultra-high vacuum (UHV) pressures of 1×10^{-12} mbar requires careful preparation of all in-vacuum parts. The general cleaning procedure was a four-stage process, using an ultrasonic bath to agitate the parts for




Figure 3.6: Render of the vacuum flange assembly, incorporating all necessary feedthroughs, and mechanically mounting all internal components of the vacuum system. The rear view illustrates the window for imaging the ion through the trap slot.

5–10 min. Between each stage the parts are rinsed in de-ionised (DI) water.

The stages are as follows:

- detergent solution²⁷
- phosphoric acid & detergent solution²⁸
- HPLC grade acetone
- DI water.

The ceramic PCBs were cleaned by gently swilling in HPLC grade acetone, then isopropanol, and finally DI water. Viewports are supplied clean and cannot be ultrasonically cleaned as they would be damaged. Vacuum valves also cannot be cleaned without disassembly because of the lubricant for the drive screw.

Outgassing from stainless steel parts can be reduced by baking in air at high temperatures before vacuum assembly [WHWM97]. This forms an oxide skin inhibiting the rate of H_2 outgassing from the bulk of the metal. We baked all stainless steel parts for a minimum of 4 h at 450 °C.

The stainless steel parts were then assembled with the pumps and vacuum gauge, along with blank flanges in place of the viewports, and “hard baked” at 350 °C under vacuum. This reduces the outgassing rate of the chamber [CL67], allowing lower ultimate pressures.

The trap itself can only be baked for a limited amount of time because of the formation of “purple plague”, where a brittle alloy is formed at the interface of the gold bond wires and aluminium pads, greatly reducing the

²⁷Decon neutracon[®], diluted to $\sim 3\%$

²⁸Trimate Deoxidine 624, diluted to $\sim 10\%$

strength of the bonds [Che67]. Thus, we perform a trial bake, incorporating all components aside the trap and the atomic source metal. We bake at 200 °C for approximately one week – this bake’s pressure is initially dominated by the water absorbed by the PCB during the cleaning process. If a satisfactory pressure (about 1×10^{-7} mbar at 200 °C) is not reached after a week, this allows us to attempt to diagnose the cause without having wasted a trap.

Finally, the trap is inserted into the ZIF socket, the atomic ovens are filled and crimped, and the vacuum system is baked again at 200 °C for no more than 7 days, or until the rate of pressure decrease has stagnated. During all bakes, the pumps and gauges were degassed appropriately as described in section 3.2.2.

3.2.2 Pumping

Multiple technologies exist for pumping at ultra-high vacuum, no single one of which can pump all gases. Sputter ion pumps are effective at pumping most gases, but ineffective for pumping hydrogen, so should be used in conjunction with another pump [Sin71]. Titanium sublimation pumps (TSPs) evaporate titanium onto the walls of the vacuum chamber, which captures chemically reactive gases, but not noble gases or methane. Non-evaporable getters (NEGs) adsorb gases at the surface, but also employ gas diffusion into the bulk of the material. NEGs do not pump noble gases or hydrocarbons.

The initial design utilised a combined ion pump and NEG unit²⁹. However, we found analysis of the residual gas traces during the bakeout difficult while activating the NEG as the hydrogen and methane partial pressures in-

²⁹SAES NEXToorr D200-5

creased greatly. Despite good pressures when used with previous systems³⁰, the ultimate pressure after the trial bake was unsatisfactory. This design contains considerably more kapton and PEEK than the previous system, so the greater outgassing of these may have been the reason for the higher pressure.

The design was modified to add an extra ion pump³¹. Unfortunately this ion pump had already seen considerable use – unbeknownst to us it contained a significant amount of titanium flakes, which caused electrical problems later on as described in section 3.1.2. However, the additional pump did result in a good ultimate pressure, below the X-ray limit of the nude Bayard-Alpert guage³² used to monitor the system pressure.

For the assembly of the second system (Bob), a third pump system was produced using an ion pump³³ and TSP³⁴, mounted to a custom welded chamber³⁵. The welded chamber is more compact than an assembly from off-the-shelf components, as well as reducing the number of potential leak points and weight while remaining lower in cost. Alice was later converted to a copy of this design, with a new ion pump³⁶.

In addition to the main pump systems, small metal strips with getter material fixed on both sides³⁷ were added, one surrounding the metalwork of the reentrant viewport and spot-welded to two “groove grabbers”³⁸, and the other screwed to the blank CF40 flange at the bottom of the system, near

³⁰Used in the most recent “blade” trap experiment

³¹Varian VacIon Plus 20 Diode (201s⁻¹)

³²Varian UHV-24p, X-ray limit 5×10^{-12} mbar

³³Varian VacIon Plus 20 Triode (201s⁻¹)

³⁴Gamma Vacuum 3-filament TSP cartridge, PN: 360819

³⁵Supplied by Scanwel Ltd

³⁶Varian VacIon Plus 20 Diode (201s⁻¹)

³⁷St707/CTAM/30D

³⁸Kimball Physics, PN: MCF600-GrvGrb-C01

the atomic ovens. These are activated passively during the bakeout.

3.2.3 Atomic ovens

A new atomic source was designed by Tim Ballance and characterised extensively in a test vacuum system [BGN⁺18]. Neutral atoms are delivered to the trap zone by resistive heating of stainless steel tubes containing either calcium or strontium. The distinguishing feature of this design is the inclusion of K-type thermocouples spot-welded close to the oven apertures. This allows closed-loop control of the oven temperature (using a dedicated microcontroller to regulate the current supplied), meaning that the ovens can be rapidly heated to their operating temperature. The typical time to produce an appreciable neutral flux is thus reduced from 2–5 min with constant current atomic ovens to around 15 s, enabling much faster loading times. Since the long-term goal of the trap systems is to be used with multi-ion chains, which can be subject to loss events every few minutes, this is vital to improving the experiment duty cycle.

The atomic sources are designed as an easily repeatable assembly. Stainless steel flexure clamps hold the ends of the oven tubes and act as the electrical contacts, while a ceramic³⁹ back piece provides the mounting point for the flexure clamps. Holes in the ceramic back act as collimation apertures for the atomic beam.

Current to the ovens is supplied through the 15-pin Micro-D feedthrough in the base flange, via a Micro-D pig tail assembly⁴⁰, which is also used for

³⁹Boron nitride was used; this has a higher thermal conductivity than Macor, undesirable for this use case, but better machining properties.

⁴⁰Winchester Interconnect, formerly SRI Hermetics, PN: SRIMD204-15-9

the thermocouple voltages. To minimise the possibility of contaminating the base flange, the wiring of the ovens was completed on an aluminium jig before transferring the completed assembly onto the base flange.

A glass slide was balanced in front of the oven during bakeout and during “cracking” of the atomic oven [All11] to reduce the risk of plating the trap. In Alice, the temperature used in the initial heat treatment of the strontium oven was too conservative and so cracking was not observed until the glass slide had already been removed (by tilting the vacuum system until the slide fell). Cracking of Alice’s strontium oven occurred at an indicated temperature of 430 °C, while Bob’s occurred at 475 °C.

The operational temperature of the ovens is approximately 230–300 °C for both calcium and strontium. We would expect that for strontium a given vapour pressure (and hence neutral atom flux) would be achieved at a substantially lower temperature [AIH84], so it seems that the temperature reading from the thermocouple is not a good absolute measure of the temperature of the source metal, probably due to variations in the distribution of source metal within the oven tube. The indicated temperature is, however, repeatable on a per oven basis and so can be used for neutral flux regulation.

We elected not to use the loading holes (see section 3.1.1) in favour of loading from the side of the trap. This did result in some plating of the trap in Bob, causing ohmic shorts to ground of some DC electrodes, the lowest resistance of which was measured (both with the trap electrode tester and a digital multimeter) as 32.3 k Ω . This resistance is high enough that the DAC supplying the DC trap voltages (of order 1 V) is not sourcing significant amounts of current and so does not experience voltage drop. However, it is

unclear what the implications are for the micromotion environment or heating rate due to surface effects. The resistance to ground was not observed to change after a significant time running the ovens at their usual temperature, so it is likely that the plating was a one-off event while attempting to crack the ovens. Plating was not observed in Alice, so it is possible that Bob's ovens were poorly aligned and plating can be avoided with good alignment, however, loading holes remain the preferred way of mitigating this issue.

In operation, a frequency offset of ~ 50 MHz in peak neutral fluorescence of strontium (and hence optimal loading frequency) between the two systems was observed. This is due to a Doppler shift from the the atomic beam not being directly perpendicular to the laser beam, as the two oven tubes are separated by 4 mm. Since both systems were constructed identically but the beams are in a mirrored configuration, the sign of this shift is reversed thus causing an offset between the systems. A rough estimate of the difference in frequency (using the modal velocity of the Maxwell-Boltzmann distribution at the oven temperature indicated) is around 60 MHz for strontium, in agreement with the measured value. This offset is larger than the linewidth of the neutral feature, and so to load each system the laser frequency must be tuned on resonance. While this is not overly arduous, the laser in question can be only be tuned globally, so loading simultaneously in both traps was not possible. Therefore, when Alice's trap was replaced, the ovens were also refilled in a mirror configuration to eliminate this offset.

3.2.4 Vacuum mounting structure

The vacuum system is mounted vertically between two 10 mm thick stainless steel plates, which are fastened to the vacuum system via eight M6 screws apiece. These plates attach to a base plate which screws directly to the optical table. The side plates have mounting points for the final delivery optics as well as the magnetic field coils, which can be mounted in two mirrored orientations.

3.3 Magnetic field coils

Magnetic field is provided by three orthogonal pairs of coils, wound in-house⁴¹ using 0.9 mm enamelled round copper wire. The windings are split into unequal windings, allowing a variety of configurations, with the properties summarised in table 3.2 For initial experiments, the inner windings of each pair were used in Helmholtz configuration to provide the quantisation field and trim its direction.

Current is supplied by a single channel current supply⁴² for the main coil and one channel each of a triple-channel supply⁴³ for the trim coils – for the output range selected (500 mA), the specified current noise is 20 μ A. The typical field strength used is 0.561 mT at a main coil current of 240 mA. The magnetic fields of the two vacuum systems are at right angles to each other in the lab frame; one of the systems (Bob) required no trim currents, suggesting that its field was aligned close to the lab’s magnetic field. The magnitude of

⁴¹With thanks to Tony Hickman

⁴²TTi QL355P

⁴³TTi QL355TP

Coil	N	d (mm)	Sensitivity (mT A ⁻¹)
Main Inner	324	65.5	2.30
Main Outer	96		0.68
Circular Trim Inner	128	52.8	1.14
Circular Trim Outer	76		0.68
Square Trim Inner	195	129.5	0.33
Square Trim Outer	96		0.17

Table 3.2: Magnetic field coil specifications. N gives a median value of the turns per coil (two pairs of each coil type were produced). The displacement d is from the centre of the coil to the ion position, and the sensitivity is calculated by integration of Biot-Savart’s Law, assuming a pair of coils wired in series.

Earth’s magnetic field in London is around 0.047 mT [BCM00] – a “square” trim coil current (see figure 3.9 and table 3.2) of 142 mA would be needed to correct this if aligned exactly in that direction. Alice’s square trim coil current was ~ 400 mA, with a circular trim coil current < 50 mA.

3.4 Experimental Control

Experimental control is entirely conducted via the ARTIQ⁴⁴ framework. This is an open-source control system consisting both of hardware and software designed specifically for quantum information experiments. A high level programming language based on Python is used to write experiments, which can contain real-time sections with strict timing requirements, called kernels, and sections with slacker requirements. The kernels are compiled and executed on an FPGA, while the slack sections are executed on the “host”

⁴⁴M-Labs, Advanced Real-Time Infrastructure for Quantum Physics, DOI: 10.5281/zenodo.1492176, version 4.0

computer CPU. We use ARTIQ’s Sinara hardware family⁴⁵ for handling real time operations in experiments.

The Sinara family consists of several commonly used generic devices of an atomic physics experiment, of which this experiment uses:

Kasli: The main FPGA carrier used to control the other Sinara devices, executing TTL pulse sequences and handling input events such as photon detection. The Distributed Real-Time Input Output (DR-TIO) protocol, transmitted over optical fibre links, allows the “core” Kasli device to delegate control to up to two “slave” Kasli devices. The core device is connected to the host computer by a Gigabit Ethernet link.

Urukul: Frequency synthesiser using four Direct Digital Synthesiser (DDS) chips⁴⁶ allowing frequency (sub-hertz resolution), phase and amplitude control of signals with an analogue bandwidth of 400 MHz, and a frequency update rate of 500 kHz. The Urukul also incorporates per-channel digital attenuation in 0.5 dB steps to -31.5 dB and RF switches. The Urukuls in this experiment are used to generate RF for the acousto-optic modulators (AOMs) and the trap RF.

Zotino: 32-channel, 16-bit Digital-to-Analogue Converter (DAC) with an update rate of 1 MSPS, used for the DC trapping voltages.

Sampler: 8-channel, 16-bit Analogue-to-Digital Converter (ADC) with an update rate of 1.5 MSPS. The Sampler can be used in combination

⁴⁵<https://m-labs.hk/experiment-control/sinara-core/>

⁴⁶AD9910 variant used

with the Urukul to stabilise laser amplitudes – this was implemented after data collection of this thesis.

TTL I/O: 8 channel isolated digital input/output device, for simple TTL signals. Used for controlling dedicated RF switches and to detect input events from photon detectors.

Clocker: Clock distribution module – to synchronise the Kasli core devices of Alice and Bob.

The above devices are all designed to mount in a Eurocard 3U subrack, and connect to the Kasli using standard IDC ribbon cables – all input/output events are ultimately managed by and synchronised to the Kasli.

Other hardware devices which are normally controlled by a computer can be used during experiments by writing ARTIQ drivers. This allows any computer that has a network connection to the host computer to provide a remote procedure call (RPC) interface to the underlying device. These RPCs are subject to network and other latencies (for example USB communication) so cannot be considered real time.

3.5 Laser systems

A combination of extended cavity diode lasers⁴⁷ and Ti:Sapphire lasers are used in our experiment. Note that the same lasers are used simultaneously for both traps, necessitating a method of independently tuning the optical frequency at the trap. A small amount of light from each laser (except the

⁴⁷Toptica DLPro with DLC pro controllers, unless otherwise noted

pulsed 422 nm and 378 nm lasers) is directed towards a diagnostic system consisting of optical spectrum analysers⁴⁸ and a wavelength meter⁴⁹. The laser to be monitored is selected using a fibre switch⁵⁰. The switch and diagnostic system are mounted in a standard 19-inch server rack to save space.

3.5.1 Diode lasers

For cooling, state preparation and readout in strontium we have 3 lasers (422 nm, 1033 nm and 1092 nm), and for calcium we have 5 lasers (393 nm, 397 nm, 850 nm, 854 nm and 866 nm). The strontium lasers, the 397 nm laser and their controllers are mounted in drawers [Nou16] in a 19-inch rack, while the remaining diode lasers were used in a pre-existing experiment and set up on a dedicated optical table [Sch18].

Six of the lasers (422 nm, 1092 nm, 393 nm, 397 nm, 850 nm and 866 nm) are frequency stabilised by Pound-Drever-Hall (PDH) locks to low-drift, tunable reference cavities⁵¹. The frequency can be tuned by adjusting the high-voltage supply⁵² for the piezo-electric cavity mirror mounts. However, we restrict this tuning to infrequent coarse adjustments since the relaxation of the piezo-electric material [Ar182] will cause slow drifts of the laser frequency. Fine-tuning is performed with the AOM networks described in section 3.5.3.

Diode lasers for photoionisation (PI) were already set up and shared between multiple experiments in the lab. The 461 nm and 423 nm lasers provide

⁴⁸Red lasers: Toptica, FPI-100-0980-1 Blue lasers: Thorlabs, SA200-3B

⁴⁹HighFinesse WS7

⁵⁰Leoni FiberSwitch 1x16

⁵¹NPL Low Drift Etalon or Stable Laser systems, SLS-NPLcav-3

⁵²Thorlabs MDT693B

the first step of the two-photon ionisation process for strontium and calcium respectively. These lasers can be coarsely frequency tuned and stabilised with a PID loop acting on the measured frequency from the diagnostic system, allowing automation of the isotope-selective loading process. A 378 nm laser⁵³ completes the excitation into the continuum for both strontium and calcium. All three of the PI lasers are coupled into the same optical fibre to the experiment, and individually switched with mechanical shutters. The single board computer (SBC)⁵⁴ operating the shutters exposes an ARTIQ RPC interface so that the loading process can be entirely ARTIQ controlled.

3.5.2 Ti:Sapphire lasers

Multiple Ti:Sapphire lasers are shared between the experiments in this lab. Of these, two are used for operations with strontium in this thesis⁵⁵.

A mode-locked Ti:Sapphire⁵⁶ produces a train of 844 nm pulses at 80 MHz, with a pulse duration of approximately 8 ps. An electro-optical “pulse picker”, built and extensively characterised by David Nadlinger [Nad17], allows selection of single pulses with good extinction at short timescales. The pulse is picked approximately 200 ns⁵⁷ after the TTL input trigger signal to the

⁵³Toptica iBeam

⁵⁴BeagleBone Black

⁵⁵A further pair of frequency doubled Ti:Sapphire lasers can be used to perform two-photon Raman transitions in both calcium and strontium. The slave is phase locked to the master with light picked off before the doubling cavity. Vera Schäfer’s thesis [Sch18] contains a more detailed overview; these lasers were not used in this work, but will be used for future mixed-species operations in this apparatus.

⁵⁶Spectra Physics Tsunami 10ps HP, pumped by Spectra Physics Millennia eV 25

⁵⁷This time is between the input signal and the pulse as measured at the exit of the pulse picker, and does not include any of the signal/optical delays present in the experimental setup used. The measured latency between the TTL signal and an optical pulse *at the trap* in the experiment is 372 ns.

pulse picker to allow a safe high-voltage switching sequence of the Pockels cell. There is an inherent 12.5 ns jitter in the timing of the pulse because of the repetition rate of the pulses; a comparator circuit monitors a bare photodiode to detect the laser pulse and provide a reference TTL signal for the experiment at the time of the pulse⁵⁸. The pulse then passes through a second-harmonic generation (SHG) stage to produce the 422 nm light needed to excite the $5^2S_{1/2} \leftrightarrow 5^2P_{1/2}$ transition. The light is split to two AOMs and coupled into fibres to each trap. The AOMs provide additional extinction at long timescales, but cannot be switched quickly enough to select individual pulses.

The 674 nm light to excite the $5^2S_{1/2} \leftrightarrow 4^2D_{5/2}$ quadrupole transition of strontium is provided by a commercially stabilised⁵⁹ Ti:Sapphire laser⁶⁰. This system is specified to have less than 10 Hz linewidth, measured over 1 s. The laser is offset sideband locked to a high-finesse cavity; long-term drifts of the cavity are compensated with a feed-forward on the offset frequency applied to the electro-optic modulator used in the lock. The absolute frequency drift was measured in September 2017 to be 17.2 kHz d⁻¹ [Nad17]. By April 2019 this had reduced to 7.27 kHz d⁻¹.

3.5.3 AOM switching and frequency tuning

Acousto-optic modulators (AOMs) offer a standard way of switching laser beams in atomic, molecular and optical physics experiments. Only when an RF tone is applied to the AOM will it deflect the laser beam so that

⁵⁸The circuit introduces a latency of 6 ns, measured at its output.

⁵⁹Stable Laser Systems

⁶⁰M Squared SolsTiS

is is coupled into the output fibre towards the trap. By varying the frequency of the RF tone, the frequency of the diffracted light can be tuned – allowing us to operate the two traps with independent laser frequencies. However, varying the frequency changes the angle of the diffracted beam. Using AOMs in a double-pass configuration allows the pointing change to be eliminated [DHL⁺05] at the expense of reducing the absolute laser power available. For single-pass configurations, coupling the output into a fibre will transform beam pointing variations into power variations. Power fluctuations can be compensated with a PID loop using photodiodes to monitor the power after the fibre and feeding back on to the AOM RF drive amplitude – this was not yet implemented for the work presented in this thesis.

A reconfigurable breadboard to house each diode laser’s AOM setup was designed by Tim Ballance. The use of compact 0.5-inch optics reduced the size scale considerably compared to previous experiments, and allowed the 10 breadboards to be mounted in a second 19-inch server rack, along with the Urukul hardware and amplifiers⁶¹ supplying the RF. Some configurations are used for simply switching the beams. Others, for example the 422 nm breadboard (shown in figure 3.7), have two outputs, one of which has two frequency components handled by separate AOMs. Focussed AOMs are used⁶² for faster switcher times, with the detrimental side effect of Rayleigh scattering of the beam in the AOM reducing extinction. The breadboard used fibre collimators with 5 degrees of freedom⁶³. These were chosen because the mechanical design allowed easy mounting, however, the unnecessary number of

⁶¹ARTIQ Sinara hardware: Booster

⁶²UV wavelengths: ISOMET 1250C-829A, IR wavelengths: ISOMET 1206C

⁶³Thorlabs FiberPort, various models for different wavelengths

degrees of freedom complicated the already time-consuming task of aligning the breadboards. A more prudent choice for future iterations would be fixed collimators, since all the required degrees of freedom are already included with components on the breadboard.

The AOM network for the 674 nm light is mounted on the optical table near to the traps, and splits the supply from the laser towards the two traps. Single- and double-pass configuration focussed AOMs⁶⁴ are used for each trap beam path; the double-pass allows frequency tuning, while the single-pass offers faster switching, with the layout shown in figure 3.8. The outputs are coupled into short fibres to the final beam delivery optics.

Another AOM network for the calcium Raman beams is mounted on a breadboard on the trap table, built by Shuoming An. This was not used within the scope of this work, although at the time of writing Raman manipulations of Calcium-43 were being performed in one of the traps.

3.6 Trap beam paths

The final stage of the beam delivery is as short and simple as possible to minimise any beam pointing drifts. A generic beam path will consist at most of the following exhaustive set of optical elements:

Fibre collimator: provides a collimated beam; with an adjustable focal length this allows fine tuning of the focus position at the ion

Polarising element: ensures stable polarisation by transforming po-

⁶⁴ISOMET 1205C-848

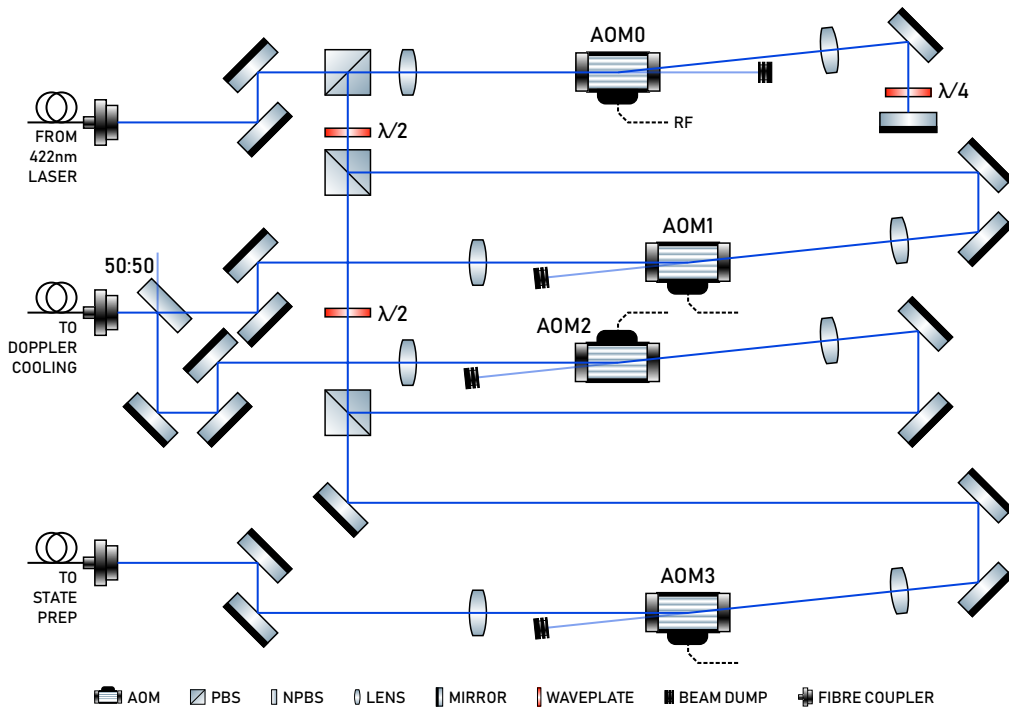


Figure 3.7: Diagram of the 422 nm AOM breadboard setup. A double-pass AOM allows for large common mode frequency tuning of the two output beam paths. Two single-pass AOMs are combined at the first output, which is used for the Doppler cooling beam – note that one of the AOMs is used in the first order and one in negative first order, to give a ~ 400 MHz frequency difference between the two components. This far detuned beam is for additional cooling far from resonance when loading multi-ion crystals. The second output is served by its own single-pass AOM and used for the σ^- -polarised state preparation beam.

