High-Fidelity Mixed Species Entanglement of Trapped Ions

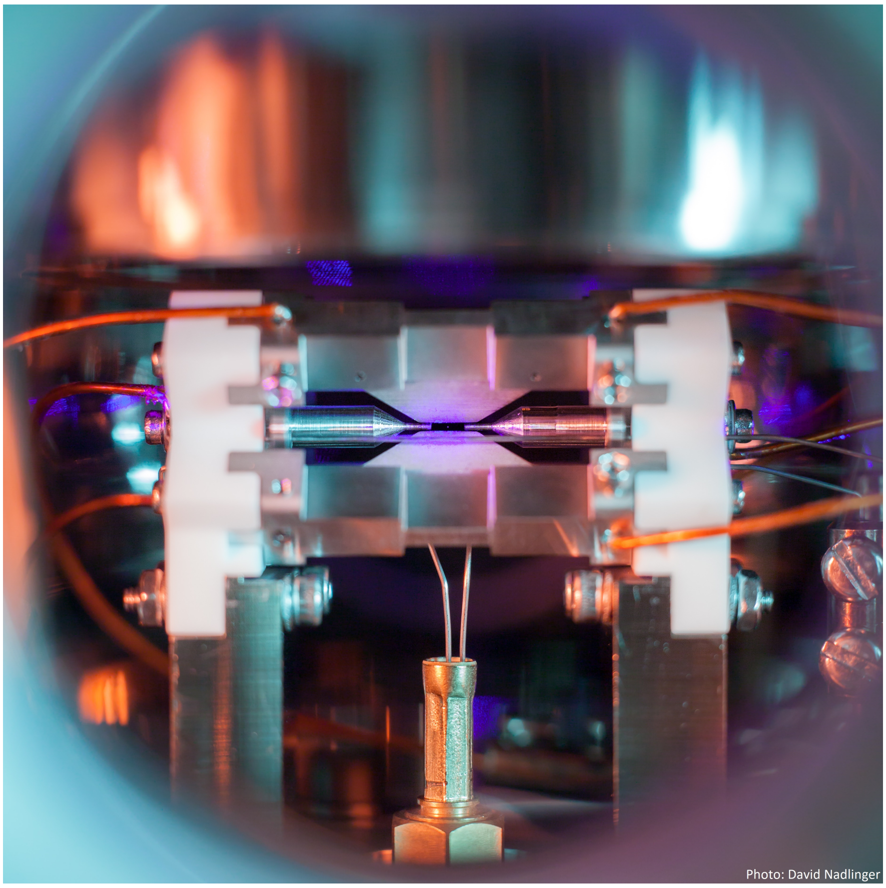

Keshav Thirumalai

A thesis submitted for the degree of Doctor of Philosophy

Oriel College
Michaelmas term, 2019

Department of Physics, University of Oxford

Abstract

High-Fidelity Mixed Species Entanglement of Trapped Ions

Keshav Thirumalai
A thesis submitted for the degree of Doctor of Philosophy
Michaelmas term 2019
Oriel College, Oxford

Harnessing quantum mechanics to solve problems intractable on classical computers is one of the foremost goals of applied physics. Trapped ions make an excellent physical platform for such a quantum computer, as they have demonstrated unparalleled qubit operation fidelities and memory times, but significant challenges remain in scaling up these systems. A promising scalable architecture is a photonically connected network of mixed element trapped ion nodes, each containing a few ions of two species. Poor quality remote entanglement links between nodes can be tolerated, if entanglement between ions in the same node can be achieved with high fidelity. While local gates have been repeatedly demonstrated between ions of the same species, the fidelity of mixed element quantum gates have not previously surpassed the fault tolerant threshold, above which errors can in principle be corrected.

In this thesis we demonstrate an entangling operation between $^{43}\mathrm{Ca}^+$, useful for its long lived qubits, and $^{88}\mathrm{Sr}^+$, well suited for remote entanglement, using a novel mechanism requiring only one laser to manipulate both species. This relatively simple scheme, which is independent of either qubit frequency, allows us to achieve Bell state fidelities of $F = 99.8(1)\%$. We assess the operation of the gate further with state and process tomography, which achieve Bell state fidelity $F = 98.1\%$ and average gate fidelity $\overline{F} = 99.05(6)\%$ respectively, though they are hampered by experimental drifts during the long data collection. Finally, we use interleaved randomised benchmarking to get the best estimate of our gate errors, finding an average gate fidelity of $\overline{F} = 99.62(3)\%$ for a sequence of 60 interleaved gates.

Along with a demonstration of entanglement between two $^{88}\mathrm{Sr}^+$ ions, achieved with fidelity $F = 95.9(4)\%$, this work paves the way towards demonstrating an elementary mixed species quantum network in future experiments.

Acknowledgements

I owe a huge debt of gratitude to my supervisor Prof. David Lucas, for giving me the opportunity to play with atoms and lasers, for providing a constant source of kind advice, guidance and funding, and for always being free to discuss the finer points. Thank you for your encouragement and constructive critique, and for believing in me throughout.