This diagram and others make extensive use of the ComponentLibrary by Alexander Franzen, licensed under CC BY-NC 3.0.




Figure 3.8: Diagram of the 674 nm AOM layout; the duplicated section for Bob is omitted.

larisation fluctuations into power fluctuations, usually a PBS⁶⁵

Photodiode pickoff: a low reflectance beam splitter to divert some power to a photodiode, for power stabilisation (not implemented for this work, but pickoffs were in place)

Two steering mirrors: enable each beam to be steered independently onto the ion

Polarisation control: *i.e.* waveplates to transform the polarisation from linear if necessary

Dichroic mirrors: combine different wavelengths of laser light along the same beam path

Focussing optics: focus the nominally collimated beams at the ion.

⁶⁵Thorlabs PBS121/PBS122/PBS123 depending on wavelength




Figure 3.9: Schematic of $^{88}\text{Sr}^+$ beam paths, shown from the high numerical aperture imaging side for Alice. Where two wavelengths are separated by a slash, they are brought to the trap table down the same optical fibre. The magnetic field coils are shown; the applied field direction lies along the teal arrow, while trim adjustments are made along the pink arrow with the “square” coils and perpendicular to the plane of the diagram with the orange “round” coils. Bob’s beam layout and field direction are mirrored so that the majority of the ‘coherent’ beams (674 nm for strontium, and Raman beams eventually for calcium) for the two traps are placed adjacently on the optical table.

For some beam paths requiring additional polarisation purity, an additional Glan-Laser polariser⁶⁶ is used. Given the large number of beams per trap and the need for compactness, a monolithic mount was designed by David Nadlinger to mount the fibre collimator, PBS cube, pickoff, photodiode⁶⁷ and the first steering mirror in an assembly approximately 130 mm in length.

Figure 3.9 shows a schematic view of the final beam geometry for Alice, with the trap orientation also shown. Imaging from both the front and rear sides of the trap necessarily places the trap plane vertical⁶⁸, posing two problems. The first is that there are three heights at which beams enter the vacuum system, and so the upper two levels are placed on breadboards, elevated from the optical table on stainless-steel posts. The second is that some beams are oriented at 45° to the horizontal; this is discussed below in section 3.6.1. Bob's beam paths (and magnetic field) are a mirror image of Alice's, in order to place as many of the beam paths for coherent manipulations as close as possible to each other and the AOM networks supplying the light.

The focussing elements for all beams (except PI beams) consist of a $f = -150$ mm meniscus lens⁶⁹, separated by 150 mm from a $f = +100$ mm plano-convex lens⁷⁰. This forms a system with an effective focal length of 75 mm, but with a back focal distance of 150 mm. This allows relatively small diameter collimated beams to be focussed to small spot sizes at the ion, while

⁶⁶Thorlabs GLB-10

⁶⁷QLNPD, designed by David Nadlinger

⁶⁸Clearly other orientations are possible, but large horizontal dimensions are generally easier to work with than vertical ones!

⁶⁹Thorlabs LF4370-UV

⁷⁰Thorlabs LA4380-A

still being able to position the final lens outside the vacuum system. The PI beams are focussed with a single $f = 150$ mm plano-convex lens⁷¹. The 422 nm and 674 nm beams have a $1/e^2$ beam radius of about $25 \mu\text{m}$, while the 1092 nm and 1033 nm beam has a radius of $\sim 40 \mu\text{m}$.

The final positioning of the pulsed laser beam is performed using a motorised “Picomotor” mirror mount⁷². This uses the stick-slip phenomenon to directly turn the actuator screws in small increments rather than mounting the mirror directly on piezo-electric actuators, which can suffer from relaxation.

3.6.1 Fixed final mirrors

The orientation of the vacuum system means that four beam paths enter the vacuum system at 45° to the optical table, two of which must be placed on breadboards to elevate them to the required position. One solution could be to use a vertical breadboard, pre-position the optical elements and finalise alignment in situ, however, this makes it difficult to reposition components. Therefore we place all elements on horizontal breadboards or the optical table and use a fixed arrangement of mirrors to create the angle required.

A single mirror can be used to take the beam from horizontal to 45° , and would be the simplest solution. However, it is convenient to keep the principal axes of the adjusting mirrors to those of the trap so that adjusting the position of the beam across the trap surface requires turning only one adjuster and not both simultaneously. If using only a single mirror, and imposing the

⁷¹Thorlabs LA4874-A

⁷²Newport New Focus 8821




Figure 3.10: Photograph of the fixed final mirror setup, including the 3D printed servo mirror to perform beam alignment described in section 3.6.2.

constraint that the angle of incidence on the mirror must not exceed 45° , this requires that the beams propagate outward from the centreline of the system to the mirror before being reflected upwards, clearly an impractical solution. Using two mirrors in the configuration shown in figure 3.10 allows the beam to originate from the front or rear side of the vacuum system. The mirrors⁷³ are glued to fixed mounts⁷⁴, which attach to a custom machined stainless steel “tombstone”. This part has mirror symmetry so can be configured in two orientations, and has dowel pins for accurate location when attaching to the vacuum system mounting structure. It also holds the final focussing lens in a slotted mount⁷⁵, with several threaded holes available for coarse positioning. A through hole allows an optional waveplate to be mounted in a high-precision rotation mount⁷⁶, or an iris for initial alignment.

Usually, one would mount waveplates after any mirrors in the beam path. This ensures that the light at the mirror surface is either p - or s -polarised,

⁷³Semrock BBDM

⁷⁴Thorlabs Polaris C1G

⁷⁵Thorlabs Polaris L1G

⁷⁶Thorlabs PRM05

so that any polarisation dependence of the mirror does not change the polarisation of the reflected beam. The fixed mirror arrangement means that it is easier to have the waveplates before the final mirrors, and so the beam may be circularly polarised at the mirror. The birefringence of the mirrors was tested for thermal stability over 10°C at 422 nm and found to have negligible thermal dependence. Unfortunately we did not check the wavelength dependence of the birefringence – this was later found to be significant when using 393 nm and 397 nm copropagating circularly polarised beams for state preparation of Calcium-43. These beams are delivered down the same optical fibre, so any polarising element affects both beams. Using the optional waveplate mounting point for a $\lambda/4$ waveplate, the polarisation of the light at the mirrors was kept close to linear on one of the mirror’s axes, mitigating this effect.

3.6.2 Alignment onto ions

The initial alignment of the beams must be conducted without any ions in the trap, the procedure for which is largely similar to that described by David Allcock [All11]. The initial alignment is performed before the insertion of the focussing lenses, using irises fixed to the final mirror assembly, ensuring that the beam is approximately parallel to the trap surface and centred on the vacuum windows. For most beams this is not critical; however, the orientation of the state preparation beams is fixed indirectly by the reentrant vacuum window⁷⁷, so we try to perform this alignment as well as possible. A

⁷⁷The imaging axis is fixed by the reentrant vacuum window, to which the high NA lens must be aligned to cancel aberrations. The applied magnetic field (and hence the circularly polarised state preparation beams) must lie perpendicular to the imaging axis.

mirror is inserted just before the vacuum to reflect the beam onto a CMOS camera⁷⁸ placed at the virtual ion position. The focussing lenses are then inserted in their nominal positions (to place the focus at the camera), taking care that the beam passes through the optical centre of each lens to minimise the aberrations in the focussed spot. The beams must now be aligned relative to the trap's features. The HOA2 trap's gold coating is highly specular, meaning that oblique illumination is not sufficient to image the trap features. Therefore we use the arrangement shown in figure 3.11 to provide illumination perpendicular to the trap surface and image the trap; alignment marks allow easy identification of the central electrodes. We position the beam so that approximately 50% is clipping the trap (by observing the exit beam), and then scan the beam across the trap until the scattered light is correctly aligned with the trap electrodes. The mirror is then replaced, and the camera used to raise the beam position by the ion height, 68 μm .

While the algorithm for this beam alignment is simple, in practice the geometry of the beam paths makes it somewhat difficult to place a mirror after the final lens. The final lens is 150 mm from the centre of the vacuum system, and thus only 45 mm from the outer edge of the vacuum window and 30 mm from the mounting structure. Conventional flip mounts could not be used due to the space restrictions, so a 3D printed adapter was used to attach a small mirror⁷⁹ to a magnetic mount⁸⁰. The rare earth magnets used in the mount are not ideal for the uniformity of the magnetic field at the trap centre – this is not important for the single-ion experiments reported here

⁷⁸Thorlabs DCC1545M

⁷⁹Edmund Optics, 15 \times 15 mm UV Enhanced Aluminum, $\lambda/4$ Mirror, Stock #45-726

⁸⁰Thorlabs KB25/M




Figure 3.11: The confocal imaging system, used to perform initial alignment of the trap beams. The fibre is deliberately not at the focus of the imaging system so that it provides illumination over a large region of the trap plane. It can also be used to image the ion during initial trapping – the fibre is replaced by a photomultiplier tube (PMT) to detect the photons from a trapped ion. That the PMT is not at the focus of the imaging system is inconsequential since it has a large active area.

but is crucial for multi-ion operation. The mounts themselves are also fairly difficult to place, with a non-negligible risk of nudging optics while trying to insert or remove them.

A second design was produced using servo motors to swing a mirror⁸¹ into the beam. This used 3D printing extensively and included a second mirror to make the beam path more compact and allow mounting the CMOS camera to the vacuum mounting structure. The servo motors still contain permanent magnets, but these are much weaker than the magnetic mounts. The motors are operated remotely using a Beaglebone, so there is no risk to the optics. A weakness of the design was that if the pivot arm holding the mirror was not mated well to the servo motor spindle, the resulting shift in angle of the

⁸¹Edmund Optics, 15 mm Dia. UV Enhanced Aluminum, $\lambda/4$ Mirror, Stock #63-168

reflected beam changed the effective distance to the camera such that the camera was no longer at the virtual ion position. However, in general the use of 3D printing here was a success – the profiling mirrors do not need long-term stability as they are generally used for comparative measurements for a few minutes at a time.

Once one beam is aligned on the ion, copropagating beams can be added trivially by ensuring that they coincide at the camera and one other place. Counterpropagating beams of the same wavelength (*i.e.* the 422σ and pulsed 422 paths) can also be added by coupling the light from the reference beam into the second beam’s coupler.

3.7 Rear imaging system

The rear imaging system⁸² (see figure 3.12) is used to collect fluorescence from both species in the trap, imaging the ions through the slot in the HOA2 trap substrate. The main objective is an NA 0.3 5-element lens⁸³ with a working distance of 59 mm. The dimensions of the trap slot restrict the actual imaged NA to 0.25. The objective is mounted on a 3-axis translation stage⁸⁴ with tip/tilt adjustment provided by a kinematic prism mount⁸⁵ with a custom adapter. The wavelength-dependence of the focal length for light from each species is compensated by splitting the two wavelengths with a dichroic mirror and using independent lenses to focus them, either recombining the two

⁸²Designed by Chris Ballance

⁸³Sill Optics GmbH

⁸⁴Thorlabs MT3/M

⁸⁵Thorlabs KM200PM/M

(with a D-shaped mirror) onto an EMCCD⁸⁶ or diverting them to independent PMTs⁸⁷ for fluorescence counting. A light tight aluminium enclosure forms the mounting structure for the imaging system's components, as well as minimising background light.

In Alice, two strontium ions were trapped at a nominal axial frequency of 1 MHz, giving an ion spacing of 4.3 μm . The spacing between the ions in the image plane of the camera was 13 pixels, or 208 μm , indicating a magnification of 48.4.

We define the overall detection efficiency as the probability of detecting a photon given that a photon was emitted; to measure it we produce single photons from the strontium atom. We pump the ion into the $4^2D_{3/2}$ level using the 422 nm laser, before extinguishing the blue laser and turning on the 1092 nm repumper beam, thus emitting a single 422 nm photon. The PMT is gated for a window 5 μs , many times longer than the $5^2P_{1/2}$ lifetime, ensuring that a photon is emitted. Background counts are monitored with interleaved windows where the ion is not prepared into the $4^2D_{3/2}$ level and so does not emit a photon, and subtracted from the result.

Ignoring losses in the collection optics, the maximum possible theoretical detection efficiency for this setup is $0.016 \cdot 30\% = 0.48\%$, where 0.016 is the fraction of solid angle subtended by an aperture of NA 0.25 and 30% is the PMT's quantum efficiency. The measured detection efficiency of 0.4% (in both systems) compares favourably with the theoretical prediction, resulting in a peak count rate for strontium of approximately 120 kHz. The background

⁸⁶Andor iXon 897 Ultra, PN: DU-897U-CSO-#EX

⁸⁷Hamamatsu H10682-210, quantum efficiency @ 400 nm $\approx 30\%$




Figure 3.12: Schematic view of the rear side imaging system described in section 3.7. The plano-convex lens^a $L1$ directly following the main objective decreases the effective focal length to make the imaging system more compact. Aperture $A1$ at the prime focus assists coarse alignment of the imaging system; best form lens^b $L2$ nominally collimates both wavelengths, up to chromatic aberration. The dichroic mirror^c $DM1$ splits the wavelengths into independent paths. Lens $L3$ focuses 422 nm light through the adjustable aperture^d $A2$ to reduce background scattered light from lasers and room lighting; lens $L4$ focuses the light onto the EMCCD. A D-shaped mirror $M5$ recombines the two paths onto the EMCCD. Mirrors $M3$ and $M3'$ can be moved by linear actuators^e to divert each wavelength to a dedicated PMT. Filters $F1$ – 3^f reduce background.

^aThorlabs LA1779-A ($f = 1000$ mm)

^bThorlabs LBF254-100A ($f = 100$ mm) – lenses $L2$, $L3$, $L4$, $L3'$ and $L4'$ are all identical

^cSemrock Di03-R405

^dOWIS SP 40

^eActuonix L12-30-50-6R

^fSemrock $F1$: FF01-433/24, $F2$: FF01-392/18, $F3$: FF01-440/SP

count rate is typically < 0.2 kHz – very little laser scatter makes it through the imaging system because we are imaging through the slot in the trap, so the background is dominated by room light.

Chapter 4

Photon collection and entangling apparatus

In the previous chapter, we described the “standard” ion trapping apparatus, similar to many other existing experiments; here we would like to introduce the apparatus and techniques used for photon collection and interference. This type of experiment has some inherent differences to the usual ion trapping experiments – namely, our control flow is branched based on the different herald patterns we can detect. The heralded events are relatively rare compared to the number of attempts, so we must maximise the probability of a successful attempt, and ensure a high repetition rate of attempts. The ability to stop immediately on detection of a herald is necessary to be able to use the resulting entangled state in some way, which in this thesis is simply performing tomography to characterise the entangled state, but in the future will include further operations to purify the entangled state.

4.1 Photon collection

We aim to collect single photons at 422 nm from a strontium ion to use as an entanglement resource. One method of collecting photons from an atom or ion is by coupling to a cavity [SCS⁺12, SCB⁺13, SLM⁺12, KES⁺16, KHR02, KLH⁺04, WWKR07, HDLL⁺08] – this allows excellent collection efficiency, since the photon is preferentially emitted into the cavity mode. The use of cavities is technically challenging, as charge accumulating on the dielectric surfaces causes problems with trapping, so we choose to collect photons in free space. Spherical [SDKB09] or parabolic [CAL⁺17] mirrors and high numerical aperture lenses [RBA⁺17] have been used in vacuum to collect fluorescence from trapped ions or atoms, achieving high collection efficiencies and the potential to couple into optical fibres. As a compromise between simplicity and collection efficiency, we chose to use an NA 0.6 lens outside the vacuum chamber, which covers approximately 10% of the solid angle around an emitter.

4.1.1 Fibre coupling procedure

Single photons from the ion are collected into an optical fibre through the front imaging system. The NA 0.6 main objective¹ is mounted to a 2-axis rotation stage² on top of a 3-axis translation stage³. The working distance from the ion is 17 mm. Cage rods on the objective's mount allow a $f =$

¹Photon Gear Inc., Atom Imager Relay, 422nm, PGI P/N: 16580

²Newport M36

³Newport M562-XYZ

1000 mm cylindrical lens⁴ to be placed in a rotation mount⁵ near (~ 3 mm) to the fibre tip to remove small amounts of astigmatism, thought to be due to uneven strain on the vacuum window, that cannot be eliminated with adjustment of the objective lens. The Z-axis of the translation stage has a stepper motor actuator⁶ allowing us to adjust the focus without touching the stage, improving stability and ease of use. The fibre is mounted in an SM1-threaded kinematic mount⁷. A 3-axis stage with additional piezoelectric actuators⁸ allows fine tuning of the fibre tip position.

In order to achieve maximal coupling of the ion light into the fibre, aberrations in the image of the ion must be minimised. The main objective is a multi-element lens designed for near diffraction-limited performance at 422 nm, compensating for the spherical aberration introduced by a 4 mm vacuum window⁹. The lens must be aligned with its axis normal to the window to properly effect this compensation.

Initial alignment consists of placing an EMCCD camera at the prime focus of the objective lens (figure 4.1a). This is sufficient for coarse alignment, but the pixel size of $16\ \mu\text{m}$ is not small enough to resolve the diffraction limited image, the Airy disk having its first dark ring at $5.7\ \mu\text{m}$ diameter.

To further reduce the aberrations, we must magnify the image of the ion such that the first Airy ring is resolvable. Magnifying the image too much will dim the image by spreading it out over many pixels. We choose to magnify

⁴Thorlabs LJ1516RM-A

⁵Thorlabs LCRM2/M

⁶Newport Conex TRA12CC Actuator

⁷Thorlabs K1T2

⁸Thorlabs Nanomax 300

⁹The window on this system is in fact $1/8''$ (3.18 mm): this difference has a negligible effect on performance.

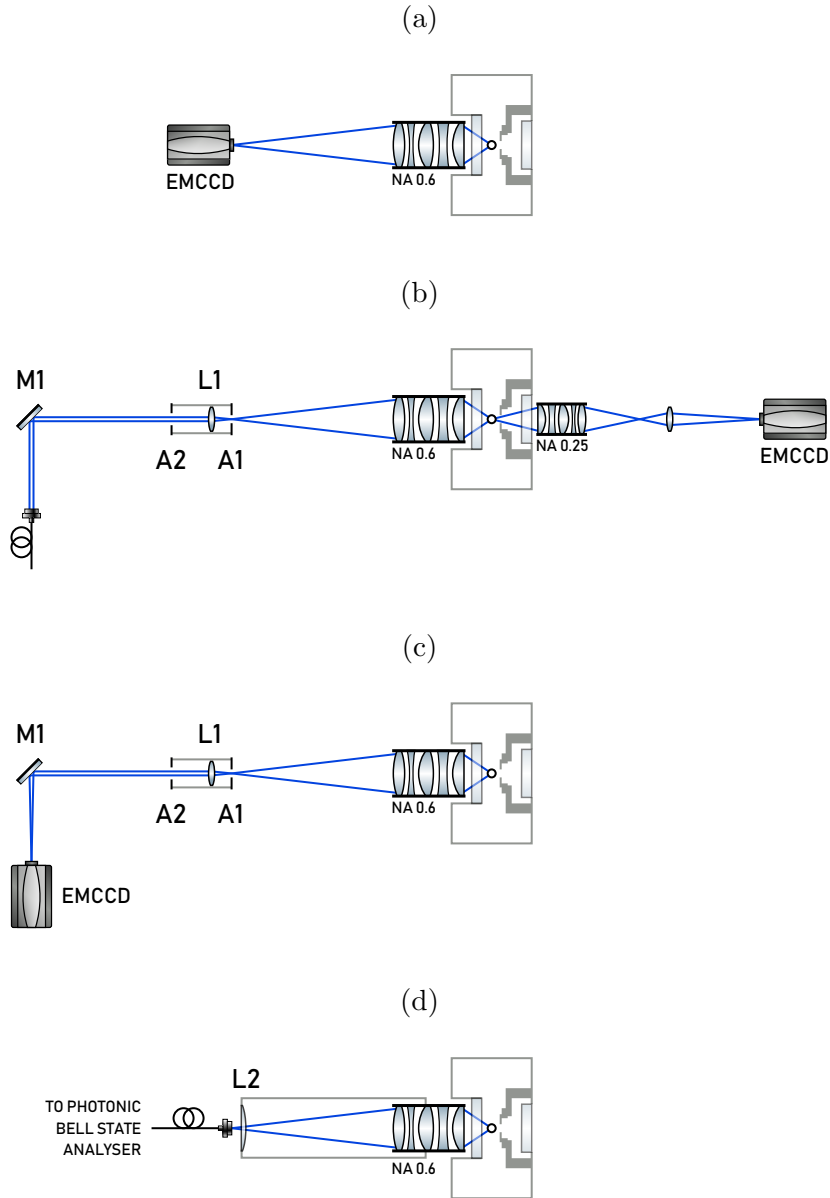


Figure 4.1: Setup steps for high numerical aperture imaging into a fibre. Firstly (a), the EMCCD is placed at the prime focus of the objective, for initial alignment of the NA 0.6 objective. A weak collimated beam is used to align the relay lens approximately (b) by overlapping the image of the fibre with that of the ion in the rear imaging system. During fine alignment (c) the relay lens must be adjusted in $\{x, y, z\}$ to keep it centred and the image focused at the camera. Mirror $M1$ allows the image to be steered onto the camera sensor. The final setup is shown in (d), with a cylindrical lens (mounted with extensions to the objective mount) ~ 3 mm from the fibre tip.

the image by ~ 60 so that the diffraction limited image has a diameter of approximately 20 pixels on the camera.

Magnification is provided by a $f = 8$ mm moulded glass aspheric lens¹⁰. This lens is held using adapters in the same SM1 kinematic mount to be used for the fibre. Irises in the same mount allow accurate positioning of the aspheric lens; the first is at the prime focus to control position¹¹ while the second is mounted on a 50 mm tube after the lens, ensuring that the relay lens axis is parallel to the objective lens axis. Propagating a low power beam of collimated light (figure 4.1b) through this system helps to locate where the ion image will be formed, by aligning the image of the fibre to the same position as the ion on the EMCCD camera in the rear imaging system. Then the fibre collimator can be replaced with the EMCCD, and the relay optic moved slightly closer to the ion to form the 60 times magnified image on the EMCCD.

An iterative process follows for correcting aberrations (figure 4.1c), reducing the dominant type with successive adjustments until another type dominates. Adjustments of the objective cause large changes in the z position of the prime focus; the aluminium base on which the kinematic stages are mounted is left unsecured and slid fore and aft against two reference points. Static friction is sufficient to keep the base in place during the alignment procedure.

Scanning through the focus helps to identify the type of aberration more clearly. Residual astigmatism that could not be corrected with the objective

¹⁰Thorlabs A240TM-A

¹¹The minimum aperture of this iris is larger than the diameter of the focussed spot, so with the iris fully closed the image should not be dimmed.

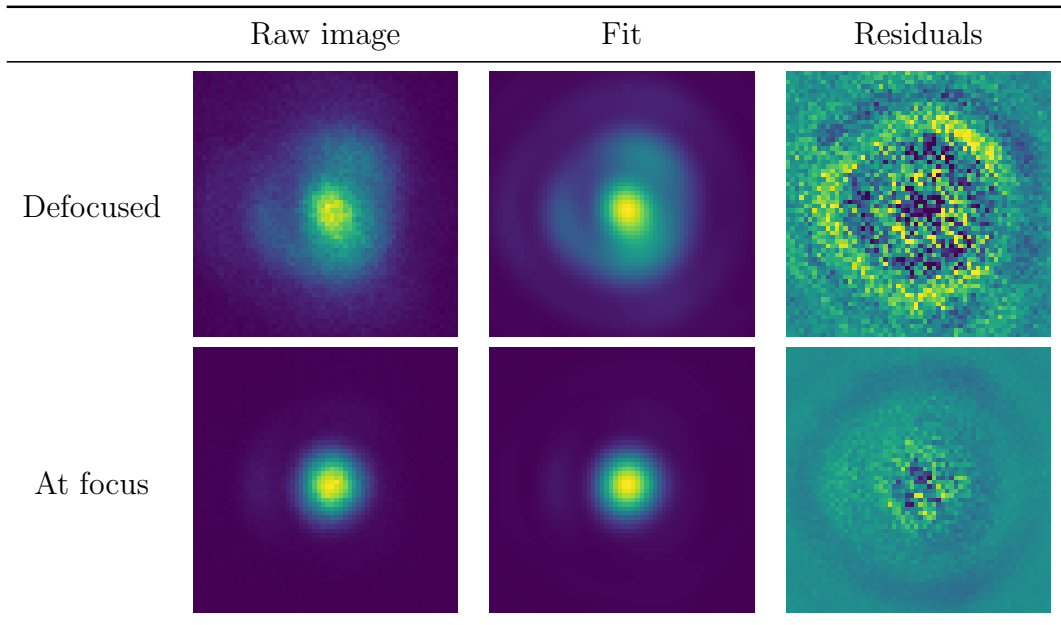


Figure 4.2: Point spread function image, fit and residuals, both for intentionally added defocus and at the focus. The color mapping used for the residuals is magnified ten times compared to the other images, with zero at the midpoint of the mapping. The Zernike coefficients for the defocused fit are shown in figure 4.3.

lens alignment was corrected using a cylindrical lens close to the prime focus. The cylindrical lens ($f = 1000$ mm) is easily aligned by moving it further from the prime focus to increase the introduced astigmatism and aligning its axes such that it is opposite in sign to the uncorrected system, before reducing the distance to the focus. Once aberrations have been visually corrected, we verify the alignment by fitting the point spread function with Zernike polynomials [WCJN⁺16]. Intentional defocus of the image increases the effect of any remaining aberrations, so improving the robustness of the fit. Example images, both with and without defocus are shown in figure 4.2, with the Zernike coefficients for the defocused image shown in figure 4.3.

The final configuration with the fibre collecting light from the ion is shown




Figure 4.3: The Zernike coefficients of the defocused point spread function – 1 ‘wave’ corresponds to a 2π phase shift, all parameters (aside the intentional defocus) are < 0.02 . Fitting the focused image is less robust.

in figure 4.1d.

The detection efficiency here is still (see section 3.7) defined by the probability of detecting a photon given that a photon was emitted; however, in this instance we must sum over the four detectors of the Bell state analyser. Thus the total detection efficiency can be written as

$$\eta_c = \frac{\Omega}{4\pi} \xi_{\max} \chi_{\text{fibre}} T_{\text{opt}} \cdot \frac{1}{4} \sum_i^4 \chi_i^{\text{b.s.a.f.}} Q_E^i, \quad (4.1)$$

where $\frac{\Omega}{4\pi} = 0.1$ is the fraction of solid angle subtended by the lens, $\xi_{\max} = 0.80$ is the maximum fraction that can be coupled into a single mode fibre¹², χ_{fibre} is the fraction of ξ_{\max} that actually enters the photon collection fibre, $T_{\text{opt}} \approx 0.95$ is the transmission through all optical elements¹³. The term $\frac{1}{4} \sum_i^4 \chi_i^{\text{b.s.a.f.}} Q_E^i$ gives the average probability of detecting a photon given its presence at an output of the Bell state analyser, incorporating the quantum efficiency $Q_E^i \approx 65\%$ of each detector, and the coupling efficiency into the single mode fibre preceding the detector $\chi_i^{\text{b.s.a.f.}} \approx 0.90$ (see section 4.2). With perfect ion-fibre coupling efficiency $\chi_{\text{fibre}} = 1$, the detection efficiency would then be 4.2%. The best measured collection efficiency in each system was 2.1% and 2.4% for Alice and Bob respectively, corresponding to estimated ion-fibre couplings of 50% and 57%¹⁴.

¹²This is discussed further in section 6.1, and is dependent on the numerical aperture.

¹³Including the Bell state analyser

¹⁴The discrepancy here is due to the finite patience of a graduate student adjusting micrometers in a darkened room...

4.2 Photonic Bell state analyser

The entanglement swapping procedure to entangle the ions relies on detecting Bell states of photons. In our case, the photons themselves are not actually *in* a Bell state, unlike typical photonics experiments where a pure Bell state is produced (by parametric downconversion for example) and subsequently measured. However, since the Bell states form a basis over the combined states of two two-level systems, we can enact a Bell state measurement to *project* the photons into such a state, destroying the photon state in the process but swapping the entanglement onto the ions. With only linear optics, a Bell state analyser of photons can only distinguish 2 of the 4 Bell states unambiguously [MMWZ96], setting a hard limit on the probability of projecting the photons into a useful state.

The Bell state analyser¹⁵ is a generalisation of the Hong-Ou-Mandel interferometer [HOM87], with the layout shown in figure 4.4. The photons are incident on a 50:50 beam splitter¹⁶; polarising beam splitters¹⁷ follow each output port. The output ports of the polarising beam splitters are monitored by avalanche photodiodes¹⁸ (APDs). Zero-order quarter¹⁹- and half-wave plates²⁰ allow arbitrary manipulation of the photon polarisation basis and are mounted in motorised stages²¹. Mechanical shutters²² allow

¹⁵Built with great help from Joe Goodwin

¹⁶Thorlabs BS10-A

¹⁷Altechna, custom coated

¹⁸COUNT BLUE

¹⁹Thorlabs WPH05Q-405

²⁰Thorlabs WPH05H-405

²¹Thorlabs DDR25/M; Thorlabs K10CR1 were used initially, but these were so slow that during the ion-photon tomography experiment an unacceptable percentage of the experiment time would have been taken up by changing waveplate angles.

²²Thorlabs SHB05T

the light from one of the traps to be blocked while performing single-trap experiments.

The outputs of the beam splitters are coupled into single mode fibres²³, increasing the contrast of the interferometer [SHKW06] by ensuring spatial mode matching of the photons, at the expense of reducing the overall detection efficiency because of imperfect coupling into the fibres and extra optical elements introducing loss. Fixed air-spaced triplet fibre collimators²⁴ and a 1:1 telescope using $f = 250$ mm plano-convex lenses²⁵ were used to mode match back into the output fibres, with fine adjustment of the lens position provided by a zoom mount²⁶ with 4 mm of travel. The fibre coupling was performed with coherent laser light, achieving end-to-end couplings of $\sim 90\%$. The interferometer contrast was not measured with classical light²⁷.

The light from the output fibres is focused onto the avalanche photodiodes' 100 μm active area using a matched achromatic doublet pair²⁸. The lens and fibre can be translated with an XY-mount²⁹ and a zoom mount³⁰ to position the focused spot onto the detector's active area. The detector quantum efficiencies were measured to be in the range 60–65% at 422 nm. The APDs have a typical latency from photon impact to TTL output pulse of 30 ns, with a resolution of 1 ns.

²³Thorlabs S405-XP custom coated: $R_{\text{avg}} < 0.75\%$ over 400–700 nm

²⁴Thorlabs TC06FC-405

²⁵Thorlabs LA1461-A

²⁶Thorlabs SM1ZM; these mounts gave slightly *too* fine adjustments, necessitating extra lens tube as a coarse adjustment. Coarser adjusters with more travel would be a better choice in future.

²⁷This was an oversight in the rush to get results – we assumed that the very high fibre coupling efficiencies would ensure good mode matching, but perhaps we were just lucky!

²⁸Thorlabs MAP051919-A

²⁹Thorlabs LM1XY/M

³⁰Thorlabs SM1ZM




Figure 4.4: Diagram of the photonic Bell state analyser layout. The shallow angle of incidence on the 50:50 beam splitter helps to minimise any polarisation dependent differences in transmission/reflection. The approximate position of the focus between the lenses is marked with a dashed circle. Fibre couplers A and B come from Alice and Bob respectively; the labels on the output couplers match the detector number. Detectors 0 and 3 (1 and 2) measure horizontally (vertically) polarised photons.

	Reflected	Transmitted	AR coating loss	R	T
S-pol	0.4653	0.5211	0.0018	0.4717	0.5283
P-pol	0.4649	0.5258	0.0018	0.4693	0.5307

Table 4.1: Properties of the selected non-polarising beam splitter. The AR coating loss is the power in the *first* reflection only. Note that there are miscellaneous losses of around 1% for both polarisations. In chapter 7 we argue that any losses affect only the rate, so can be ignored – the final two columns show R and T assuming the beam splitter were lossless.

4.2.1 Beam splitters

The Bell state analyser’s effectiveness depends on the qualities of the beam splitters used, as will be discussed in chapter 7 – imperfect beam splitters cause the projective measurements on the photons to no longer be exactly in the Bell basis, and thus the resulting ion state is not a Bell state either. Characterisation of the beam splitters enables us to bound this source of infidelity.

We can minimise the polarisation dependence of the 50:50 beam splitter by using it close to normal incidence; then the only free parameter is the reflectance, which we would like to be as close to 50% as possible. Several nominally identical non-polarising beam splitters from different batches were evaluated for their reflection/transmission balance and anti-reflection coating quality on the non-beam-splitting face. The chosen beam splitter’s properties³¹ are listed in table 4.1.

The polarising beam splitters (PBSs) are used to measure the polarisation of the photons, so their polarisation purity dictates the fidelity of this

³¹The non-polarising beam splitter was not tested for any birefringence; this would affect fidelity but appears to be small.

	PBS1	PBS2
S-pol (reflected)	12 500:1	3000:1
P-pol (transmitted)	700:1	1900:1

Table 4.2: Extinction ratios of the two Altechna PBSs used in the entangler board, as measured in situ. The labels correspond to those given in figure 4.4.

measurement. Cube PBSs from different suppliers were evaluated for best polarisation purity *ex situ*³². The best performing³³ were then mounted in a cube mount³⁴. A Glan-Laser polariser³⁵ in a rotating mount was used to maximise and minimise the power in each arm, and so measure the *in situ* performance. The clamping force significantly affected the beam splitters' performance, necessitating a reduction in the mounting screws' tightness. Final measurements are shown in table 4.2.

4.2.2 Waveplates

The waveplates of the Bell state analyser are used to rotate the photon state, both to perform tomography on an ion-photon state as in chapter 6 and to rotate the photon into an appropriate basis in chapter 7. Characterising the waveplates ensures that we can calculate the required waveplate angles for a given rotation on the Bloch sphere.

The waveplates' retardance and the angle of the fast axis were determined using a polarisation analyser³⁶. A cage mount assembly held a fibre coupler,

³²Wollaston prisms offer excellent extinction in both polarisations (routinely specified at 100 000:1), and were only omitted due to oversight.

³³Supplied by Altechna

³⁴Thorlabs CCM1-4ER/M with BS127CAM adapter

³⁵Thorlabs GLB10, extinction ratio of transmitted beam measured at 45 000:1

³⁶Schäfter and Kirchhoff SK010PA-UV

the motorised waveplate mount and the polarisation analyser colinear to ensure the waveplate was perpendicular to the beam direction. The input light polarisation was measured by removing the waveplate mount entirely, and then replacing it to measure the output state. The 6 input states of $\{|H\rangle, |V\rangle, |A\rangle, |D\rangle, |L\rangle, |R\rangle\}$ form an over-complete basis. Measuring the output states with the mount set to 0° and 45° allows us to unambiguously determine the waveplate's fast axis relative to the mount.

An in situ measurement can also be performed using the polarisers on the analyser with an additional crossed polariser at the input. The action of a generic waveplate of retardance α at an angle β to the coordinate axes is given by a Jones matrix

$$U(\alpha, \beta) = \begin{pmatrix} \cos^2 \beta + e^{-i\alpha} \sin^2 \beta & (1 - e^{-i\alpha}) \cos \beta \sin \beta \\ (1 - e^{-i\alpha}) \cos \beta \sin \beta & e^{-i\alpha} \cos^2 \beta + \sin^2 \beta \end{pmatrix}. \quad (4.2)$$

If the waveplate is between crossed polarisers, we select one of the off diagonal elements; taking the modulus squared to find the transmitted intensity yields

$$I = I_0 \sin^2 2\beta \frac{1 - \cos \alpha}{2}. \quad (4.3)$$

Note that this method cannot distinguish between the fast and slow axes of the waveplates – this must be done separately using the polarisation analyser.

For the in situ measurement, a Glan-Laser polariser is inserted directly after the input fibre coupler, with the waveplates removed. The polariser is rotated until the transmission of laser light through the output polariser is minimal, so that the two polarisers are crossed. Both the transmitted and

	Retardance (rad)	Fast axis (deg)
QWP1	1.364(1)	29.8
QWP2	1.519(2)	165.8
HWP1	2.819(4)	15.4
HWP2	2.957(5)	123.7

Table 4.3: Bell state analyser waveplate retardances and fast axis angles as measured in situ.

reflected power from the output polariser are monitored to enable normalisation of the transmitted to total power. The first waveplate is put into place and rotated while monitoring the transmitted powers. This waveplate can be set so that its axes are aligned with the input polariser, and so it has no effect on the polarisation state; then inserting the second waveplate, the same measurement can proceed. Results from the in situ measurements are summarised in table 4.3.

4.2.3 Optical fibres

The photon’s polarisation state must be faithfully transported from the collection optics to the Bell state analyser. Single-mode (SM) fibres have no engineered birefringence, although mechanical stresses will inevitably introduce some small birefringence. The birefringence, and hence the phase accumulated between orthogonal polarisations, varies on thermal timescales. These variations are small, with the phase not drifting appreciably over several hours, allowing us to perform tomography to characterise the fibre and use the waveplates to undo the effect of the fibre.

Anti-reflection coated fibres³⁷ are used – as explained later in chapter 7,

³⁷Thorlabs S405-XP custom coated: $R_{\text{avg}} < 0.75\%$ over 400–700 nm

the rate of ion-ion entanglement heralds is dependent on the *square* of the single photon detection efficiency. With four fibre facets between ion and detector, uncoated fibres would give 16 % loss in detection efficiency and thus 32 % lower rate. All of the fibres use FC/APC connectors; this lowers the coupling from the ion into the fibre marginally, but avoids any unintentional optical cavities formed by normal incidence on air-glass interfaces.

4.3 Experimental control

In this thesis, we are presenting the results of “heralded” events; that is, we attempt a procedure that has a low probability of success, and only on detection of a certain combination of photon(s) do we stop our attempts and move to the analysis stage. Therefore we must (a) maximise the speed of our decision making process and (b) minimise the time taken to execute our entanglement attempts. The methods to achieve this are described below.

4.3.1 The “entangler core”

ARTIQ offers the flexibility of controlling an FPGA using just-in-time compiled code, however, the overheads introduced make it unsuitable for the fast loop used to attempt the generation of entangled photons, as several tens of nanoseconds latency are introduced to generate output events. It is also difficult to synchronise the two FPGAs controlling Alice and Bob as is necessary for the two-trap experiment. The “entangler core” was the solution to this problem, and is a section of precompiled logic gates on the Kasli used to control the inputs, outputs and decision branching of the entanglement

sequence.

The entangler core has a counter which advances with the system clock. It takes in several values at runtime, which are stored in registers for use in logic operations:

- on and off times of TTL outputs, namely the TTLs controlling the 422 nm and 1092 nm state preparation beam RF switches (see section 4.3.2) and the pulse picker trigger (see section 3.5.2)
- open and close time of the pulsed laser photodiode input gate
- open and close *offsets* of the APD input gates – note each APD has its own offset to allow for different latencies, both in the APD itself and the beam path to it
- a list of the herald patterns on which to exit, each stored as a 4-bit one-hot encoded number with one bit per detector
- the time at which to restart the cycle if no herald is detected
- The total attempt duration before forcing an exit due to a timeout

Each time the counter advances, a series of comparisons are made. The counter is compared for equality with the TTL on (off) time and the outputs set high (low) correspondingly. The counter is compared with the open and close time of the pulsed laser photodiode input gate and used to set a flag to accept a rising edge. If a rising edge is detected during this acceptance window, the counter value at the input time is recorded. This counter value is then added to each of the APD gate offsets to give the APD input gate

open and close times, which similarly set a flag per detector and record the counter value on a detection event. If the pattern of APD flags matches any of the herald patterns, a ‘stop’ flag is set, the matched herald pattern and the counter value of the last input event stored. Once the counter is equal to the cycle time, the stop flag is checked – if set, the entangler loop exits with the reason for stopping (*i.e.* the herald pattern) and counter value of the last detection event. If the stop flag is not set, the counter is reset to zero and the cycle starts again. A second counter is not reset at the end of the sequence, but compared with the total attempt duration. If this second counter exceeds the total attempt duration, the stop flag is set. The ‘herald pattern’ stored in this case is a value reserved for indicating timeouts.

On exit of the entangler core, it generates an ARTIQ input event, with the timestamp at the exit time of the core and the herald pattern that caused the exit. Different analysis sequences can then be chosen in the ARTIQ experiment based on the herald pattern recorded. The APD input time can be retrieved later in the experiment, and is used for setting the phase of the 674 nm analysis pulses.

For the two trap experiments, the master (Alice’s core Kasli) and slave (Bob’s Kasli) must be synchronised to begin the entanglement sequence simultaneously. The two devices are connected with a ribbon cable, to allow each device to set a ready flag that the other device can monitor. When both ready flags are set, the entanglement sequence can begin.

4.3.2 AOM switching latency

A typical atomic physics experiment consists of control pulses which trigger events, such as a laser beam switching on, with some latency associated with that action. To make a sequence occur as quickly as possible, the latencies can be measured so that the trigger pulses occur with as little delay between them as desired. Latencies that occur at the beginning of an experiment run can usually be “pre-triggered”, *i.e.* the trigger pulses are sent *before* the end of the previous experiment. This experiment is designed to entangle ions across a network, an inherently probabilistic process with a relatively low chance of success. We must be able to stop on detection of entanglement, which rules out pre-triggering as a method to reduce the cycle time, so latencies must be reduced wherever possible.

In this experiment, we use AOMs to switch the laser beams; switching with AOMs has greater latency than with electro-optic modulators, but typically has a superior extinction ratio, lower insertion loss, and a lower setup complexity.

The time-critical section of the entanglement generation experiment involves only two beam paths³⁸ – the 422 nm σ -polarised and 1092 nm beams used for state preparation. The AOMs for these beam paths are situated in the server rack, controlled by the slave Kasli and Urukuls. Controlling the Urukul’s RF switches via the slave Kasli and fibre DRTIO link introduces several hundred ns latency, so dedicated RF switches controlled by TTL channels on the core Kasli are used for these beam paths.

³⁸The 372 ns latency between the pulse picker trigger and transmitted pulse is not significant, as there is sufficient state preparation time before the pulse is needed.

With the dedicated RF switches, an initial latency of about $1\ \mu\text{s}$ was measured for the 422 nm beam. The slave server rack is situated around 8 m from both the core server rack and the experiment table. A ~ 10 m BNC cable connects to the TTL input of the RF switch, then ~ 7 m of SMA cabling³⁹ within the rack carries the RF signal from the switch to the amplifier and then to the AOM. Finally, a 17 m optical fibre⁴⁰ takes the light to the trap delivery optics. This means that a rough estimate for unavoidable signal latency (from the propagation of electrical signals in cables and light in optical fibres) accounts for around 170 ns of the total $1\ \mu\text{s}$ latency.

Clearly there is an additional latency source in play; there is a delay in between the RF signal appearing at the transducer of the AOM and light being diffracted off the acoustic wavefront, which propagates at $4.2\ \text{mm}\ \mu\text{s}^{-1}$ in the TeO_2 AOM crystal. By realigning the AOM with the beam focus position much closer to the transducer, we were able to drastically reduce the latency of the 422 nm beam paths. The 1092 nm beam paths use fibre AOMs, so this optimisation was not possible. The measured latencies after optimisation are shown in table 4.4.

³⁹A suprising amount of length is required to route the cables around the rack, although there is some slack

⁴⁰This was ordered before the final layout of the experiment was decided, so includes several metres spare.

	On (ns)	Off (ns)
Alice 1092 nm	440	431
Alice 422 nm	309	304
Bob 1092 nm	380	372
Bob 422 nm	288	283

Table 4.4: On and off latencies for each of the time-critical beams in both traps. Since the rise time of the AOM is typically tens of nanoseconds, the times quote here are from the control pulse edge to 50 % beam intensity. The optimisation of the 422 nm latencies is discussed in section 4.3.2.

Chapter 5

Ion state manipulation and readout

In this chapter we present the initial characterisation work on single strontium ions that is a necessary precursor to both the ion-photon and ion-ion entangling experiments. This is focused on the ion manipulations discussed in section 2.1; initialisation, rotations and readout of the qubit. Preparation of the excited state that produces entangled photons is covered in chapter 6.

5.1 State preparation

To initialise the Zeeman qubit, we optically pump the ion into $|\downarrow\rangle$ using σ^- -polarised 422 nm light, as outlined in section 2.1.3. The fidelity of the state preparation is directly set by the polarisation purity of the 422 nm beam; a Glan-Taylor polariser¹ before a quarter-wave plate ensures control over

¹Thorlabs GLB10-A, extinction ratio measured at 50 000:1

the polarisation. The state preparation beam and the applied field must be colinear to create pure σ^- polarisation.

The state preparation beam's direction is constrained by the single photon collection scheme. The beam and field must be colinear, while the field must be perpendicular to the imaging axis; The imaging axis is constrained by the need to minimise aberrations from the recessed vacuum window, and so is colinear with the normal to the window, but this is unknown a priori. The best we can do is to assume the window is well registered to the vacuum system, and fix the beam direction relative to the vacuum system². The beam direction is set with the use of irises fixed relative to the vacuum chamber (see section 3.6.1).

To set the polarisation and field, the ion is alternately cooled using the usual Doppler cooling beams (422 nm and 1092 nm) and with the σ^- -polarised 422 nm beam and 1092 nm repumper. The ion fluorescence is monitored during application of the σ^- beam, which is set to a high intensity ($\sim 100I_S$, where I_S is the saturation intensity); if the field and polarisation are correctly set then the ion will be pumped into $|\downarrow\rangle$ and cease to fluoresce. This cessation occurs with a time constant equal to the $5p^2P_{1/2}$ state lifetime (7.39 ns) multiplied by the inverse of the branching ratio from $5p^2P_{1/2,-1/2}$ to the dark state $5s^2S_{1/2,-1/2}$ ($1/3$), approximately 25 ns. Monitoring the fluorescence over a timescale much longer than this (*i.e.* 15 μ s) renders the transient fluorescence as the ion is pumped into the dark state insignificant over the collection window. The increased power allows polarisation impurities in the

²This is a reasonable assumption – a 0.5 mm difference in vacuum gasket compression would be easily visible over the 152 mm flange diameter, corresponding to less than 0.2° misalignment.

beam to have a greater effect, since the main transition is saturated. By adjusting the waveplate rotation, tilt and trim coil currents, we can minimise the fluorescence, and so maximise the state preparation fidelity.

For normal operation, the intensity is set to around $0.5I_S$ and the σ^- -polarised beam applied for some 15 μs . Faster state preparation is achievable with higher intensities, at the cost of decreased preparation fidelity (see section 6.1).

5.2 Readout in the computational basis

To read out the computational (Z) basis of the Zeeman qubit, we transfer population from $|\uparrow\rangle$ out of the Doppler cooling cycle to $|D\rangle$, so that when the Doppler cooling lasers are applied the ion only fluoresces if the ion remains in $|\downarrow\rangle$. The transfer is achieved by applying a π -pulse from $5s^2S_{1/2,+1/2} \leftrightarrow 4d^2D_{5/2,-3/2}$ with the 674 nm laser; the narrow linewidth of the transition makes off-resonant excitation of $|\downarrow\rangle$ negligible. After transferring population to $D_{5/2}$, the 422 nm and 1092 nm lasers are applied and fluorescence photons counted for 250 μs . At maximal fluorescence (called the ‘bright’ state), the count rate in both traps is typically 80–100 kHz. The count rate with the ion in a non-fluorescing (‘dark’) state³ is < 1 kHz.

Photon count histograms for the bright and dark states are collected, shown in figure 5.1. This is done independently from the 674 nm laser operation to avoid the contribution of shelving errors. To measure the bright state,

³Incorporating the PMT’s intrinsic dark rate, laser scatter and room light; imaging through the slot in the trap substrate substantially cuts down the laser scatter, so this is mostly room lighting.

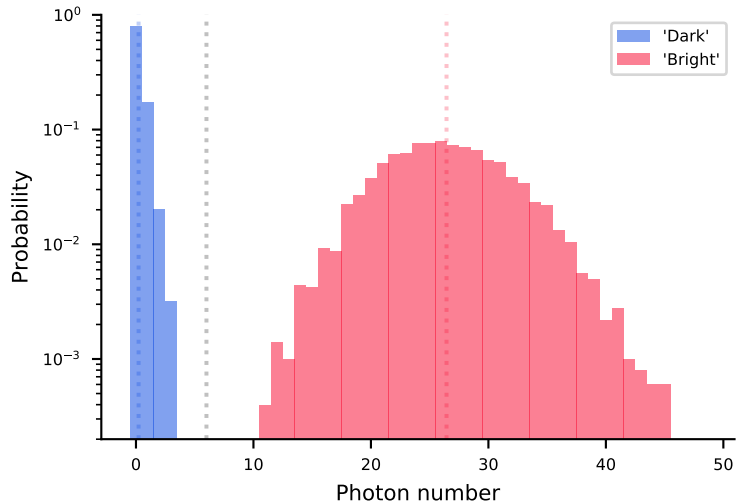


Figure 5.1: Dark (blue) and bright (red) readout histograms obtained with 5000 shots of $250\ \mu\text{s}$ windows, normalised to unit probability. The dashed blue and red lines show the mean dark state count $n_d = 0.2$ and bright count $n_b = 26.4$ respectively, and the dashed grey line shows the chosen threshold value of 6. (This data is from the ‘Alice’ system.)

photons are counted for $250\ \mu\text{s}$ with both the 422 nm and 1092 nm lasers illuminating the ion. For the dark state, only the 422 nm beam is used, so that the ion is pumped into $D_{3/2}$ and becomes dark⁴.

We wish to identify a threshold value n_{th} such that, for a number of photons n counted during the collection window, the ion will be designated ‘bright’ if $n \geq n_{th}$ and ‘dark’ if $n < n_{th}$. Assuming that the number of photons collected in each case follows a Poisson distribution, with n_b the mean bright count and n_d the mean dark count, the threshold value is simply the lowest value at which the probability mass function with mean n_b exceeds

⁴Some photons will of course be scattered by the ion during this process, however, since the collection efficiency is only 0.4% the expected number of these photons counted from this is $\ll 1$.

that of mean n_d ;

$$n_{th} = \left\lceil \frac{n_b - n_d}{\ln(n_b/n_d)} \right\rceil. \quad (5.1)$$

Errors in discrimination between the two states in this model are purely statistical, due to the overlap of the two Poisson distributions [Bur10]; for the data in figure 5.1, this error is 3×10^{-7} .

The assumption of a Poisson distribution only holds if the ion does not transfer between the two states during readout, but the dark state $4d^2D_{5/2}$ has a finite decay lifetime to the bright state of 435 ms. This introduces a second source of error, the full derivation of which can be seen in refs. [Bur10, KGA⁺11]. Calculating the discrimination error (using the same threshold as before) gives $\epsilon_b = 4.3 \times 10^{-7}$ for the bright state and $\epsilon_d = 5 \times 10^{-4}$ for the dark state, and an average error $\epsilon \equiv (\epsilon_b + \epsilon_d)/2 = 2.5 \times 10^{-4}$. The asymmetry between the calculated errors for the bright and dark states is because the threshold used does not account for the decays from the dark state, and so is not completely optimal. However, the threshold is close to optimal, and the discrimination error is negligible compared to the readout error, so the minor correction needed is neglected here.

An estimate of the shelving error can be deduced from the Rabi oscillations on the optical qubit shown in figure 5.2, although these are on the “main” rather than “map” transition (see section 2.1.4 for the definition of these terms). The frequency and duration of the 674 nm shelving π -pulse is calibrated regularly during operation, so shelving errors are typically 4×10^{-3} , and dominate the readout error.

After readout is complete, the $D_{5/2}$ level is cleared back to the $S_{1/2}$ mani-




Figure 5.2: Rabi oscillations on the $|\downarrow\rangle \leftrightarrow |D\rangle$ transition. The population in $|\downarrow\rangle$ is plotted against the duration of the 674 nm laser pulse. The π -pulse error is 3.6×10^{-3} .

fold with the 1033 nm laser.

5.3 Readout in other bases

In order to perform state tomography of the ion, we must be able to measure the ion qubit in bases other than the computational bases. We first map the Zeeman qubit into the optical qubit with a π -pulse on the $|\uparrow\rangle \leftrightarrow |D\rangle$ transition (the “map” transition), and then a $\pi/2$ -pulse between $|\downarrow\rangle$ and $|D\rangle$ (the “main” transition) rotates the qubit state around an axis on the equator of the Bloch sphere. The axis of the measurement is dependent on the phase relationship between the map π -pulse and the main $\pi/2$ -pulse – we define the X basis such that the phase difference is zero, and the Y basis with a phase difference of $\pi/2$. Then the 422 nm and 1092 nm lasers are applied and fluorescence photons counted to exact a measurement, as with Z basis

readout.

Achieving phase referencing is trivial within the ARTIQ framework; the phase of the Urukul DDSs can be directly controlled within the experiment program. During the experiment, the DDS phase offset at a given reference time (which may be in the past or future relative to the current time) can be set. The phase is adjusted before turning on the RF switch controlling the output, so there are no discontinuities in the RF applied to the AOM.

Rabi oscillations of the optical qubit are shown in figure 5.2; the contrast of these oscillations is 99.6%. For this data, population is prepared first to $|\downarrow\rangle$, so does not include the Zeeman qubit shelving error from the map π -pulse, but does include state preparation errors and discrimination errors.

Readout of the Zeeman qubit in the X and Y bases is subject to the shelving π -pulse error and the subsequent $\pi/2$ -pulse error; this is then approximately 5.4×10^{-3} .

5.4 Errors in preparation and readout

Errors in ion state preparation and readout cannot be separated for the readout scheme used, since the readout does not form a quantum non-demolition measurement [HAB⁺14]. Characterisation of the coherent manipulations can in principle be corrected for state preparation and readout errors; a thorough error analysis was not performed, since we expect magnetic field fluctuations to dominate our errors as the field is not stabilised.

Chapter 6

Ion-photon entanglement

In any envisioned application of quantum communication technologies, eventually the need to share quantum information between parties separated by some distance will arise. Sending a single qubit over a quantum channel can be used for quantum key distribution, as in the BB84 [BB84] protocol, but more sophisticated schemes like E91 [Eke91] rely on sharing *entanglement*. Indeed, while sending a single flying qubit a significant distance is possible, at some point the losses will become unacceptable, and to transmit the quantum state further a ‘quantum repeater’ is needed [SSdRG11]. The quantum repeater makes use of entangled pairs to teleport a quantum state from one place to another.

Although electrons have been used as flying qubits over short distances [YTB⁺12], photons remain the obvious choice for a long-distance transport qubit, as they interact only weakly with the environment. Photons have been used to faithfully transmit entanglement over extreme distances both in optical fibres [IMT⁺13] and in free space [UTSM⁺07, YRL⁺12], even using a satel-

lite to achieve a separation of 1200 km [YCL⁺17]. Photons have several degrees of freedom that can be used to encode quantum information, including polarisation [FC72, BPM⁺97], time-bin [Fra89], orbital angular momentum [MVWZ01] and spatial modes [HSZ89].

Arguably the easiest degree of freedom of the photon to manipulate is the polarisation degree of freedom – successive waveplates can be used to arbitrarily rotate the polarisation state, while polarising beam splitters can be used to transfer orthogonal polarisations into different spatial modes and so direct them towards detectors. The spontaneous decay of an excited electronic state of an ion provides a natural pair of entangled qubits, one stationary (the trapped ion), and one flying (the photon). If there are multiple possible decay channels, then with a suitable choice of excitation scheme the polarisation degree of freedom of the photon will be entangled with the ion.

Leaving aside our eventual goal of entangling pairs of ion qubits in separate traps, the production of a high-quality Bell pair of this nature is useful in and of itself; it can be used to perform a “blind” quantum computation. This is where a simple client (simple because its only capability is measuring the polarisation encoded photons in different bases) can perform a computation on a complex quantum server (with the capability to perform single qubit and two qubit entangling gates) using shared entangled pairs and classical communications only [MF13]. Crucially, the server gains no knowledge of the algorithm that was used, the input or the output values.

In this chapter, we describe the theory of producing an entangled ion-photon pair, present our experimental procedure and the measurement of entanglement, and conclude with a discussion of error sources.

6.1 Collecting photons from ions

An atom in an excited level will undergo spontaneous decay, in general giving a state of the form

$$|\psi\rangle = \sum_{i,j} c_{ij} |S_i\rangle |\mathbf{k}_{ij}\rangle, \quad (6.1)$$

where we sum over all the excited and ground states. The atom decays to a superposition of states $|S_i\rangle$ of the ground level, each associated with a photon mode $|\mathbf{k}_{ij}\rangle$ whose polarisation, spatial amplitude and frequency depends on the transition from the excited state. The coefficients c_i are the Clebsch-Gordan coefficients linking the ground states to the excited state.

For some choice of pure excited state, the detection of a given photon mode projects the ion into a known state. This destroys the entanglement between ion and photon (because the photon has been measured) but allows us to post-select events where an entangled ion-photon pair was present¹.

The detection of the photon is inherently probabilistic – there are no existing detectors with 100 % quantum efficiency². For our ion-photon pairs to be useful, we must maximise the rate at which we can create them, so that the application qubits (for example, the memory qubits in a blind quantum server) do not decohere over the time it takes to generate entanglement. We must maximise the rate at which we attempt entanglement generation, as well as maximising the collection efficiency.

Multiple excitations of the ion to the upper level are unacceptable; if we detect the first of two photons emitted by the ion, its state will be uncorre-

¹Neglecting dark counts in the detector, obviously.

²Yet! We live in eternal hope...

lated with the final state of the ion.

For ion-photon entanglement, the photon can be polarisation [BMDM04], or frequency [Hay12] encoded. We have already asserted that polarisation-encoded photons are easier to manipulate than frequency-encoded, so we will look only at these methods. We must have a one-to-one mapping between the polarisation of the photon and the qubit state, so there should be only two detectable decay pathways, possibly after filtering polarisation states of the photon – two such example decay schemes are shown in figures 6.1a and 6.1b. Both of these schemes produce photons from σ^\pm and π decays. Trivially, the photons from π decays do not propagate along the quantisation axis, so for these schemes collecting along this axis seems like the obvious choice. However, to gather any photons, we must collect over some finite solid angle, which would mean some mixing of the unwanted π emission, decreasing the purity of the resulting state. Happily for these cases, symmetry considerations mean that the π emission does not couple into the TEM-00 Gaussian mode of an optical fibre [LHM⁺09], thus performing the required filtering, and so the ion-photon state is maximally entangled after filtering.

We elect to use a third scheme presented in figure 6.1c. We coherently excite into only the $P_{1/2, 1/2}$ state, so that the combined state of the ion and photon after emission is

$$|\psi\rangle = \sqrt{\frac{2}{3}} |\downarrow\rangle |\mathbf{k}_\sigma\rangle + \sqrt{\frac{1}{3}} |\uparrow\rangle |\mathbf{k}_\pi\rangle, \quad (6.2)$$

where we use $\{\sigma, \pi\}$ to denote the different decays. The (unnormalised)

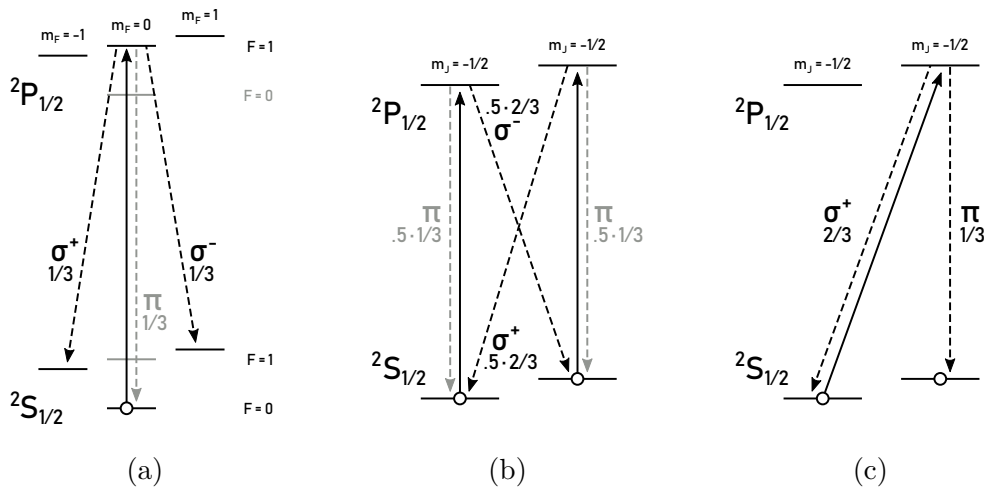


Figure 6.1: Three polarisation-encoded photon production schemes with polarisations and relative branching ratios shown. A nuclear spin-1/2 atom can be used for the scheme in (a) as in [VHI⁺14, HIV⁺15]. In (b), a superposition of the ground state Zeeman levels of a nuclear spin-0 atom is excited coherently to the upper levels, before decaying. For (c), the ion is excited to an upper level before decaying back to the $S_{1/2}$ manifold, with photons collected *perpendicular* to the quantisation axis. The lower states can be a hyperfine states of the ground level as in [BMDM04], or a Zeeman qubit as in this work.

polarisation states of the two photon modes are

$$|\mathbf{k}_\pi\rangle = -\sin\vartheta \left| \hat{\vartheta} \right\rangle \quad (6.3)$$

$$|\mathbf{k}_\sigma\rangle = \frac{e^{i\varphi}}{\sqrt{2}} \left(\cos\vartheta \left| \hat{\vartheta} \right\rangle + i \left| \hat{\varphi} \right\rangle \right), \quad (6.4)$$

where ϑ and φ are the polar and azimuthal axes of the photon with the magnetic field axis, and $\hat{\vartheta}$ and $\hat{\varphi}$ are the respective unit vectors. If we collect along an axis perpendicular to the quantisation axis (*i.e.* $\vartheta = \pi/2$, φ arbitrary) then the σ and π decays have orthogonal polarisation, and the $1/\sqrt{2}$ prefactor in the $|\mathbf{k}_\sigma\rangle$ dipole emission pattern exactly cancels the difference in Clebsch-Gordan coefficients. Again, however, we must consider that we will collect over some finite solid angle. The photons from σ decays have some elliptical polarisation as we move away from the equator defined by the quantisation axis, and so are not orthogonal to the linearly-polarised photons from π decays. This will cause polarisation mixing if detecting in free space [MMO⁺07]. Once more, single mode fibres rescue us from a seemingly impossible situation; the non-orthogonal part of the σ decay does not couple into an optical fibre aligned orthogonal to the quantisation axis [Nad18]. The full calculation is non-trivial, but it should be obvious³ that about the equator $\vartheta = \pi/2$, $\cos\vartheta$ is an odd function, and so when integrated over the (even) Gaussian fibre mode, the integral associated with the $\hat{\vartheta}$ term is zero for the σ decay. Therefore once coupled into the fibre, the σ decay retains only the $\hat{\varphi}$ term and so the two emission modes are both linear and orthogonal in the

³With the benefit of hindsight, almost anything can be said to have been obvious. . .

fibre⁴. The σ emission also retains the prefactor of $1/\sqrt{2}$.

After collection into the fibre, the ion-photon state is [Nad18]

$$|\psi\rangle = \frac{1}{\sqrt{2}} \{ |\downarrow\rangle |H\rangle + |\uparrow\rangle |V\rangle \}, \quad (6.5)$$

where we have mapped $|\mathbf{k}_\pi\rangle \rightarrow |\hat{\vartheta}\rangle \rightarrow |V\rangle$ and $|\mathbf{k}_\sigma\rangle \rightarrow |\hat{\varphi}\rangle \rightarrow |H\rangle$ as detailed above, with $|V\rangle \parallel B$ and $|H\rangle \perp B$ ⁵. Trivially we can see that (6.5) describes a maximally entangled state, as desired.

The rate of detecting the photons will depend on the rate and probability with which we prepare the excited state, as well as the end-to-end collection efficiency of our fibre-coupled detection system. The per-attempt probability p_1 of detecting a single photon in *any* detector of the Bell state analyser (see section 4.2) is

$$p_1 = \eta_c \cdot F_{\text{s.p.}} F_{\text{exc}} B_{S_{1/2}}, \quad (6.6)$$

where $F_{\text{s.p.}}$ and F_{exc} are the fidelities of ground state preparation to $|\downarrow\rangle$ and excitation to $P_{1/2}$ respectively, and $B_{S_{1/2}} = 0.9449(5)$ [LTGG16] is the branching ratio from $P_{1/2}$ to $S_{1/2}$. The total detection efficiency η_c , defined in equation 4.1, is 2.0% and 2.4% for Alice and Bob respectively. The rate is then simply $R_1 = \bar{R}_a \cdot p_1$, where \bar{R}_a is the average attempt rate.

⁴It may appear that we have somehow collected more of the emission by collecting both σ and π – it must be remembered that the amplitude of the σ emission is reduced along our choice of axis, so that both the schemes outlined in figures 6.1b and 6.1c collect the same amount of photons regardless.

⁵Of course in the lab, the quantisation axes are actually at $\pm 45^\circ$ to the vertical, just to make things easier.

6.2 Ion-photon entanglement

We control the attempted creation of ion-photon entangled pairs using the “entangler core” described in detail in chapter 4. This preprogrammed set of logic gates allows reduction of the cycle time to $1\ \mu\text{s}$, by controlling the small subset of TTL inputs and outputs required during the photon production sequence. The fast sequence, shown in figure 6.2, consists of:

- state preparation to $|\downarrow\rangle$,
- pulsed excitation to $P_{1/2}$,
- the photon detection window,
- decision branching, conditioned on the detection of a photon.

Note that there is no laser cooling within the fast sequence; the ion is Doppler cooled before entering the loop. The sequence is executed up to 500 times if no photon is detected, after which $100\ \mu\text{s}$ Doppler cooling is performed and the loop is reentered.

Typically optical pumping to the state $|\downarrow\rangle$ (see section 5.1) uses a low power ($0.5 \cdot I_S$, where I_S is the saturation intensity) for the σ^- -polarised 422 nm beam to avoid off resonant excitations and ensure maximum state preparation fidelity. However, the duration of this preparation is $\sim 15\ \mu\text{s}$ which would severely limit our repetition rate, so we increase the power of the σ -polarised 422 nm beam at the cost of an increased population in the other qubit state. At this increased power ($\sim 10 \cdot I_S$), 300 ns optical pumping is used. Errors in state preparation and excitation lead to no photons being

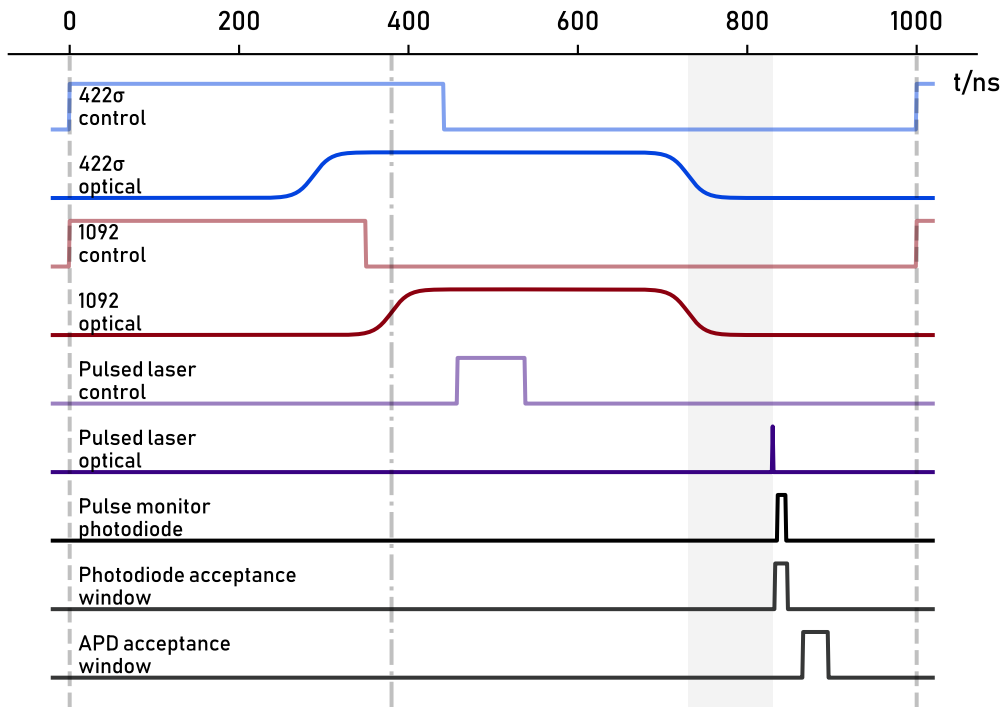



Figure 6.2: Graphical representation of the timeline of control and optical pulses at the ion, to scale, using timings from Bob. Around 40% of the sequence is occupied by latency (everything up to the vertical dot-dash line), and a significant amount of padding is introduced between the nominal off time of the state preparation and the excitation (shaded grey) to ensure that the state preparation beams are fully extinguished. The time after the APD acceptance window closes is used to decide if a photon event (in the ion-photon experiment, a single photon at the given detector, or in ion-ion experiments, a valid herald pattern) was detected, in which case at the end of the 1 μ s long sequence the entangler core exits. The sequence is repeated for up to 500 μ s if no photons are detected, after which the ion is recooled with 100 μ s of Doppler cooling.

generated, so they do not affect the ion-photon fidelity. The population remaining in $|\uparrow\rangle$ is not excited to $P_{1/2}$ since the σ^+ -polarised pulsed excitation beam does not couple to it. This does have an effect on rate though; the probability of exciting to the upper level is directly proportional to the state preparation fidelity. After 300 ns the state $|\downarrow\rangle$ is prepared with fidelity $F_{\text{s.p.}} \approx 0.99$, as shown in figure 6.3.

Note that, given a sufficiently well polarised excitation beam, there can be some trade off between the state preparation duration and fidelity to optimise the excitation rate – if we can reduce the *total* cycle time by 10%, then provided the reduction in fidelity is less than this we will expect to detect photons at a greater rate⁶. At present the cycle time is limited by the thermal limit of the pulse-picker, so there are no gains to be made, but after subtracting the initial latency the state preparation takes up around 50% of the cycle, meaning this could give a few percent speed-up in future iterations.

The ion is excited from $|\downarrow\rangle$ to $|P_{1/2}, m_J = 1/2\rangle$ using the σ^+ -polarised 422 nm ps-pulsed laser. As discussed above, the polarisation impurity in this beam affects only the entanglement rate to first order – the fidelity is decreased only on erroneous excitations to $|P_{1/2}, m_J = -1/2\rangle$. Since the duration of the ultrafast pulse is fixed by the laser itself, we vary the Rabi frequency of the interaction with the ion by varying the pulse *energy* [Nad17]. Practically, this means varying the drive amplitude of the AOM before the fibre to the trap. The population transfer is measured by preparing the state $|\downarrow\rangle$, applying the ultrafast pulse and waiting for the decay to the ground state.

⁶For the two-photon detection events, the reduction in cycle time must be at least twice the reduction in preparation fidelity, as we are dependent on p_1^2 .



(a) State preparation error against optical pumping duration, with latency subtracted. After 300–350 ns, the state preparation error is typically 1–2% – the second dotted vertical line is at 350 ns.

(b) State preparation error against frequency detuning of the σ^- -polarised 422 nm optical pumping beam. The detuning is measured relative to the fluorescence peak of the π -polarised Doppler cooling beam.

Figure 6.3: Assessment of the state preparation to $S_{1/2, m = -1/2}$; the ion is Doppler cooled, then the state preparation beams applied. Population in $|\downarrow\rangle$ is then transferred to $|D\rangle$ with a π -pulse on the “main” 674 nm transition, then the cooling beams applied to readout the ion state. These plots show the population *remaining* in the $S_{1/2}$ manifold, due to both the state preparation and π -pulse error. With no state preparation, the expected population in $|\uparrow\rangle$ is $1/3$ because of the polarisation of the Doppler cooling lasers.

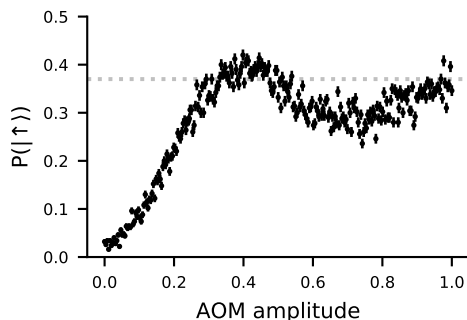


Figure 6.4: Population in $|\uparrow\rangle$ following a pulsed excitation and spontaneous decay, plotted against the RF drive amplitude of the pulsed laser AOM. The RF drive amplitude is proportional to the pulse energy (up to the linearity of the RF amplifier and AOM response). Population in $|\downarrow\rangle$ is shelved to $|D\rangle$ before fluorescence measurement. The 5.5% of population that decays from $P_{1/2}$ to $D_{3/2}$ is within the Doppler cooling cycle and will fluoresce, so the expected bright fraction for perfect transfer to $|P_{1/2}, m_J = 1/2\rangle$ is $1/3 \cdot 0.945 + 0.055 = 0.370$, shown by the dashed horizontal line. The bright fraction may exceed this with imperfect $|\downarrow\rangle$ state preparation and shelving errors. Points in this graph also lie above the line because of shot noise.

No emitted photons need be measured; any population in $|\uparrow\rangle$ means that an excitation occurred⁷. The outcome is shown in figure 6.4, and is consistent with $F_{\text{exc}} \approx 1$ at an AOM amplitude of 0.4. However, beam pointing drift (causing intensity variations as the ion is not at the waist of the excitation beam) means that typically $F_{\text{exc}} \approx 0.95$.

The emitted photons from spontaneous decay are collected into the fibre and directed to the Bell state analyser (see figure 4.4) for tomography. A detector following a polarising beam splitter enacts a polarisation measurement on the photon; rotating the photon's polarisation with the $\lambda/2$ and $\lambda/4$ waveplates allows this measurement to be made for an arbitrary polarisation⁸. Multiple sets of waveplate rotation angles measure the photons in

⁷Neglecting state preparation error

⁸We could in principle use all four of the Bell state analyser's detectors to maximise the

bases spanning the Bloch sphere⁹, with different rotations applied to the ion to perform standard two-qubit tomography.

The density matrix of the ion-photon state is reconstructed via maximum-likelihood estimation (MLE)¹⁰. The polarisation state of the photon is subject to a unitary transformation in the fibre due to its birefringence, meaning that the photon basis is not aligned with our nominal polarisation measurement basis provided by the PBS; so the density matrix is not obviously maximally entangled. Single qubit rotations are applied in analysis¹¹ to the photon to maximise the fidelity $\mathcal{F} \equiv \langle \Phi^+ | \rho | \Phi^+ \rangle$ to the ideal (maximally entangled) state (6.5). The rotated density matrices are shown in figure 6.5; the fidelity is 97.90(12) % for Alice and 97.70(12) % for Bob.

The maximum-likelihood estimate of the ion-photon state also includes an estimate of the probability of there having been a photon present before the measurement PBS¹² – given that the collection efficiency can be measured independently, we can use this to infer the per-attempt probability of producing a photon. The average per-attempt probability of detecting a photon at a *given* detector was $p = 4.04 \times 10^{-3}$ ($p = 5.71 \times 10^{-3}$) in Alice

rate of events and detect orthogonal photon polarisations instead of only one polarisation. However, the data analysis in this setup is non-trivial, and the rate of photon detection with only one detector is already sufficient for good statistics, so we use the simplest setup. The Bell state analyser also has the complication that photons from an ion the other trap could be detected; shutters on the Bell state analyser allow us to test each trap in turn.

⁹Note that the angles are chosen such that ideal $\lambda/2$ and $\lambda/4$ waveplates would produce measurements of $\{|H\rangle, |V\rangle, |D\rangle, |A\rangle, |R\rangle, |L\rangle\}$. However, the waveplates' design wavelength is 405 nm so the retardance is not exactly π or $\pi/2$, and the fibre adds a (fixed) unitary such that the measurements are neither aligned with the photon basis nor even orthogonal – none of that matters for tomography so long as we span the Bloch sphere adequately.

¹⁰Data analysis performed by David Nadlinger

¹¹Equivalently, what additional waveplate rotations would cancel the fibre unitary to align the photon basis to the measurement basis.

¹²A photon which is rejected on average 50 % of the time by the PBS before the detector

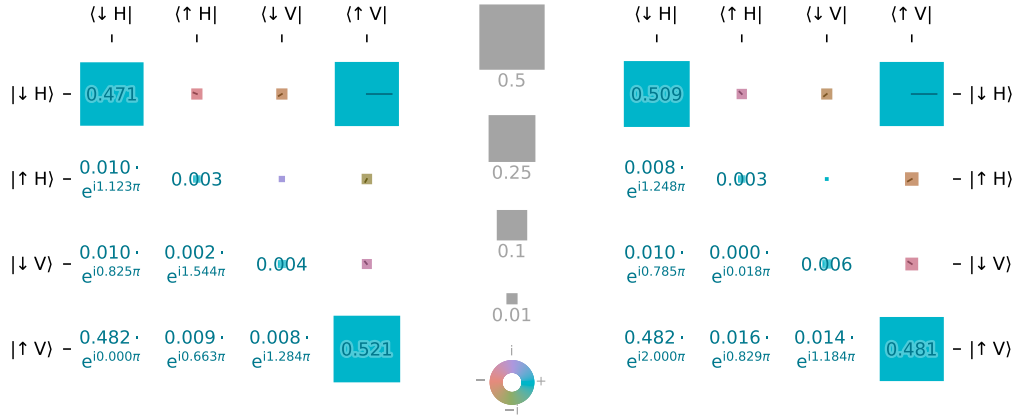


Figure 6.5: Density matrices reconstructed for the ion-photon state in Alice (left) and Bob (right), with additional post factum rotation to maximise the fidelity to $|\Phi^+\rangle$. The *area* of the square shows the magnitude of the matrix element, with the colour representing its complex phase.

(Bob). After n attempts, the probability of at least one photon having been detected is the outcome of a binomial experiment with success probability p , so $P = 1 - (1 - p)^n$. After 500 attempts (*i.e.* before we have performed any recooling), $P = 0.87$ ($P = 0.94$) for Alice (Bob) and so the average attempt rate \bar{R}_a approaches the 1 MHz repetition rate of the fast loop¹³. The average rate at which entangled photons are detected is then $3.98 \times 10^3 \text{ s}^{-1}$ ($5.68 \times 10^3 \text{ s}^{-1}$) for Alice (Bob). The disparity in performance between the two systems is due only to the difference in collection efficiency noted in chapter 4.

Taking into account that only one of the four detectors is used for the tomography experiment, the detection probability measured was $p = p_1/4$,

¹³Demonstrating for Alice, the worse performing of the two: we neglect cases where more than one recooling cycle occurs (1.7%) so that the worst-case attempt rate during the second attempt cycle is $(2 \cdot 500 \mu\text{s} + 100 \mu\text{s})/1000$ attempts = 0.909 MHz. The overall average attempt rate is then roughly $0.87 \cdot 1 \text{ MHz} + 0.13 \cdot 0.909 \text{ MHz} = 0.988 \text{ MHz}$.

and so the rate at which entangled ion-photon pairs can be produced and detected $4 \cdot p\bar{R}_a = 16.0 \times 10^3 \text{ s}^{-1}$ (22.7 s^{-1}) for Alice (Bob).

The advantage of using the Bell state analyser to perform ion-photon tomography is that we will be using the same apparatus during the ion-ion entangling experiment, and so we do not disturb the optical fibre (changing its unitary effect on the photon) between experiments.

6.3 Error sources

Our experiment developed from a nascent idea in October 2016, to a single vacuum system on an optics table in October 2017, to performing our first remote entanglement operations between two traps less than a year later. Such a pace of evolution meant little time was dedicated to characterisation beyond basic checks, so our discussion of errors is limited to a very approximate sense.

Errors in any experiment are divided into two types, the fundamental limitations due to the physics used and imperfections in the experimental apparatus. In this case, the fundamental limitations come mainly from the single qubit operations on the strontium ion; however, other experiments using optical qubits have achieved single qubit errors of order 5×10^{-5} [BXN⁺17], far below the $\sim 2.3 \times 10^{-2}$ errors we observe here. The only other potential fundamental issue is polarisation mixing of the collected photon. Although we have not fully quantified the expected errors, we believe that the majority of our polarisation errors are technical rather than fundamental.

6.3.1 Ion dephasing

Once a photon has been detected, the ion is projected into some superposition of $|\downarrow\rangle$ and $|\uparrow\rangle$, the relative phase of which evolves at the qubit frequency ω . A delay of 60 μs occurs before any analysis is performed¹⁴, during which the ion qubit phase evolves freely. Since we use a Zeeman qubit, the qubit frequency is directly proportional to the applied field, and therefore highly sensitive to fluctuations which will cause dephasing.

We can probe the coherence time with Ramsey experiments; preliminary results suggest a coherence time of $\tau \approx 0.4\text{ ms}$ in Alice and $\tau \approx 0.2\text{ ms}$ [Drm19], well below even the typical 10–40 ms values for a Zeeman qubit [HOTU11, SNM⁺13]. This can be extended with a spin-echo sequence (see for example [RSK⁺16]) to $T_2^* \approx 2.5\text{ ms}$, but this is still rather poor. Line triggering the Ramsey experiment at 50 Hz significantly improves coherence time, suggesting that a large portion of the noise is coherent rather than broadband.

The fidelity of a generic pure state $|\psi\rangle$ subject to dephasing (due to coherent noise) after time t can be written as in terms of a Gaussian decay [Mon11]

$$F = \langle\psi| e^{-\frac{1}{2}\frac{t^2}{\tau^2}} \rho |\psi\rangle + \langle\psi| \left(1 - e^{-\frac{1}{2}\frac{t^2}{\tau^2}}\right) \sigma_z \rho \sigma_z |\psi\rangle, \quad (6.7)$$

where $\rho = |\psi\rangle\langle\psi|$ is the density matrix associated with $|\psi\rangle$. If we substitute a generic state $|\psi\rangle = \cos\theta/2 |\downarrow\rangle + e^{i\phi} \sin\theta/2 |\uparrow\rangle$ and integrate over the Bloch

¹⁴Due to the entangler core taking some time to exit

sphere, the average fidelity to the original state is

$$F = \frac{2}{3} + \frac{1}{3}e^{-\frac{1}{2}\frac{t^2}{\tau^2}}. \quad (6.8)$$

The ion qubit measurements take place 60 μ s after the photon detection, due to technical limitations in how the entangler core exits. With a coherence time of 0.2–0.4 ms (we stress that these results were very preliminary), this gives an error of 0.4×10^{-2} – 1.5×10^{-2} .

Rough estimates place the magnetic field noise at 60 nT [Drm19] – the specified 20 μ A noise on the current supplies and calculated coil sensitivity (see chapter 3) produce an expected field noise of ~ 46 nT. The addition of a feedforward/feedback stabilisation circuit as demonstrated previously in the group [MTT⁺19] would significantly improve coherence times.

With optical qubits, the local oscillator for tracking the qubit phase is the laser itself, so the additional dephasing time from the oscillator is inversely proportional to the laser linewidth [ANK⁺15]. The linewidth of the 674 nm laser is specified as < 10 Hz, so this is not expected to limit the coherence time at present.

6.3.2 Ion rotation errors

Another significant contribution to our reported error is with the optical qubit rotations used to analyse the qubit. The optical qubit π -pulse fidelity is only around 0.997, well below what is possible, in part because the ion is not in the motional ground state during the analysis phase of the experiment – no sub-Doppler cooling is used, so even without heating effects this would

not be the case. Work is being done to implement sub-Doppler cooling based on electromagnetically induced transparency (EIT) [RLM⁺00], which would reduce the initial temperature before entanglement generation attempts. The ion is subject to heating while attempting entanglement generation, both by the scattering of resonant photons, as well as anomalous heating – the magnitude of this heating was not thoroughly explored. This effect is partially mitigated by improvements in collection efficiency since fewer entanglement attempts are required (meaning fewer photons scattered, and less time before entanglement), and so the ion is heated less.

Ion rotation errors account for $\approx 6 \times 10^{-3}$ of the total error, at $\approx 3 \times 10^{-3}$ per rotation. These errors only appear during the analysis of ion-photon entanglement, rather than during its creation, so in principle we could renormalise for the readout error and claim higher entanglement fidelity. However, for the entanglement to be useful (either in entanglement distillation protocols with ion-ion entanglement, or blind quantum computing with ion-photon entanglement) we wish to be able to utilise the state of the $^{88}\text{Sr}^+$ ion after its creation, so renormalisation is perhaps misleading.

To mitigate the heating of the $^{88}\text{Sr}^+$ ion causing gate errors, a second species could be used to sympathetically cool it during the entanglement generation process.

6.3.3 Laser leakages

Any further 422 nm photons scattered between the detection of the entangled photon and subsequent analysis leads to decoherence of the ion state.

The first iteration of the entanglement experiment extinguished the 422 nm Doppler cooling beam using its single pass AOM (see the diagram in figure 3.7), but left the double-pass AOM on to allow use of the state preparation beam. Leakage through the Doppler AOM was found to cause significant photon scattering, so another TTL controlled RF switch was installed and connected to the entangler core to immediately extinguish the double-pass AOM on detection of an entangled photon. This reduced the decoherence due to laser leakage to negligible levels.

6.3.4 Photon polarisation errors

Photon polarisation errors come in two forms. The first is that the photon polarisation we collected was mixed and there was no way of distinguishing the photon without error, which we hope to avoid with our collection scheme. However, if the collection axis is not exactly perpendicular to the quantisation axis, this will cause some error. Aberrations in the ion's point spread function coupled into the fibre could cause polarisation mixing.

Measuring the classical correlations of the photon with the ion can help to verify the effectiveness of the cancellation of polarization mixing by fibre coupling. Detectors 0 and 3 correspond to $|H\rangle$, and 1 and 2 to $|V\rangle$. We measure $P(|\downarrow\rangle | \Pi_{0,3}) \approx 0.995$ and $P(|\uparrow\rangle | \Pi_{1,2}) \approx 0.995$; the readout π -pulse infidelity is $\sim 3 \times 10^{-3}$ so we can bound polarisation mixing effects to $\sim 2 \times 10^{-3}$. Note that this measures the fidelity with which we can remotely prepare the ion in the computational basis.

The second type of error is the misidentification of a given photon po-

larisation on measurement, either because the polarising beam splitter does not have sufficient purity, or we have not applied the expected unitary to the photon with the waveplates. The unitary does drift because of changes in the fibre birefringence, however, the classical correlations are stable to $< 1e - 3$ over tens of minutes to hours, so this is not a significant contribution, and is calibrated between experiment runs.

6.3.5 Summary of errors

The total ion-photon infidelity measured was 2.1×10^{-2} and 2.3×10^{-2} in Alice and Bob respectively. This consists of: polarization error $\approx 2 \times 10^{-3}$, readout error $\approx 6 \times 10^{-3}$, and ion dephasing during the delay between photon detection and ion readout, for which a coherence time of ≈ 0.2 ms would give the remaining 1.3×10^{-2} – 1.5×10^{-2} error. Note that errors in preparation of the initial ground state of the ion before excitation do not contribute appreciably, since they do not result in a photon being emitted by the ion.

The ion-photon infidelity is dominated by effects on the ion rather than the ion-photon interface; the leading cause is fluctuations in the magnetic field. Active stabilisation of the magnetic field should reduce these errors, making ion-photon entanglement fidelities above 99% an achievable goal.

Chapter 7

Ion-ion entanglement

The sharing of entangled qubit pairs over macroscopic distances is crucial for the realisation of distributed quantum information technologies, such as quantum cryptography and quantum networking. The ability to create a link between stationary qubits (also known as quantum memories in this application) also leads to the concept of a quantum repeater, where many such links are leveraged to create entanglement over distances much larger than would otherwise be possible [SSdRG11], making the idea of a quantum internet possible. Such a network could be used for quantum communications, distributed quantum computing or leveraging the properties of entanglement for enhanced quantum metrology [Kim08].

However, it is important to note that we don't necessarily need kilometre-scale quantum networks to say that they are useful; high quality metre-, centi- or even millimetre-scale connections can help us to overcome one of the main problems with trapped ion quantum computing. The number of ions in a single trap we can use is limited to of order 10; spectral crowding means that

off resonant excitations of the wrong motional mode become problematic the more ions are included in the trap. Physically transporting qubits between locations of a multi-zone trap [KMW02, LWF⁺17] offers one solution, but a trap with many zones also poses difficulties; laser access to each zone becomes challenging in a large 2-dimensional array of surface traps. Using a photonic interface between trap zones allows us to break up the “quantum charge-coupled device” into many smaller, *simpler*, unit nodes, each of which is fairly simple [NFB14].

The key requirements of any quantum networking link are

- **Fidelity:** The remote entangled state should be as close to a maximally entangled state as possible to be useful for quantum applications. Ideally we would generate this in one step, but if the entanglement fidelity of the process is not optimal, we can use purification.
- **Purification:** We can theoretically condense multiple, low fidelity, remote entangled pairs into a single high fidelity state. In practice, for ion trap systems, this requires at least two ion species in the trap [NBBB16], one for remote applications, and the other to store states locally.
- **Rate and coherence time:** We would like to actually use the entanglement, and therefore the coherence time must exceed the time taken to establish a link. With two ion species, this becomes tractable; we use a species which is well suited to making photonic links for that purpose, and we can tolerate its poor coherence time by using a second species as a memory qubit.

Photonic interfaces are *heralded*, meaning that a successful entanglement

event is indicated by the detection of two photons in a Bell state. However, the heralding means that the generation of entanglement takes non-deterministic time; each entanglement attempt has a certain failure probability, and so if we perform a number of attempts, stopping immediately on success, there is a finite probability that all of these attempts fail and no entanglement is generated. In this case though, we know that no entanglement was generated, and we can abort whatever procedure we are undertaking. This stands in contrast to methods where an entangled state is generated with a certain probability on each attempt [PAVW07], but without a herald to indicate success. With the probabilistic methods, the entanglement process takes only one attempt, but the resultant state of the ensemble is a mixed state, part of which is entangled and part is not.

Quantum networking applications require that the time taken to establish entanglement be less than the coherence time of the entangled state, and thus again one of our main goals must be to maximise the rate of entanglement generation. The rate of two photon entanglement herald detections is given by

$$R_2 = \bar{R}_a \cdot p_{\text{Bell}} \cdot p_1^2, \quad (7.1)$$

where \bar{R}_a and p_1 are the average attempt rate and per-attempt probability of single photon detection introduced in the previous chapter. Note that the rate is conditional on the *square* of the single photon probability, so any improvements that can be made here are vital. We have already discussed how the repetition rate was maximised in chapter 4, where the reduction of latencies cut time from the experiment, and chapter 6 contains a detailed

description of the fast excitation sequence.

In this chapter, we will examine some of the theory for our case of entanglement swapping; rather than developing an all-encompassing theory, we present several considerations in turn. The author hopes that this makes physical interpretation of the theory easier – it also more closely matches the thought process used during feasibility studies at each stage.

7.1 Beamsplitter theory

We implement the photonic interface between ion traps by making use of *entanglement swapping*, where two elements of an ensemble can become entangled without directly interacting with each other, via the interactions of the rest of the ensemble. We generate two photons, each of which is (ideally) maximally entangled with the state of an ion (see chapter 6), and measure the photons in a two-particle basis which is maximally entangled, *i.e.* the Bell basis. The initial ensemble state of the two ions and two photons can be described as follows,

$$|\psi\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle|V\rangle + |\downarrow\rangle|H\rangle)_a \otimes \frac{1}{\sqrt{2}} (|\uparrow\rangle|V\rangle + |\downarrow\rangle|H\rangle)_b, \quad (7.2)$$

where the ion states are described in the spin basis $\{|\uparrow\rangle, |\downarrow\rangle\}$ and the photon states in the polarisation basis $\{|H\rangle, |V\rangle\}$, and subscripts denote Alice and Bob; at this point the photons are distinguishable because they are in

different spatial modes. Denoting the photon Bell basis by

$$|\Phi^\pm\rangle_P = \frac{1}{\sqrt{2}}\{ |H\rangle_a |H\rangle_b \pm |V\rangle_a |V\rangle_b \} \quad (7.3)$$

$$|\Psi^\pm\rangle_P = \frac{1}{\sqrt{2}}\{ |H\rangle_a |V\rangle_b \pm |V\rangle_a |H\rangle_b \}, \quad (7.4)$$

and with the basis defined similarly for the ions as

$$|\Phi^\pm\rangle_I = \frac{1}{\sqrt{2}}\{ |\downarrow\rangle_a |\downarrow\rangle_b \pm |\uparrow\rangle_a |\uparrow\rangle_b \} \quad (7.5)$$

$$|\Psi^\pm\rangle_I = \frac{1}{\sqrt{2}}\{ |\downarrow\rangle_a |\uparrow\rangle_b \pm |\uparrow\rangle_a |\downarrow\rangle_b \}, \quad (7.6)$$

we can rearrange equation (7.2) into

$$|\psi\rangle = \frac{1}{2}\{ |\Phi^+\rangle_I |\Phi^+\rangle_P + |\Phi^-\rangle_I |\Phi^-\rangle_P + |\Psi^+\rangle_I |\Psi^+\rangle_P + |\Psi^-\rangle_I |\Psi^-\rangle_P \}. \quad (7.7)$$

Thus if we can effect a perfect Bell measurement on the photons, we can project the ions into a corresponding maximally entangled Bell state. If we restrict ourselves to linear optics, the optimal number of Bell states that we can distinguish is 2 from 4 [VY99, LCS99], so the probability p_{Bell} in (7.1) is equal to $1/2$.

Bell state measurements on photons are generally effected with 50:50 beam splitters; the interaction of identical photons on a beam splitter was first investigated by Hong, Ou and Mandel [HOM87]. In their experiment, two identical photons (generated through parametric down-conversion of a single photon) incident on the input ports of a 50:50 beamsplitter were observed ‘bunching’ at the outputs; both photons exited together from either

of the output ports, and never exited through different ports.

In general our beam splitter will not be the idealised 50:50 splitter often treated in theory; what follows is a fairly rough and ready investigation of errors introduced by imperfect beam splitters. A much more thorough presentation of the error analysis is given in Ref. [Drm19], which incorporates the beam splitter imperfections, temporal profile of the photon, timing jitter and dark counts in the detector, and also dephasing and spin-flip errors on the ion.

We can express the action of a generic lossless¹ beam splitter with four ports $\{a, b, c, d\}$ in terms of a unitary matrix acting on the creation operators for each port;

$$\begin{pmatrix} a_c^\dagger \\ a_d^\dagger \end{pmatrix} = \begin{pmatrix} t & ir \\ ir & t \end{pmatrix} \begin{pmatrix} a_a^\dagger \\ a_b^\dagger \end{pmatrix}, \quad \{t, r\} \in \mathbb{R}, \quad t^2 + r^2 = 1, \quad (7.8)$$

where we have here explicitly shown the i phase shift on reflection so that t and r can be real numbers. The ports are labeled to match figure 7.1.

It is simple to show that with one photon in each input arm

$$\begin{aligned} |1_a 1_b\rangle &= a_a^\dagger a_b^\dagger |0\rangle \\ &= \left(t a_c^\dagger - ir a_d^\dagger \right) \left(-ir a_c^\dagger + t a_d^\dagger \right) |0\rangle \\ &= \left[(t^2 - r^2) a_c^\dagger a_d^\dagger - itr \left(a_c^\dagger a_c^\dagger + a_d^\dagger a_d^\dagger \right) \right] |0\rangle, \end{aligned}$$

and so by setting $t = r = 1/\sqrt{2}$, we recover the expected Hong-Ou-Mandel

¹For our purposes, all beamsplitters can be considered lossless – since we only herald when two photons are detected, by definition no loss event occurred.

result

$$\begin{aligned} |1_a 1_b\rangle &= \frac{-i}{2} (a_c^\dagger a_c^\dagger + a_d^\dagger a_d^\dagger) |0\rangle \\ &= \frac{-i}{\sqrt{2}} (|2_c 0_d\rangle + |0_c 2_d\rangle) \end{aligned}$$

where the two photons are “bunched” and always appear at the same output.

However, we are not working with identical photons in the same way as the original Hong-Ou-Mandel experiment – each photon is entangled with another particle, the ion it was emitted from. Thus we must consider the whole ensemble, and that the photons may have different polarisations. We introduce h_i^\dagger and v_i^\dagger as the creation operators for horizontal and vertical polarisations respectively. The beam splitter can be considered as two independent beam splitters, one acting on each polarisation, *i.e.* some unitary $\mathbb{B} = \mathbb{B}_h \otimes \mathbb{B}_v$. For the following calculation we consider only a beam splitter that is imbalanced but non-polarising², matching our measurement of the real beam splitter characteristics.

Rewriting equation (7.2) in terms of the output port creation operators,

²The calculation is no more difficult, just more tedious to represent, and because of the symmetry of the problem gives very little extra information to interpret.


and neglecting normalisation for simplicity, we now have

$$\begin{aligned}
|\psi\rangle &= (|\uparrow\rangle_a v_a^\dagger + |\downarrow\rangle_a h_a^\dagger) \otimes (|\uparrow\rangle_b v_b^\dagger + |\downarrow\rangle_b h_b^\dagger) |0\rangle \\
&= \left(|\uparrow\rangle_a (t v_c^\dagger - ir v_d^\dagger) + |\downarrow\rangle_a (t h_c^\dagger - ir h_d^\dagger) \right) \\
&\quad \otimes \left(|\uparrow\rangle_b (-ir v_c^\dagger + t v_d^\dagger) + |\downarrow\rangle_b (-ir h_c^\dagger + t h_d^\dagger) \right) |0\rangle \\
&= \left\{ |\downarrow\downarrow\rangle \left(-irt [h_c^\dagger h_c^\dagger + h_d^\dagger h_d^\dagger] + [t^2 - r^2] h_c^\dagger h_d^\dagger \right) \right. \\
&\quad + |\downarrow\uparrow\rangle \left(-irt v_c^\dagger h_c^\dagger - r^2 v_c^\dagger h_d^\dagger + t^2 h_c^\dagger v_d^\dagger - irt v_d^\dagger h_d^\dagger \right) \\
&\quad + |\uparrow\downarrow\rangle \left(-irt v_c^\dagger h_c^\dagger + t^2 v_c^\dagger h_d^\dagger - r^2 h_c^\dagger v_d^\dagger - irt v_d^\dagger h_d^\dagger \right) \\
&\quad \left. + |\uparrow\uparrow\rangle \left(-irt [v_c^\dagger v_c^\dagger + v_d^\dagger v_d^\dagger] + [t^2 - r^2] v_c^\dagger v_d^\dagger \right) \right\} |0\rangle . \quad (7.9)
\end{aligned}$$

We can rearrange equation (7.2) to group terms by the creation operators, giving

$$\begin{aligned}
|\Psi\rangle &= [h_c^\dagger h_c^\dagger + h_d^\dagger h_d^\dagger] |0\rangle (-irt |\downarrow\downarrow\rangle) + [v_c^\dagger v_c^\dagger + v_d^\dagger v_d^\dagger] |0\rangle (-irt |\uparrow\uparrow\rangle) \\
&\quad + h_c^\dagger h_d^\dagger |0\rangle (t^2 - r^2) |\downarrow\downarrow\rangle + v_c^\dagger v_d^\dagger |0\rangle (t^2 - r^2) |\uparrow\uparrow\rangle \\
&\quad + v_c^\dagger h_c^\dagger |0\rangle (-irt) \{|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle\} + v_d^\dagger h_d^\dagger |0\rangle (-irt) \{|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle\} \\
&\quad - v_c^\dagger h_d^\dagger |0\rangle \{r^2 |\downarrow\uparrow\rangle - t^2 |\uparrow\downarrow\rangle\} + v_d^\dagger h_c^\dagger |0\rangle \{t^2 |\downarrow\uparrow\rangle - r^2 |\uparrow\downarrow\rangle\} , \quad (7.10)
\end{aligned}$$

Terms in the first and second lines of (7.10) describe states where two photons of the same polarisation are detected – these are two photon events but *not* valid heralds producing entanglement. Our detectors are not photon-number resolving so in any case we cannot identify terms in the first line experimentally. States where photons of the same polarisation travel down



(a) The two detector analyser; this can only identify $|\Psi^-\rangle$ photon states, and requires that the photons are identical so that anti-bunching is suppressed.

(b) The four detector analyser; this can identify $|\Psi^+\rangle$ and $|\Psi^-\rangle$ photon states, as well as lifting the requirement on photon identity.

Figure 7.1: Two types of photonic Bell state analyser.

opposite arms of the beamsplitter (second line) are Hong-Ou-Mandel suppressed by $t^2 - r^2$, disappearing when $t = r = 1/\sqrt{2}$. These events can be used as a diagnostic; if the rate at which they occur exceeds our expectation from measurement of the beam splitter, then some other factor (for example, spatial mode mismatch) must be causing these photons to be non-identical.

The ion-ion states where photons of different polarisations travel down the same output arm of the beamsplitter (third line of (7.10)), are clearly equal (up to normalisation) to the maximally entangled Bell state $|\Psi^+\rangle$. The states produced when photons of different polarisation travel down different arms of the beam splitter (fourth line of (7.10)), $r^2 |\downarrow\uparrow\rangle - t^2 |\uparrow\downarrow\rangle$ and $t^2 |\downarrow\uparrow\rangle - r^2 |\uparrow\downarrow\rangle$, are clearly similar to the Bell state $|\Psi^-\rangle$. Their fidelity to $|\Psi^-\rangle$ is $\mathcal{F} = 1/2 + r^2 t^2 / (r^4 + t^4)$, which is 1 if the beamsplitter is balanced; in our experiment $r^2 = 0.47$ and so the expected error $\epsilon \equiv 1 - \mathcal{F} \sim 3.6 \times 10^{-3}$. We can consider

the two differing situations logically in terms of information; an entangled state is one where we have perfect knowledge of the joint state of a system, but no knowledge of the state of an individual particle. The beamsplitter's intended action is to erase the which-way information about which photon originated from which ion, and thus the knowledge of which state each ion is in. In cases where the both photons travel down the same arm of the beamsplitter, no amount of imbalance can give us any information about the origin of those photons, so the fidelity of the entangled state remains unity³, although the rate of these events will drop as the prefactor rt is maximal for $r = t$. In cases where the photons travel down different arms and the beamsplitter is imbalanced, clearly we have gained information about the origin of the photons; in the limiting cases of a perfectly reflective mirror, or no beamsplitter at all (perfect transmission), then coincidences still occur but create perfectly separable states rather than entangled states, since we know exactly which photon came from which ion, and hence the state that each ion is in.

Assuming our beamsplitter is perfectly balanced, then the ion-ion state produced on detection is

$$|\psi\rangle = \frac{1}{\sqrt{2}} \{ |\downarrow\uparrow\rangle + e^{i\phi_P} |\uparrow\downarrow\rangle \} , \quad (7.11)$$

where ϕ_P is 0 when clicks are registered on the same side of the Bell state analyser, and π if registered on opposite sides. The phase relationship holds for imbalanced beamsplitters.

³Until dark counts of the detector become a significant cause of coincidence detections

We should also consider the case where our polarising beam splitters are not perfect, *i.e.* we do not measure perfectly in the $\{h, v\}$ basis. If the two photons have the same polarisation, then propagation down different arms of the beamsplitter is Hong-Ou-Mandel suppressed, so we ignore polarisation errors in this case as a second order effect. Polarisation errors in the nominally heralded states (where photons have orthogonal polarisation) will cause losses, either as an unheralded two-photon event, or two photons incident on the same detector. A polarisation error with two photons of the same polarisation travelling down the same arm results in a falsely heralded two-photon event. Note that, with perfect measurement, these were the states that could tolerate any amount of imbalance without fidelity loss, so the fidelity of the ion-ion states produced by different heralds is (to first order) affected by either the balance of the nominal 50:50 NPBS or the purity of the measurement PBS, but not both.

We assume that the PBS error is small, so that we only consider errors that are first order in the imperfections, and also that the errors are symmetric⁴. Consider arm c of the Bell state analyser shown in figure 7.1b. If we write the action of the imperfect PBS as a map from the $\{h_c, v_c\}$ creation operators to Π_i^\dagger , $i \in \{0, 1\}$ the creation operators (now polarisation insensitive) at the detectors, then to first order in the imperfection $\epsilon \ll 1$, the PBS's actions are

$$\mathbb{B}_h = \begin{pmatrix} 1 & i\epsilon \\ i\epsilon & 1 \end{pmatrix}, \quad \mathbb{B}_v = \begin{pmatrix} \epsilon & i \\ i & \epsilon \end{pmatrix}.$$

⁴This is not the case in practice, but will suffice for an estimate

Applying this to the cases where the photon polarisations are the same,

$$\begin{aligned}
v_c^\dagger v_c^\dagger &= \left(\epsilon \Pi_0^\dagger + i \Pi_1^\dagger \right)^2 \\
&= -\Pi_1^\dagger \Pi_1^\dagger + 2i\epsilon \Pi_0^\dagger \Pi_1^\dagger + \mathcal{O}(\epsilon^2) \\
v_c^\dagger v_c^\dagger &= \left(\Pi_0^\dagger + i\epsilon \Pi_1^\dagger \right)^2 \\
&= \Pi_0^\dagger \Pi_0^\dagger + 2i\epsilon \Pi_0^\dagger \Pi_1^\dagger + \mathcal{O}(\epsilon^2).
\end{aligned}$$

The terms in $\Pi_0^\dagger \Pi_1^\dagger$ will contribute an error, as they are falsely marked as a herald. A detection of $|0011\rangle$, *i.e.* detectors 0 and 1 clicking, gives an (unnormalised) ion-ion state

$$\begin{aligned}
|\psi\rangle &= |\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle + 2i\epsilon(|\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle) \\
&= |\Psi^+\rangle + 2i\epsilon |\Phi^+\rangle,
\end{aligned}$$

by inspection of the relevant terms in (7.10). The fidelity of this state to $|\Psi^+\rangle$ is then $\mathcal{F} \approx 1 - 2\epsilon^2$. The PBS error ϵ is related to its polarisation extinction ratio (PER) by $\text{PER} = 1/\epsilon^2$, so assuming a PER of 1000 (see table 4.2 for the measured values), the expected infidelity in the states affected is 2×10^{-3} .

7.2 Entangled state phase

We have thus far ignored the phase evolution of the ensemble state throughout the detection process, effectively making the assumption that all the states are degenerate in energy. In general, this is of course not the case. In this calculation, we allow any of the levels involved to have different energies,




Figure 7.2: Diagram of the energy levels of an arbitrary system. The qubit states are split by an energy ω , while the average energy of the qubit states is given by $\tilde{\omega}$. Two *different* such systems labelled A and B can be used to enact an entanglement swapping protocol with photons. We calculate the phase of the entangled state created by the protocol in section 7.2 to illustrate and interpret the phases in the case of two nominally identical systems.

as illustrated in figure 7.2 in order to determine all possible phases that the ion-ion state can accrue – we then make the argument that all of these phases are either stable or negligible in this experiment.

After excitation of the two subsystems, but before either decay, the state is separable

$$|\psi\rangle = e^{i\omega_e^A t} |e\rangle \otimes e^{i\omega_e^B t} |e\rangle$$

and so all phase evolution is global, and can be factored out.

If we allow the subsystems to decay but do not detect the photons, the total energy in each subsystem is still the same, because the photon incorporates the energy difference between the excited state and either qubit level. However, there will be an additional phase dependent on the propagation of the photon,

$$\begin{aligned} |\psi\rangle &= e^{i\omega_e^A t} \left(e^{-ik_H^A Z^A} |\downarrow\rangle |H\rangle + e^{-ik_V^A Z^A} |\uparrow\rangle |V\rangle \right) \otimes e^{i\omega_e^B t} \left(e^{-ik_H^B Z^B} |\downarrow\rangle |H\rangle + e^{-ik_V^B Z^B} |\uparrow\rangle |V\rangle \right) \\ &= e^{i\omega_e^A t} \left(|\downarrow\rangle |H\rangle + e^{-ik^A Z^A} |\uparrow\rangle |V\rangle \right) \otimes e^{i\omega_e^B t} \left(|\downarrow\rangle |H\rangle + e^{-ik^B Z^B} |\uparrow\rangle |V\rangle \right), \end{aligned}$$

where we have defined $k^{A,B} = k_{\uparrow}^{A,B} - k_{\downarrow}^{A,B}$ the difference in wavenumber between the decays associated with each qubit state and factored out global phases in a way that will be convenient. Note also that this assumes that the paths for the two photon polarisations are identical *i.e.* $Z_H^{A,B} = Z_V^{A,B} = Z^{A,B}$; this would not be the case if using polarisation maintaining fibres as discussed in section 4.2.3.

Now imagine we detect an H -polarised photon at time $t = -\Delta t$, projecting one of the subsystems into $|\downarrow\rangle$, and starting its energetic evolution.

However, our beamsplitter (assumed perfect here) has erased our path information, so we don't know which ion was projected, and so the total state is

$$|\psi\rangle = e^{i\omega_{\downarrow}^A(t+\Delta t)} |\downarrow\rangle \otimes e^{i\omega_e^B(t+\Delta t)} \left(|\downarrow\rangle |H\rangle + e^{-ik^B Z^B} |\uparrow\rangle |V\rangle \right) \\ + e^{i\omega_e^A(t+\Delta t)} \left(|\downarrow\rangle |H\rangle + e^{-ik^A Z^A} |\uparrow\rangle |V\rangle \right) \otimes e^{i\omega_{\downarrow}^B(t+\Delta t)} |\downarrow\rangle .$$

The V -polarised photon is detected at $t = 0$; there may be an additional phase $\phi_P = \{0, \pi\}$ associated with the specifics of the photon detection, as discussed in section 7.1 – we omit this for clarity here. Now both subsystems will evolve with the energy of the qubit states, and there is a fixed phase introduced from the period between the decays

$$|\psi\rangle = e^{i\omega_{\downarrow}^A(t+\Delta t)} |\downarrow\rangle \otimes e^{i\omega_e^B \Delta t} e^{i\omega_{\uparrow}^B t} e^{-ik^B Z^B} |\uparrow\rangle \\ + e^{i\phi_P} e^{i\omega_e^A \Delta t} e^{-ik^A Z^A} e^{i\omega_{\uparrow}^A t} |\uparrow\rangle \otimes e^{i\omega_{\downarrow}^B(t+\Delta t)} |\downarrow\rangle .$$

We can rewrite this in the form $|\psi\rangle = |\uparrow\downarrow\rangle + e^{i\phi} |\downarrow\uparrow\rangle$, where ϕ is given by

$$\phi = \Delta\omega^{AB} t + \Delta\omega^{AB} \frac{\Delta t}{2} - \Delta\tilde{\omega}^{AB} \Delta t + \Delta\omega_e^{AB} \Delta t - \Delta\omega^{AB} \frac{Z}{c} + \frac{\omega^A + \omega^B}{2} \frac{\Delta Z}{c}, \quad (7.12)$$

where the following relations and new symbols have been used in order to

symmetrise this equation with respect to the upper and lower qubit levels

$$\begin{aligned}
\omega^{A,B} &= \omega_{\uparrow}^{A,B} - \omega_{\downarrow}^{A,B} \\
\omega_{\uparrow,\downarrow}^{A,B} &= \tilde{\omega}^{A,B} \pm \omega^{A,B}/2 \\
\Delta\tilde{\omega}^{AB} &= \tilde{\omega}^A - \tilde{\omega}^B \\
\Delta\omega^{AB} &= \omega^A - \omega^B \\
\Delta\omega_e^{AB} &= \omega_e^A - \omega_e^B \\
k^{A,B} &= \omega^{A,B}/c \\
Z^{A,B} &= Z \pm \Delta Z/2.
\end{aligned}$$

We emphasise to the reader that the reference time chosen for this phase is the detection time of the V -polarised photon, and *not* the detection time of the second (or first) photon to be detected. Note also that the excitation time does not appear in this equation – this is to be expected, since spontaneous emission should carry no memory of the time of excitation, only the time of emission.

Terms in (7.12) multiplied by Δt (the difference in time between photon detections) can affect the entanglement fidelity, because we select events in which the photons can be detected anywhere within a 30 ns window – over an ensemble of events the spread in Δt will lower the entanglement fidelity [VHI⁺14]. Because we are using Zeemann qubits in otherwise identical ions, the difference between the two subsystem average qubit energies $\Delta\tilde{\omega}^{AB}$ is identically zero and this term can be ignored. The constant difference in qubit frequencies $\Delta\omega^{AB} \sim 2\pi \times 14$ kHz in our experiment is due to the

uncertainty in our measurement of the magnetic field in each trap⁵. The maximum difference in photon detection times is $\Delta t = 30$ ns, so the product of these, $\Delta\omega^{AB}\Delta t < 3 \times 10^{-3}$, is negligible. The difference in excited state energies $\Delta\omega_e^{AB}$ is also set by a Zeemann splitting in the magnetic field, and is therefore the same order of magnitude as the $\Delta\omega^{AB}$ term⁶, so we also neglect this term.

In (7.12), we rewrote the path length phase term $k^B Z^B - k^A Z^A$ into common and differential path length terms. The common path length is multiplied by the difference in qubit splittings, and so extremely large common mode path fluctuations (\sim metre!) would be needed to introduce even a 1×10^{-3} phase jitter. Indeed, the total path itself is only ~ 3 m so this term is completely negligible. The differential path length term depends on the average qubit splitting, and so is more sensitive to the length; this would still would have to vary on the scale of \sim mm to cause 1×10^{-3} phase jitter, and so though while this term may not be negligible (the paths may differ by up to a few cm), its variations certainly are.

Finally we consider the free phase evolution term $\Delta\omega^{AB} t$; we have already discussed that the qubit frequency difference is small, so this term evolves slowly. Fluctuations in $\Delta\omega^{AB}$ are exactly the independent fluctuations of magnetic field in each trap that cause dephasing of a single ion's state; these are not negligible, and are considered separately in section 6.3.1. Our analysis occurs $60 \mu\text{s}$ after the creation of the entangled state; as long as the jitter of the analysis time is $\mathcal{O}(10\text{ns})$, the constant part of this phase is essentially

⁵This uncertainty is because we do not do long Ramsey experiments to determine the qubit frequency.

⁶actually smaller, because $g_J = 2/3$ for the $P_{1/2}$ level

fixed.

7.3 Indistinguishability criteria

Complete indistinguishability of the photons is necessary for a 2-detector scheme like Ref. [MMO⁺07] (see figure 7.1) to suppress anti-bunching, but once we can measure the polarisation of the photons as in the 4-detector scheme this restriction is lifted [VHI⁺14] as we saw in section 7.1. However, there are still some restrictions that apply in order to maximise the fidelity of the entangled state, and the rate at which it is produced.

We have already seen that the main action of the beamsplitter is to erase the which-way information of the photons. If the distribution of photon arrival times from each ion is different, then in principle with time resolving detectors we can say that one of the photons was more likely to have been emitted by one of the ions than the other – hence we have gained some information, causing a loss in fidelity. Therefore we must make sure that the *distributions* of photons arriving from each ion are identical.

We also know that the phase of the entangled state depends on the difference between the two ion qubit frequencies, as well as the difference in detection time of the photons. If we can eliminate the qubit frequency difference, the phase of the entangled state is fixed in time, and no phase jitter is introduced from the difference in detection times.

7.3.1 Temporal matching

In order that the photons be temporally indistinguishable, the temporal profiles of the photons from each trap must be identical up to their amplitude (*i.e.* the collection efficiency). The ion is decaying in free space unmodified by a cavity, so the photon's temporal profile is simply an exponential decay with time constant equal to the upper state lifetime 7.39 ns [PBL95]. After detection, this appears as an exponential decay convoluted with the jitter both of the detector itself and the circuitry registering the input event, so the distribution of measured photon arrival times has a smooth leading edge as seen in figure 7.3. The exponential decay is self-similar at all times, so if we choose our window such that the initial time bins (before the distribution is in the exponential decay regime) are excluded, then the distributions from each trap will be identical. However, this initial section contains a large part of the integrated probability of photon detection, so by excluding it we drastically reduce the rate of two-photon events. Therefore we must endeavour to align the photon arrival times as well as possible to be able to maximise both fidelity and rate.

The detector chain has order ns jitter, but we can achieve almost arbitrary timing accuracy by collecting many photons, then fitting the resulting histogram of arrival times relative to the pulsed excitation. The same detector and input electronics are used for to gather histograms of the arrival time for each trap. The histograms are fitted with the convolution of the detector response⁷, incorporating the latency and jitter of the detection chain, and the

⁷Modeled as a Gaussian response to enable an analytic form for the convolution

exponential decay of the ion [Drm19]. The fit is constrained such that the detector parameters and decay time constant (*i.e.* the physical properties of the ion and the detection chain) are identical while allowing the start time of the decay to vary between the two systems; an example fit for one of the traps is shown in figure 7.3. Any difference in start time can be compensated by changing the path length of the pulsed laser to the ion – note that although the phase of the entangled ion state depends on the path difference from the ions to the detector, the entanglement fidelity does not, provided the photons arrive at the detector simultaneously as discussed. After one such path adjustment, the arrival time difference between the two traps was reduced to 13 ps. This is equivalent to a path length difference of 4 mm but note that this is the combined optical path length from the pulsed laser to the ion and to the detector, and not simply the length from ion to detector.

We also wish to note that the use of a single pulse split between the two traps is a matter of convenience; the pulses serve only to produce the excited state of the ion, and there is no need for these to have any phase stability with each other. This is important for entanglement generation across a real network, as provided Alice and Bob have sufficiently synchronised clocks, there is no need to propagate a short pulse through long optical fibres.

7.3.2 Field equalisation

Ensuring that the ion qubits are identical is *not* a necessary condition for the creation of entanglement [VHI⁺14]. However, if the qubit frequencies match exactly, the phase of the entangled state is fixed and does not depend




Figure 7.3: Maximum likelihood fit and residuals for the distribution of 255 449 photon arrival times, normalised as a probability distribution function. The fit indicates a $P_{1/2}$ lifetime of 7.47(2) ns, agreeing with the literature [PBL95]. The dotted line in the main figure is at the maximum of the fitted distribution, while further lines at integer multiples of 12.5 ns (the pulsed laser period) in the residuals guide the eye in looking for multiple excitations. Residuals are normalised by the square root of the histogram value. The model does not include a mechanism for dark counts.

on the time elapsed between or after photon detection, resulting in a higher fidelity when averaged over all events. If the photons are identical, Hong-Ou-Mandel suppression helps to minimise errors introduced by imperfect polarising beam splitters (see section 7.1), making the 2-photon events that are not valid heralds more useful as channels for identifying errors.

The absolute frequency of the photon and the ion qubit frequency is set by the Zeemann splittings due to the magnetic field (see figure 2.1), so we must make sure the applied magnetic field is the same for each ion. With an absolute frequency reference it would take only one measurement to exactly determine the field, however, the exact frequency of the 674 nm laser is also unknown, so we must make two frequency measurements. The difference in AOM offset frequency between the “main” and “map” transitions allows us to calculate the magnetic field exactly; once the field is known then calculation of the unshifted frequency of the 674 nm laser follows⁸.

The measured field in both trap systems is $0.5612(5) \text{ mT}$ ⁹ – with this level of uncertainty and the ion qubit splitting sensitivity of 28 MHz mT^{-1} , the difference in qubit frequencies is bounded to $\sim 2\pi \times 14 \text{ kHz}$ ¹⁰. The maximum arrival time difference of the photons is 30 ns by our choice of detection window – although because of the exponential decay profile the vast majority of events will occur with a much smaller difference – and so we expect this to contribute an error in phase of at most 3×10^{-3} (see section 7.2. If we wish to

⁸In practice, our calculations use the change in offset frequencies; the same principles apply

⁹That is to say, they are identical and stable over days up to the measurement uncertainty. The uncertainty is very approximate, from observing the changes when re-running the field calculation experiment.

¹⁰This is roughly as expected, given that the maximum interrogation time of the transition frequency is some 200 μs

also consider the identity of the photons, then we can compare with their natural linewidth 21 MHz – clearly frequency differences of 14 kHz are not resolved, and therefore the photons can be considered identical in frequency. By counting the occurrence of the two-photon events that should be Hong-Ou-Mandel suppressed¹¹, we can measure the photon’s non-identity in spatial mode.

7.4 Fibre birefringence compensation

The single mode fibre used to collect photons from the ion exhibits some residual birefringence, due to small anisotropies in the core or cladding and mechanical stresses. Thus, the polarisation basis of the photon at the Bell state analyser after propagation through the fibre is unknown. The ion-photon tomography results from chapter 6 can be used to predict the waveplate angles needed to maximally correlate the photon basis with the ion’s computational basis for each trap.

The waveplate angles are then optimised to ensure the correlations are maximal, giving correlations $P(|\downarrow\rangle | \Pi_{0,3}) \sim 0.995$ and $P(|\uparrow\rangle | \Pi_{1,2}) \sim 0.990$ in both traps. Thus, measuring the photon polarisation performs a state preparation operation on the ion, exactly as required for the remote entanglement scheme.

¹¹These can and do occur because our NPBS is imbalanced; we want to count the excess above our predicted frequency.

7.5 Remote entanglement

The excitation scheme used for remote entanglement is the same as that described in chapter 6 – the pulse sequence to produce an entangled photon from a single trap is shown in figure 6.2. The control systems of the two nodes are synchronised before the entanglement sequence commences. For operation during entanglement attempts, the master node (Alice) controls both the pulse picker, setting the timing of the 422 nm ps pulse, and monitors the inputs from the APDs. On detection of a valid herald, the master node communicates a stop flag to the slave so that the entangled state can be analysed.

The probability of successfully detecting a herald is given by

$$P = p_{\text{Bell}} \cdot p_1^2, \quad (7.13)$$

where $p_{\text{Bell}} = 1/2$ is the probability of projecting the photons into a Bell state (discussed in section 7.1), and p_1 the probability of producing and detecting a photon with the pulsed excitation.

After successfully detecting a heralding pair of photons, we can verify the entanglement using standard two-qubit tomography protocols. We measure the ions in an over-complete basis so that the density matrix for each click pattern can be reconstructed using the maximum-likelihood estimation (MLE) technique – see [Nad17] for further explanation for the ion-photon case. The pure state created on herald detection contains a stable inter-node phase introduced in section 7.2, such that the pure state $|\psi\rangle = |\uparrow\downarrow\rangle + e^{i\phi} |\downarrow\uparrow\rangle$. This phase can in principle be adjusted with local operations and classical

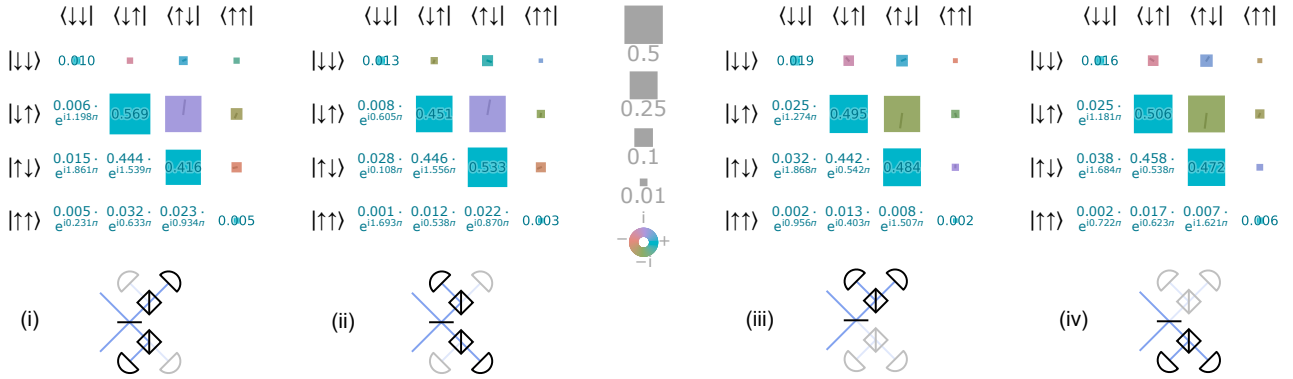


Figure 7.4: The density matrices reconstructed for each herald pattern. The herald patterns with the detector clicks on opposing sides of the 50:50 beam splitter ($\phi_P = \pi$) are shown to the left of the key, and patterns with clicks on the same side ($\phi_P = 0$) are to the right. The average fidelity to the closest maximally entangled state ($|\Psi^+\rangle$ on the left, $|\Psi^i\rangle$ on the right) is 94.0(5)%. The total probability of any herald occurring per attempt is 2.2×10^{-4} . The *area* of the square shows the magnitude of the matrix element, with the colour representing its complex phase.

communication (LOCC) between Alice and Bob, so in the analysis we allow the use of such transformations to individually maximise the fidelity of the density matrix for each click pattern to the closest Bell state (either $|\Psi^+\rangle$ or $|\Psi^-\rangle$). The resulting density matrices after local operations are shown in figure 7.4¹². Note that this implies a different set of local operations for each detector pattern; this is not problematic at all, since the detector pattern is known immediately, any local operations to be used can be chosen with decision branching in the experiment. The average fidelity of states resulting from the four click patterns to the closest Bell state is 94.0(5)%¹³.

The total per-attempt probability of a heralded event is 2.2×10^{-4} . The

¹²Data analysis by David Nadlinger

¹³Strictly we should weight by the probability of each click pattern, but these (and the fidelities) are all close enough that it makes little difference

tight loop of entanglement attempts runs at 1 MHz for 500 μs , before breaking for 100 μs of Doppler cooling if unsuccessful. Given that many more than 500 attempts are needed on average, we can consider the average attempt rate to be 833 kHz, resulting in an entanglement generation rate of 182 s^{-1} . Considering this as performing a number of Bernoulli trials with only a small probability of success, if we attempt for 12.6 ms we have a 90% chance of having produced entanglement – another 12.6 ms of attempts increases this to 99%.

7.6 Error sources

The error sources in the ion-ion entanglement are largely shared with the ion-photon experiment in chapter 6 – the largest errors are due to gate errors and dephasing of the ion qubit. New error channels introduced are

- the accuracy with which we match the photon basis to the ion computational basis;
- imperfect beam splitters, both 50:50 and polarising;
- spatial overlap of the photon modes at the beam splitter;
- arrival time mismatch of the photon distributions;
- phase jitter of the entangled state; and
- dark counts contributing false heralds.

The matching of the photon basis to ion basis consists of the polarisation mixing discussed in chapter 6, drifts in the single mode collection fibre bire-

fringe, errors in setting the waveplates for each trap, and polarising beam splitter errors. These errors combine to give the error in observing correlations between measurements of the photon and ion state. Our correlations are typically 0.995, including the readout error of 3×10^{-3} . We thus bound polarisation mixing errors to around 2×10^{-3} .

We discussed the errors due to beam splitters extensively in section 7.1, concluding that the errors are on average $< 2 \times 10^{-3}$. The infidelity caused by an imbalanced beam splitter is quadratic with the imbalance and affects only the 2/4 heralds where the photons are detected on opposite sides of the beam splitter, so small imbalances are not greatly important. The PBSs contribute an error to the heralds where the photon is detected on the same side of the beamsplitter; this is small, and is considered as part of the polarisation mixing effects.

The error due to temporal mismatch of the photons can be calculated analytically [Drm19]. The maximally-entangled fidelity is given by

$$\mathcal{F} = \frac{1}{2} [1 + \langle \vartheta_A, \vartheta_B \rangle^2] ,$$

where $\langle \vartheta_A, \vartheta_B \rangle$ is the overlap integral between the temporal wavefunctions of photons from Alice and Bob respectively. If we ignore the detector response

and consider only the exponential decay¹⁴, the overlap is

$$\begin{aligned} \int_{-\infty}^{\infty} \vartheta_A^*(t) \vartheta_B(t) dt &= \frac{1}{\tau} \int_0^{\infty} e^{-\frac{(t+\Delta t)}{2\tau}} e^{-\frac{t}{2\tau}} dt \\ &= e^{-\frac{\Delta t}{2\tau}} \\ &\approx 1 - \frac{\Delta t}{2\tau} \end{aligned}$$

where Δt is the temporal misalignment, and we have preempted that it is much smaller than the excited state lifetime τ . With $\Delta t = 13$ ps and $\tau = 7.39$ ns [PBL95], our estimate of the infidelity is 9×10^{-4} .

Any phase jitter in the entangled state causes a loss in fidelity over the ensemble of created states. Recalling that we can remove any fixed phase offsets with local operations, the infidelity of $|\psi\rangle = |\downarrow\uparrow\rangle \pm e^{i\phi} |\uparrow\downarrow\rangle$ to the state $|\Psi^\pm\rangle$ for small ϕ is $\phi^2/4$; clearly with phase jitter $\Delta\phi \lesssim 1 \times 10^{-3}$ the infidelity from this term is completely negligible. The phase uncertainty caused by fluctuations in the magnetic fields for each trap is considered separately as the dephasing of each individual ion.

Dark counts will contribute an infidelity to the final state – if we detect a heralded pattern but one of the clicks is simply detector noise, then the state of one of the ions is uncorrelated with this click. We can consider this process by looking at only one herald pattern. We neglect events consisting of two dark counts because the dark count probability in a given detector during the 30 ns detection window $p_{\text{dark}} = R_{\text{dark}}\tau = 1.8 \times 10^{-6}$, assuming a dark rate of 60 Hz per detector, is many orders of magnitude smaller than

¹⁴This assumption is the reason for the discrepancy between the figure quoted here and Ref. [Drm19]

detecting a ‘real’ photon emitted by an ion. The probability of detecting a herald pattern i with photons from both ions is $1/4$ that of the total success probability $P_i = 1/8p_1^2$. For each herald pattern, there are four ways in which a dark count can produce a false herald – *i.e.* a photon from Alice is detected in detector 0, with a dark count in detector 1, and so on permuting the location of the photon from an ion and which ion it originated from, so the probability of any of these situations occurring is

$$P(\text{dark herald}) = 4 p_{\text{dark}} \cdot \frac{1}{4} p_1 ,$$

where the probability of detecting a photon from an ion in a *given* detector is $(1/4)p_1$. The relative probability of any dark heralded event occurring is then $8 p_{\text{dark}}/p_1$ compared to a ‘real’ event.

We can examine the density matrix of each of the dark herald situations to calculate their fidelity to the maximally entangled state. For a dark click in the vertical polarization detector, and the horizontal photon originating from Alice, the density matrix is (neglecting decays to the $D_{3/2}$ level)

$$\rho(\text{dark } V, \text{ Alice } H) = \frac{1}{3} |\downarrow\uparrow\rangle \langle\downarrow\uparrow| + \frac{2}{3} |\downarrow\downarrow\rangle \langle\downarrow\downarrow| ,$$

whereas with the detections reversed we have

$$\rho(\text{Alice } V, \text{ dark } H) = \frac{1}{3} |\uparrow\uparrow\rangle \langle\uparrow\uparrow| + \frac{2}{3} |\uparrow\downarrow\rangle \langle\uparrow\downarrow| ,$$

and the expressions for detections from Bob follow logically. We have considered each of these separately because of the different branching ratios to

the qubit states for an isotropic decay, *i.e.* one where the photon is not gathered into the fibre and detected. It is easy to see that the average fidelity ($F = \langle \Psi^\pm | \rho | \Psi^\pm \rangle$) of the four possible states is 0.25^{15} . If we assume the fidelity of the state is 1 when two photons from ions are detected, then for small dark count probabilities the overall fidelity is

$$\begin{aligned} F &= [1 - P(\text{dark herald})]F(\text{real herald}) + P(\text{dark herald})F(\text{dark herald}) \\ &= 1 - 8 \frac{p_{\text{dark}}}{p_1} (1 - F(\text{dark})), \end{aligned}$$

and thus the infidelity introduced is $6p_{\text{dark}}/p_1 \sim 5 \times 10^{-4}$.

The spatial overlap of the photon modes at the beam splitter affects fidelity; this is best considered in terms of which-way information, as with misaligned photon modes we can vary the relative likelihood of a photon arriving at a given detector, and any which-way information reduces the fidelity of the entangled state. We mitigate this by fibre coupling the output modes of the beamsplitter before detection (see section 4.2) – this costs some 20 % in entanglement rate as the coupling efficiency into the detector fibre is $\sim 90\%$. To obtain an estimate for the spatial mismatch error, we can count the occurrences of Hong-Ou-Mandel suppressed two-photon events (which do not herald entangled ion-ion states). Preliminary results¹⁶ suggest an error of order 1×10^{-2} – 2×10^{-2} , so we attribute the remaining 1.3×10^{-2} infidelity to spatial mismatch of the photon modes.

¹⁵Strictly, we should include some of the more unlikely possibilities: a photon being produced, collected into the fibre and subsequently *not* detected; and the state being prepared into $|\downarrow\rangle$ but not excited. The average fidelity for all these cases is 0.25 anyway.

¹⁶Ideally these statistics would be gathered at the same time as the ‘real’ heralds, but this was not implemented and so this estimate is not contemporaneous with the ion-ion

	Infidelity
Ion gate errors	$6 \times 10^{-3} + 6 \times 10^{-3} = 1.2 \times 10^{-2}$
Ion dephasing	$1.3 \times 10^{-2} + 1.5 \times 10^{-2} = 2.8 \times 10^{-2}$
Polarisation mixing	$2 \times 10^{-3} + 2 \times 10^{-3} = 4 \times 10^{-3}$
50:50 NPBS imbalance	1.7×10^{-3}
Arrival time mismatch	9×10^{-4}
Entangled state phase jitter	$\ll 1 \times 10^{-4}$
Dark counts	5×10^{-4}
Spatial mode mismatch	1.3×10^{-2}
Total	6×10^{-2}

Table 7.1: A (very) approximate error budget. The first group of errors are also present in the ion-photon experiment, and are largely related to analysis of the entangled state, while the second group are the errors discussed in this section. Note that the spatial mode mismatch is inferred from the remaining infidelity.

The estimates for each contribution to the infidelity are summarised in table 7.1; errors pertaining to the analysis of the entangled ion-ion state (*i.e.* ion gate errors during analysis, and dephasing of the individual ions during the wait period before analysis) total 4×10^{-2} , accounting for a large portion of the measured infidelity. As discussed in section 6.3.5, active stabilisation of the magnetic field should drastically improve the measured fidelity. The next largest contributor to the infidelity is likely to be the mismatch of the photon's spatial mode at the beamsplitter and requires a more detailed investigation to confirm that this is the case.

entanglement data.

Chapter 8

Conclusions

We have demonstrated heralded entanglement between strontium ions in two traps separated by several metres. The average fidelity to the nearest maximally entangled state was 94.0(5) %, with entangled ion-ion states produced at a rate of 182 s^{-1} . This constitutes a high-rate, high-fidelity entanglement resource of the type required for entanglement distillation. The apparatus constructed within this thesis is equipped with

- multiple ion species ($^{88}\text{Sr}^+$ and $^{43}\text{Ca}^+$),
- a microfabricated segmented trap with shuttling capability, and
- photonic interconnects,

so a distillation protocol as proposed in Ref. [NBBB16] is certainly possible within this experiment.

As an intermediate result, we also showed entanglement between a single spontaneously emitted photon and an ion in each of the two traps used. The maximally-entangled fidelity and rate were 97.90(12) % and $3.98 \times 10^3 \text{ s}^{-1}$ in

the trap designated “Alice”, and 97.70(12) % and $5.68 \times 10^3 \text{ s}^{-1}$ in “Bob”. This result should allow demonstration of a blind quantum computing experiment where a client, making only measurements on the photon, can perform a computation on the server, the ion trap, without the server gaining any knowledge about the algorithm used [MF13].

The principle source of error in our experiments is decoherence of the ion qubit due to magnetic field fluctuations. No field stabilisation was used; this is expected to provide a significant improvement in fidelity and will allow us to probe other error sources once implemented. These results were achieved using only Doppler cooling – ground state cooling should improve the optical qubit manipulations. The error contribution from polarisation mixing of the collected photon was not a limiting factor, and could not be probed beyond an upper bound of 0.2 % infidelity because of the other errors present.

The current rate is not fast enough for multiple entangled pairs to be generated before the strontium Zeeman qubits decohere; however, we argue that this is not the most useful metric for a mixed species remote entanglement experiment. The coherence time should be compared to the SWAP gate duration with the planned calcium memory, and the remote entanglement rate compared to the calcium coherence time. The coherence time of the strontium qubit is currently below the tens of ms [HOTU11] expected for a Zeeman qubit, but still exceeds typical gate times of 50–100 μs .

8.1 Comparison to other experiments with ions

Identifying the ‘best’ results from the literature is inherently difficult for this type of experiment, as there are many potential figures of merit. Clearly the fidelity is important, but what about rate? The detected rate is manifestly dependent on both the quantum efficiencies of the detectors used and the probability of a photon reaching the detector. Should we prioritise protocols that have high photon generation rate but unfavourable detection rates with currently existing detectors and optical fibre materials, or the converse with excellent transmission and detection but limited rate? These are open questions; the author’s opinion is that each scheme can offer advantages at different length scales of entanglement distribution in the medium term, while advances in detectors, optical fibres, and frequency conversion of the photons may lead to one candidate emerging as the favourite.

8.1.1 Ion-photon entanglement

The closest comparable experiment, using ultra-fast excitation in $^{171}\text{Yb}^+$ and collecting photons in free space, reports an ion-photon fidelity of 92% at a rate of $2.1 \times 10^3 \text{ s}^{-1}$ [HIV⁺15]. When compared with weak excitation to the upper level, ultra-fast excitation offers faster repetition rates (because the interaction time is much shorter) as well as a higher probability of producing an entangled photon, limited only by the branching ratio into the lower state. The wavelength used for ytterbium experiments is 369 nm, which

¹Inferred from the quoted repetition rate and two-photon detection probability

does not have favourable transmission properties in optical fibres². Barium’s wavelength is certainly better in this regard at 493 nm, allowing researchers at the University of Washington to observe entanglement with a fidelity of 84(1) % at 2.5 s^{-1} [ACNB14]. Work is also done with the UV transitions of calcium – entangled ion-photon pairs produced with 84.6(2) % fidelity at 140.5 s^{-1} [KES⁺16].

Other free-space experiments have used different schemes; using the decays into low-lying D -states is possible, although the probability of creating a photon is reduced because of the lower branching fraction. These photons are then in the near-infrared, and can be frequency converted to allow for even better transmission through optical fibres. 854 nm photons have been converted to 1310 nm, resulting in an ion-photon entanglement fidelity of 94.8(2) % at 24.8 s^{-1} [BEK⁺18].

In principle, optical cavities can offer higher collection efficiencies. One recent notable work collected 854 nm photons from $^{40}\text{Ca}^+$ into an optical cavity with 16 % efficiency³. The entanglement fidelity at this stage is comparable to our experiment, at 96.7(6) %. The entangled photons were subsequently downconverted to telecoms wavelength (1550 nm) and transmitted through 50 km of optical fibre [KMS⁺19], with a maximally-entangled fidelity of 90(3) % and at a rate of $\sim 1\text{ s}^{-1}$ – while both rate and fidelity are lower than the work in this thesis⁴, this is clearly an important result in ion trap networking given the distance involved.

²Another technical issue is that high numerical aperture objectives are harder to align at shorter wavelengths.

³The *collection* efficiency in this thesis is $\sim 3.5\%$.

⁴Of course with 50 km of optical fibre, they have to wait some time before being able to detect the photon!

8.1.2 Ion-ion entanglement

To the author’s knowledge, the only other group currently performing remote entanglement with ions is at the University of Maryland. Their latest published results are an ion-ion fidelity of 78(3)% at 4.5 s^{-1} entanglement rate [HIV⁺15] with a polarisation-encoded photon scheme, as well as previously using a frequency encoded scheme [Hay12]. This group have demonstrated several extensions to the remote entanglement protocol, including deliberately making the photons from each trap maximally distinguishable and still demonstrating entanglement [VHI⁺14]⁵. They have also performed a local two-qubit gate (between identical species) after having generated the entanglement [HIV⁺15], which (with mixed species) is the next experiment that we hope to accomplish.

Both the University of Maryland and the University of Washington are working towards remote entanglement systems using $^{138}\text{Ba}^+$, whose wavelength of 493nm is more favourable in this regard, but these systems have not yet been completed.

8.2 Comparison with alternative platforms

Trapped ions are not the only possible qubit platform, and other platforms can also perform photonic entanglement. This can be done with a heralded scheme similar to that presented in this work, or by directly mapping the quantum state of a matter qubit onto a photon and reabsorbing the photon

⁵We did the same experiment by accident during initial data acquisition – we were puzzled why the fidelity was so low until we realised our error!

at a different location.

With neutral atoms, the exchange of a single photon has been used to create a remotely entangled state at a rate of 30 s^{-1} [RNH⁺12]. However, the readout step to verify the entangled state (with an entanglement fidelity of 85 %) has a low success probability, reducing the rate to some $5 \times 10^{-2} \text{ s}^{-1}$. Exchange of a single photon was also used to create entanglement between a single atom and a Bose-Einstein condensate with a fidelity of 95 %, but the detected rate was low at $2.4 \times 10^{-3} \text{ s}^{-1}$ [LMR⁺11].

Heralded entanglement has been demonstrated in nitrogen vacancy centres, with high fidelity (92 % at a rate of $3 \times 10^{-4} \text{ s}^{-1}$ [HBD⁺15]) and high rate (39 s^{-1} , with fidelity 60 %), but the combination of high fidelity and rate has not yet been achieved.

Quantum dots have also been used to perform a heralded entanglement scheme, with fidelity 61.6 % and a rate of $7.3 \times 10^3 \text{ s}^{-1}$. While this rate is impressive, the time between entanglement events compares poorly with the typical quantum dot coherence times of around a microsecond.

8.3 Outlook

How close are we to building a useful quantum network of trapped ions?

All of the required building blocks for constructing such a network have been demonstrated in ion trap experiments. Local same-species two-qubit gates are possible with high-fidelity with both lasers [BHL⁺16, GTL⁺16] and microwaves [HSA⁺16]. Two-qubit gates between different ion elements are continually improving with fidelities of 97.9 [TGL⁺15] reported between

${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$, and preliminary results suggesting a gate fidelity of 99.7% between ${}^{88}\text{Sr}^+$ and ${}^{43}\text{Ca}^+$ here in the Oxford group (in another apparatus). Several mixed species gate operations have been combined with ion shuttling to perform quantum gate teleportation between separated qubits [WKE⁺19]. The work presented in this thesis represents a substantial improvement in both remote entanglement fidelity and rate – so what can we do with this?

With the current remote entanglement error, we would only need 8 raw entangled pairs to distil a remote entangled state to 99.7% maximally entangled fidelity, assuming the worst case scenario where the raw state is a depolarised Werner state⁶ [NBBB16]. In this type of algorithm, if we have to wait too long for an entangled pair, the whole process must be abandoned and restarted. The entanglement generation rate reported here results in a 95% probability of generating the 8 pairs needed within 72 ms. If we assume the coherence time of the calcium qubit is 50 s [HAB⁺14], then over the time waiting for the entangled pairs, the application qubit’s fidelity has only dropped from 1 to ~ 0.999 , so the photonic link’s *rate* is not the limiting factor, at least for generating a single high-fidelity Bell pair for use between two locations. Moreover, we assert that the infidelities reported in this thesis are largely technical, and can be improved upon in the very near future in the same apparatus.

Building a fault-tolerant quantum computer may not be possible today; however, perhaps the first demonstration of a fully-operational scalable node of a quantum computer is nearer than we might dare to think.

⁶with other states, more efficient algorithms can be devised for higher fidelities.

Bibliography

- [ACNB14] C. Auchter, C.-K. Chou, T. W. Noel, and B. B. Blinov. Ion-photon entanglement and Bell inequality violation with $^{138}\text{Ba}^+$. *Journal of the Optical Society of America B*, 31(7):1568, 2014.
- [AIH84] C. B. Alcock, V. P. Itkin, and M. K. Horrigan. Vapour Pressure Equations for the Metallic Elements: 298–2500K. *Canadian Metallurgical Quarterly*, 23(3):309–313, 1984.
- [AKR⁺18] K. J. Arnold, R. Kaewuam, A. Roy, T. R. Tan, and M. D. Barrett. Blackbody radiation shift assessment for a lutetium ion clock. *Nature Communications*, 9(1):1650, 2018.
- [All11] D. T. C. Allcock. *Surface-Electrode Ion Traps for Scalable Quantum Computing*. Ph.D. Thesis, University of Oxford, 2011.
- [ANK⁺15] N. Akerman, N. Navon, S. Kotler, Y. Glickman, and R. Ozeri. Universal gate-set for trapped-ion qubits using a narrow linewidth diode laser. *New Journal of Physics*, 17(11):113060, 2015.
- [Arl82] G. Arlt. Piezoelectric relaxation. *Ferroelectrics*, 40(1):149–157, 1982.
- [ASS⁺10] D. T. C. Allcock, J. A. Sherman, D. N. Stacey, A. H. Burrell, M. J. Curtis, G. Imreh, N. M. Linke, D. J. Szwer, S. C. Webster, A. M. Steane, and D. M. Lucas. Implementation of a symmetric surface-electrode ion trap with field compensation using a modulated Raman effect. *New Journal of Physics*, 12(5):053026, 2010.
- [BB84] C. Bennett and G. Brassard. Quantum cryptography: public key distribution and coin tossing. pages 175–179, 1984.
- [BBP⁺96] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters. Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels. *Physical Review Letters*, 76(5):722–725, 1996.
- [BCM00] D. R. Barraclough, J. G. Carrigan, and S. R. C. Malin. Observed geomagnetic field intensity in London since 1820. *Geophysical Journal International*, 141(1):83–99, 2000.

- [BEK⁺18] M. Bock, P. Eich, S. Kucera, M. Kreis, A. Lenhard, C. Becher, and J. Eschner. High-fidelity entanglement between a trapped ion and a telecom photon via quantum frequency conversion. *Nature Communications*, 9(1):1998, 2018.
- [Ben80] P. Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*, 22(5):563–591, 1980.
- [Ben95] C. H. Bennett. Quantum Information and Computation. *Physics Today*, 48(10):24–30, 1995.
- [Ber02] D. J. Berkeland. Quantum Information with Trapped Strontium Ions. *Los Alamos Science*, (27):178–183, 2002.
- [BGN⁺18] T. G. Ballance, J. F. Goodwin, B. Nichol, L. J. Stephenson, C. J. Ballance, and D. M. Lucas. A short response time atomic source for trapped ion experiments. *Review of Scientific Instruments*, 89(5):053102, 2018.
- [BHL⁺16] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits. *Physical Review Letters*, 117(6):060504, 2016.
- [BKM16] K. R. Brown, J. Kim, and C. Monroe. Co-designing a scalable quantum computer with trapped atomic ions. *npj Quantum Information*, 2(1):16034, 2016.
- [BLW⁺07] M. Brownnutt, V. Letchumanan, G. Wilpers, R. Thompson, P. Gill, and A. Sinclair. Controlled photoionization loading of 88Sr^+ for precision ion-trap experiments. *Applied Physics B*, 87(3):411–415, 2007.
- [BMDM04] B. B. Blinov, D. L. Moehring, L. M. Duan, and C. Monroe. Observation of entanglement between a single trapped atom and a single photon. *Nature*, 428(6979):153–157, 2004.
- [BPM⁺97] D. Bouwmeester, J.-W. Pan, K. Mattle, M. Eibl, H. Weinfurter, and A. Zeilinger. Experimental quantum teleportation. *Nature*, 390(6660):575–579, 1997.
- [BSC15] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. Measurement of ion motional heating rates over a range of trap frequencies and temperatures. *Physical Review A*, 91(4):041402, 2015.
- [Bur10] A. H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. Ph.D. Thesis, University of Oxford, 2010.

- [BW11] M. Berglund and M. E. Wieser. Isotopic compositions of the elements 2009 (IUPAC Technical Report). *Pure and Applied Chemistry*, 83(2):397–410, 2011.
- [BXN⁺17] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller. Assessing the Progress of Trapped-Ion Processors Towards Fault-Tolerant Quantum Computation. *Physical Review X*, 7(4):041061, 2017.
- [CAL⁺17] C.-K. Chou, C. Aughter, J. Lilieholm, K. Smith, and B. Blinov. Note: Single ion imaging and fluorescence collection with a parabolic mirror trap. *Review of Scientific Instruments*, 88(8):086101, 2017.
- [Che67] G. Chen. On the Physics of Purple-Plague Formation, and the Observation of Purple Plague in Ultrasonically-Joined Gold-Aluminum Bonds. *IEEE Transactions on Parts, Materials and Packaging*, 3(4):149–153, 1967.
- [CL67] R. Calder and G. Lewin. Reduction of stainless-steel outgassing in ultra-high vacuum. *British Journal of Applied Physics*, 18(10):1459–1472, 1967.
- [Coe04] B. Coecke. The logic of entanglement. pages 1–8, 2004.
- [CZ95] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Physical Review Letters*, 74(20):4091–4094, 1995.
- [Deu85] D. Deutsch. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 400(1818):97–117, 1985.
- [DHL⁺05] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. Double-pass acousto-optic modulator system. *Review of Scientific Instruments*, 76(6):063112, 2005.
- [DiV00] D. P. DiVincenzo. The Physical Implementation of Quantum Computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.
- [Drm19] P. Drmota. *Detection of remote entanglement across an elementary trapped-ion quantum network*. Transfer of status report, University of Oxford, 2019.
- [Eke91] A. K. Ekert. Quantum cryptography based on Bell’s theorem. *Physical Review Letters*, 67(6):661–663, 1991.

- [Elt13] A. M. Eltony. *Sensitive, 3D micromotion compensation in a surface-electrode ion trap*. Msc thesis, University of British Columbia, 2013.
- [EMSKB03] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt. Laser cooling of trapped ions. *Journal of the Optical Society of America B*, 20(5):1003, 2003.
- [FC72] S. J. Freedman and J. F. Clauser. Experimental Test of Local Hidden-Variable Theories. *Physical Review Letters*, 28(14):938–941, 1972.
- [Fey82] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 1982.
- [Fra89] J. D. Franson. Bell inequality for position and time. *Physical Review Letters*, 62(19):2205–2208, 1989.
- [GC99] D. Gottesman and I. L. Chuang. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402(6760):390–393, 1999.
- [GRB⁺01] S. Gulde, D. Rotter, P. Barton, F. Schmidt-Kaler, R. Blatt, and W. Hogervorst. Simple and efficient photo-ionization loading of ions for precision ion-trapping experiments. *Applied Physics B*, 73(8):861–863, 2001.
- [Gro97] L. K. Grover. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Physical Review Letters*, 79(2):325–328, 1997.
- [GTL⁺16] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-Fidelity Universal Gate Set for 9Be^+ Ion Qubits. *Physical Review Letters*, 117(6):060505, 2016.
- [HAB⁺14] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit. *Physical Review Letters*, 113(22):220501, 2014.
- [Hay12] D. L. Hayes. *Remote and Local Entanglement of Ions using Photons and Phonons*. Ph.D. Thesis, University of Maryland, 2012.
- [HBD⁺15] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson. Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686, 2015.

- [HDLL⁺08] P. Herskind, A. Dantan, M. B. Langkilde-Lauesen, A. Mortensen, J. L. Sørensen, and M. Drewsen. Loading of large ion Coulomb crystals into a linear Paul trap incorporating an optical cavity. *Applied Physics B*, 93(2-3):373–379, 2008.
- [HGH⁺07] R. Hendricks, D. Grant, P. Herskind, A. Dantan, and M. Drewsen. An all-optical ion-loading technique for scalable microtrap architectures. *Applied Physics B*, 88(4):507–513, 2007.
- [HIV⁺15] D. Hucul, I. V. Inlek, G. Vittorini, C. Crocker, S. Debnath, S. M. Clark, and C. Monroe. Modular entanglement of atomic qubits using photons and phonons. *Nature Physics*, 11(1):37–42, 2015.
- [HOM87] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of subpicosecond time intervals between two photons by interference. *Physical Review Letters*, 59(18):2044–2046, 1987.
- [HOTU11] S. Haze, T. Ohno, K. Toyoda, and S. Urabe. Measurement of the coherence time of the ground-state Zeeman states in 40Ca^+ . *Applied Physics B*, 105(4):761–765, 2011.
- [HSA⁺16] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas. High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves. *Physical Review Letters*, 117(14):140501, 2016.
- [HSZ89] M. A. Horne, A. Shimony, and A. Zeilinger. Two-particle interferometry. *Physical Review Letters*, 62(19):2209–2212, 1989.
- [IMT⁺13] T. Inagaki, N. Matsuda, O. Tadanaga, M. Asobe, and H. Takesue. Entanglement distribution over 300 km of fiber. *Optics Express*, 21(20):23241, 2013.
- [ITU11] Y. Ibaraki, U. Tanaka, and S. Urabe. Detection of parametric resonance of trapped ions for micromotion compensation. *Applied Physics B*, 105(2):219–223, 2011.
- [KES⁺16] C. Kurz, P. Eich, M. Schug, P. Müller, and J. Eschner. Programmable atom-photon quantum interface. *Physical Review A*, 93(6):062348, 2016.
- [KGA⁺11] A. Keselman, Y. Glickman, N. Akerman, S. Kotler, and R. Ozeri. High-fidelity state detection and tomography of a single-ion Zeeman qubit. *New Journal of Physics*, 13(7):073027, 2011.
- [KHR02] A. Kuhn, M. Hennrich, and G. Rempe. Deterministic Single-Photon Source for Distributed Quantum Networking. *Physical Review Letters*, 89(6):067901, 2002.

- [KHT⁺00] N. Kjaergaard, L. Hornekaer, A. Thommesen, Z. Videsen, and M. Drewsen. Isotope selective loading of an ion trap using resonance-enhanced two-photon ionization. *Applied Physics B*, 71(2):207–210, 2000.
- [Kim08] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–1030, 2008.
- [Kit97] A. Y. Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):53–112, 1997.
- [KLH⁺04] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature*, 431(7012):1075–1078, 2004.
- [KMS⁺19] V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. P. Lanyon. Light-matter entanglement over 50 km of optical fibre. *npj Quantum Information*, 5(1):72, 2019.
- [KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, 2002.
- [KPBM15] J. Keller, H. L. Partner, T. Burgermeister, and T. E. Mehlstäubler. Precise determination of micromotion for trapped-ion optical clocks. *Journal of Applied Physics*, 118(10):104501, 2015.
- [LCS99] N. Lütkenhaus, J. Calsamiglia, and K.-A. Suominen. Bell measurements for teleportation. *Physical Review A*, 59(5):3295–3300, 1999.
- [LHM⁺09] L. Luo, D. Hayes, T. Manning, D. Matsukevich, P. Maunz, S. Olmschenk, J. Sterk, and C. Monroe. Protocols and techniques for a scalable atom-photon quantum network. *Fortschritte der Physik*, 57(11-12):1133–1152, 2009.
- [LMR⁺11] M. Lettner, M. Mücke, S. Riedl, C. Vo, C. Hahn, S. Baur, J. Bochmann, S. Ritter, S. Dürr, and G. Rempe. Remote Entanglement between a Single Atom and a Bose-Einstein Condensate. *Physical Review Letters*, 106(21):210503, 2011.
- [LRH⁺04] D. Lucas, A. Ramos, J. Home, M. McDonnell, S. Nakayama, J.-P. Stacey, S. Webster, D. Stacey, and A. Steane. Isotope-selective photoionization for calcium ion trapping. *Physical Review A*, 69(1):012711, 2004.
- [LTGG16] J.-P. Likforman, V. Tugayé, S. Guibal, and L. Guidoni. Precision measurement of the branching fractions of the 5p 2P_{1/2} state in

- 88Sr+ with a single ion in a microfabricated surface trap. *Physical Review A*, 93(5):052507, 2016.
- [LWF⁺17] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger. Blueprint for a microwave trapped ion quantum computer. *Science Advances*, 3(2):e1601540, 2017.
- [LWGS05] V. Letchumanan, M. A. Wilson, P. Gill, and A. G. Sinclair. Lifetime measurement of the metastable 4d 2D_{5/2} state in 88Sr+ using a single trapped ion. *Physical Review A*, 72(1):012509, 2005.
- [LWL⁺19] K. A. Landsman, Y. Wu, P. H. Leung, D. Zhu, N. M. Linke, K. R. Brown, L. Duan, and C. Monroe. Two-qubit entangling gates within arbitrarily long chains of trapped ions. *Physical Review A*, 100(2):022332, 2019.
- [Mau16] P. L. W. Maunz. High Optical Access Trap 2.0. Technical report, 2016.
- [MF13] T. Morimae and K. Fujii. Blind quantum computation protocol in which Alice only makes measurements. *Physical Review A*, 87(5):050301, 2013.
- [MK13] C. Monroe and J. Kim. Scaling the Ion Trap Quantum Processor. *Science*, 339(6124):1164–1169, 2013.
- [MLN⁺99] S. Mannervik, J. Lidberg, L.-O. Norlin, P. Royen, A. Schmitt, W. Shi, and X. Tordoir. Lifetime Measurement of the Metastable 4d 2D_{3/2} Level in Sr+ by Optical Pumping of a Stored Ion Beam. *Physical Review Letters*, 83(4):698–701, 1999.
- [MMBM04] D. L. Moehring, M. J. Madsen, B. B. Blinov, and C. Monroe. Experimental Bell Inequality Violation with an Atom and a Photon. *Physical Review Letters*, 93(9):090410, 2004.
- [MMO⁺07] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe. Entanglement of single-atom quantum bits at a distance. *Nature*, 449(7158):68–71, 2007.
- [MMWZ96] M. Michler, K. Mattle, H. Weinfurter, and A. Zeilinger. Interferometric Bell-state analysis. *Physical Review A*, 53(3):R1209–R1212, 1996.
- [Mon11] T. Monz. *Quantum information processing beyond ten ion-qubits*. Ph.D. Thesis, University of Innsbruck, 2011.

- [MRR⁺14] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, 89(2):022317, 2014.
- [MSW⁺08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-Fidelity Readout of Trapped-Ion Qubits. *Physical Review Letters*, 100(20):200502, 2008.
- [MTT⁺19] B. Merkel, K. Thirumalai, J. E. Tarlton, V. M. Schäfer, C. J. Ballance, T. P. Harty, and D. M. Lucas. Magnetic field stabilization system for atomic physics experiments. *Review of Scientific Instruments*, 90(4):044702, 2019.
- [MVWZ01] A. Mair, A. Vaziri, G. Weihs, and A. Zeilinger. Entanglement of the orbital angular momentum states of photons. *Nature*, 412(6844):313–316, 2001.
- [Nad17] D. Nadlinger. *Entanglement between Trapped Strontium Ions and Photons*. Msc thesis, ETH Zürich, 2017.
- [Nad18] D. Nadlinger. *Towards Remote Entanglement Experiments With Two Mixed-Species Ion Trap Nodes*. D.phil transfer report, University of Oxford, 2018.
- [NBBB16] R. Nigmatullin, C. J. Ballance, N. D. Beaudrap, and S. C. Benjamin. Minimally complex ion traps as modules for quantum communication and computing. *New Journal of Physics*, 18(10):103028, 2016.
- [NFB14] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin. Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links. *Physical Review X*, 4(4):041041, 2014.
- [Nou16] R. J. Nourshargh. *Miniaturised laser systems for ion trap quantum computing*. Msc thesis, University of Oxford, 2016.
- [Par70] J. L. Park. The concept of transition in quantum mechanics. *Foundations of Physics*, 1(1):23–33, 1970.
- [Pau90] W. Paul. Electromagnetic traps for charged and neutral particles. *Reviews of Modern Physics*, 62(3):531–540, 1990.
- [PAVW07] G. Puentes, A. Aiello, D. Voigt, and J. P. Woerdman. Entangled mixed-state generation by twin-photon scattering. *Physical Review A*, 75(3):032319, 2007.

- [PBL95] E. H. Pinnington, R. W. Berends, and M. Lumsden. Studies of laser-induced fluorescence in fast beams of Sr⁺ and Ba⁺ ions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 28(11):2095–2103, 1995.
- [RBA⁺17] C. Robens, S. Brakhane, W. Alt, F. Kleiβler, D. Meschede, G. Moon, G. Ramola, and A. Alberti. High numerical aperture (NA = 0.92) objective lens for imaging and addressing of cold atoms. *Optics Letters*, 42(6):1043, 2017.
- [RLM⁺00] C. F. Roos, D. Leibfried, A. Mundt, F. Schmidt-Kaler, J. Eschner, and R. Blatt. Experimental Demonstration of Ground State Laser Cooling with Electromagnetically Induced Transparency. *Physical Review Letters*, 85(26):5547–5550, 2000.
- [RNH⁺12] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe. An elementary quantum network of single atoms in optical cavities. *Nature*, 484(7393):195–200, 2012.
- [RSK⁺16] T. Ruster, C. T. Schmiegelow, H. Kaufmann, C. Warschburger, F. Schmidt-Kaler, and U. G. Poschinger. A long-lived Zeeman trapped-ion qubit. *Applied Physics B*, 122(10):254, 2016.
- [San12] J. E. Sansonetti. Wavelengths, Transition Probabilities, and Energy Levels for the Spectra of Strontium Ions (Sr II through Sr XXXVIII). *Journal of Physical and Chemical Reference Data*, 41(1):013102–013102–119, 2012.
- [SCB⁺13] A. Stute, B. Casabone, B. Brandstätter, K. Friebe, T. E. Northup, and R. Blatt. Quantum-state transfer from an ion to a photon. *Nature Photonics*, 7(3):219–222, 2013.
- [Sch18] V. M. Schäfer. *Fast gates and mixed-species entanglement with trapped ions*. Ph.D. Thesis, University of Oxford, 2018.
- [SCS⁺12] A. Stute, B. Casabone, P. Schindler, T. Monz, P. O. Schmidt, B. Brandstätter, T. E. Northup, and R. Blatt. Tunable ion–photon entanglement in an optical cavity. *Nature*, 485(7399):482–485, 2012.
- [SDKB09] G. Shu, M. R. Dietrich, N. Kurz, and B. B. Blinov. Trapped ion imaging with a high numerical aperture spherical mirror. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 42(15):154005, 2009.
- [SHKW06] C. Schuck, G. Huber, C. Kurtsiefer, and H. Weinfurter. Complete Deterministic Linear Optics Bell State Analysis. *Physical Review Letters*, 96(19):190501, 2006.

- [Sho97] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [SI03] C. Simon and W. T. M. Irvine. Robust Long-Distance Entanglement and a Loophole-Free Bell Test with Ions and Photons. *Physical Review Letters*, 91(11):110405, 2003.
- [Sin71] J. H. Singleton. Hydrogen Pumping by Sputter-Ion Pumps and Getter Pumps. *Journal of Vacuum Science and Technology*, 8(1):275–282, 1971.
- [SLM⁺12] J. D. Sterk, L. Luo, T. A. Manning, P. Maunz, and C. Monroe. Photon collection from a trapped ion-cavity system. *Physical Review A*, 85(6):062308, 2012.
- [SNM⁺13] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt. A quantum information processor with trapped ions. *New Journal of Physics*, 15(12):123012, 2013.
- [SNN⁺20] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance. High-rate, high-fidelity entanglement of qubits across an elementary quantum network. 2020.
- [SSdRG11] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin. Quantum repeaters based on atomic ensembles and linear optics. *Reviews of Modern Physics*, 83(1):33–80, 2011.
- [Ste96] A. M. Steane. The Ion Trap Quantum Information Processor. *Applied Physics B*, 642(6):623–642, 1996.
- [TGL⁺15] T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland. Multi-element logic gates for trapped-ion qubits. *Nature*, 528(7582):380–383, 2015.
- [TMA⁺12] U. Tanaka, K. Masuda, Y. Akimoto, K. Koda, Y. Ibaraki, and S. Urabe. Micromotion compensation in a surface electrode trap by parametric excitation of trapped ions. *Applied Physics B*, 107(4):907–912, 2012.
- [UTSM⁺07] R. Ursin, F. Tiefenbacher, T. Schmitt-Manderbach, H. Weier, T. Scheidl, M. Lindenthal, B. Blauensteiner, T. Jennewein, J. Perdigues, P. Trojek, B. Ömer, M. Fürst, M. Meyenburg, J. Rarity, Z. Sodnik, C. Barbieri, H. Weinfurter, and A. Zeilinger.

- Entanglement-based quantum communication over 144 km. *Nature Physics*, 3(7):481–486, 2007.
- [VCLB06] K. Vant, J. Chiaverini, W. Lybarger, and D. J. Berkeland. Photoionization of strontium for trapped-ion quantum information processing. pages 1–7, 2006.
- [VHI⁺14] G. Vittorini, D. Hucul, I. V. Inlek, C. Crocker, and C. Monroe. Entanglement of distinguishable quantum memories. *Physical Review A*, 90(4):040302, 2014.
- [VY99] L. Vaidman and N. Yoran. Methods for reliable teleportation. *Physical Review A*, 59(1):116–125, 1999.
- [WBID80] D. J. Wineland, J. C. Bergquist, W. M. Itano, and R. E. Drullinger. Double-resonance and optical-pumping experiments on electromagnetically confined, laser-cooled ions. *Optics Letters*, 5(6):245, 1980.
- [WCJN⁺16] J. D. Wong-Campos, K. G. Johnson, B. Neyenhuis, J. Mizrahi, and C. Monroe. High-resolution adaptive imaging of a single atom. *Nature Photonics*, 10(9):606–610, 2016.
- [WHWM97] L. Westerberg, B. Hjörvarsson, E. Wallén, and A. Mathewson. Hydrogen content and outgassing of air-baked and vacuum-fired stainless steel. *Vacuum*, 48(7-9):771–773, 1997.
- [WKE⁺19] Y. Wan, D. Kienzler, S. D. Erickson, K. H. Mayer, T. R. Tan, J. J. Wu, H. M. Vasconcelos, S. Glancy, E. Knill, D. J. Wineland, A. C. Wilson, and D. Leibfried. Quantum gate teleportation between separated qubits in a trapped-ion processor. *Science*, 364(6443):875–878, 2019.
- [WMI⁺98] D. Wineland, C. Monroe, W. Itano, D. Leibfried, B. King, and D. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259, 1998.
- [WUZ⁺17] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J. N. Zhang, L. M. Duan, D. Yum, and K. Kim. Single-qubit quantum memory exceeding ten-minute coherence time. *Nature Photonics*, 11(10):646–650, 2017.
- [WWKR07] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe. Single-Atom Single-Photon Quantum Interface. *Science*, 317(5837):488–490, 2007.
- [WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.

- [YCL⁺17] J. Yin, Y. Cao, Y.-H. Li, S.-k. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai, G.-B. Li, Q.-M. Lu, Y.-H. Gong, Y. Xu, S.-L. Li, F.-Z. Li, Y.-Y. Yin, Z.-Q. Jiang, M. Li, J.-J. Jia, G. Ren, D. He, Y.-L. Zhou, X.-X. Zhang, N. Wang, X. Chang, Z.-C. Zhu, N.-L. Liu, Y.-A. Chen, C.-Y. Lu, R. Shu, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan. Satellite-based entanglement distribution over 1200 kilometers. *Science*, 356(6343):1140–1144, 2017.
- [YRL⁺12] J. Yin, J.-g. Ren, H. Lu, Y. Cao, H.-l. Yong, Y.-p. Wu, C. Liu, S.-k. Liao, F. Zhou, Y. Jiang, X.-d. Cai, P. Xu, G.-S. Pan, J.-j. Jia, Y.-m. Huang, H. Yin, J.-y. Wang, Y.-a. Chen, C.-z. Peng, and J.-w. Pan. Quantum teleportation and entanglement distribution over 100-kilometre free-space channels. *Nature*, 488(7410):185–188, 2012.
- [YTB⁺12] M. Yamamoto, S. Takada, C. Bäuerle, K. Watanabe, A. D. Wieck, and S. Tarucha. Electrical control of a solid-state flying qubit. *Nature Nanotechnology*, 7(4):247–251, 2012.