My DPhil has been profoundly shaped by Dr. Chris Ballance, who has taught me almost everything I know about ion trapping, and whose leadership and insight has helped me grow and learn. Thank you for infecting me with the joy of experimental physics, and tolerating my incessant questions with unwavering patience and enthusiasm. I don't think I will ever meet anyone else who somehow seems to always have the answers.

Thank you to Dr. Vera Schäfer, for building so much of the experiment and working tirelessly to keep it running. Along with Amy Hughes, who made a fine and friendly addition to the 'Blade Trap Team', I am indebted to you for your help collecting the data in this thesis.

The Lab 1 community has grown over the years, bringing more joy into an otherwise dingy basement. Thank you to David Nadlinger for your skills in photography and pretty much everything else, to Laurent Stephenson for giving us the gift of music with your speakers, and to Beth Nichol for always making time for a chat. I am grateful for the work of Sarah Woodrow for designing the trap, Benjamin Merkel for stabilising our magnetic field, and Rustin Nourshargh for finding a way to squeeze even more lasers into our lab.

Jochen Wolf and Dr. James Tarlton never failed to make me laugh and were never afraid to mess around, thank you both for your camaraderie and companionship. Dr. Martin Sepiol always had the jokes, even if he didn't want to tell them. Thank you to Dr. Norbert Linke, Dr. Diana Aude Craik and Dr. Hugo Janacek, for making the office a warm and amiable place and helping me prepare for teaching. Thank you also to everyone else in the group, you have all helped me at some point or brightened my day with cake club.

I also appreciate the funding I received from DSTL, and the help of all the departmental support staff who work in the background making this possible. In particular Mathew Newport, who far too often helped me extract a hex screw I'd rounded. I must also thank Prof. Tanja Mehlstäubler and Dr. Jonas Keller for introducing me to this field in the first place, for which I will always be grateful.

I would like to thank all my friends and family, who have kept me going and kept me sane over the years, though rarely both at the same time. Thank you to all my housemates, who always ensured I had a good reason to leave the lab when things weren't going my way. Ellie, your steadfast friendship and polite entertainment of my jabbering are very much appreciated. I am dreading the day when we have to break up our board game collection.

The Oriel College community perennially improved my social life and gave me a place to befriend other nerds, but of a non-physics variety. Thank you also to my Cambridge friends, who kept me grounded by demonstrating how to do things faster and better (I'm talking about you Tom).

Knowing that my family all over the world are so proud of me, has given me immense support. I have learnt so much from them, especially the examples set by my grandparents, on whose shoulders I stand to see so far.

To Amma and Appa, thank you for nurturing my childhood inquisitiveness and always encouraging me to pursue my interests (even if that meant dropping biology!). You have always put Aditya and me first, and yet shown us how to achieve anything no matter the obstacles. Thank you for all your love and support. To my brother Aditya, thank you for being my longest friend, and never letting me take myself too seriously.

To Nisha, I simply could not have done this without you. Meeting you and marrying you has been the best and happiest adventure of my life, and I count myself lucky to have your love and friendship. Thank you for your patience and understanding, and for your limitless practical and emotional support. Your kindness carried me through my most testing months. You are my better half, and without you this thesis would have taken even longer. Thank you for everything.

Contents

1

Introduction

The rise of the computer over the last 70 years, from the first transistors in the late 1940s to supercomputers today executing in excess of 10^{17} floating point operations per second, has revolutionised almost every sphere of life. Calculations that were inconceivable a few decades ago are now routinely done on devices that can fit in your pocket. Despite this fantastic progress however, there are still classes of problems that are intractable with today's computers, which no amount of development will help solve without a fundamental shift in the paradigm of their operation.

The basic building block of the modern computer is the bit; a physical encoding of either a 1 or 0. These bits can appear as two distinct voltage levels in a circuit, or two levels of electrical charge stored in a capacitor, and are the unit of information on which a computer applies logic to carry out a calculation. This system inherently obeys the rules of classical physics and as such can never fully model a fundamentally quantum world. So, since Feynman, Deutsch and others in the 1980s, the idea of using quantum bits, or qubits, at the heart of a quantum simulator or quantum computer has taken root [1,2]. In the decade that followed, algorithms were found for such computers that used the quantum phenomena of superposition and entanglement to speed up particular computations beyond that possible on classical computers [3,4]. Most notably, Shor in 1994 found an algorithm to quickly factorise large numbers, the difficulty of which the widely used RSA encryption scheme relies [5]. This, along with results showing that errors during a quantum computation were not fatal but could be corrected, as long as they were below the 'fault tolerant' threshold and at the expensive of needing more qubits, spurred on rapid progress [6, 7]. A fundamental interest in understanding and gaining precise control of quantum systems, along with the promise of tangible benefits to society, e.g. speeding up drug discovery or machine learning, have led a growing community of research toward developing practical quantum computation [8, 9].

1.1 The Ion Trap Quantum Computer

The first step to building a quantum computer is to choose the physical implementation of the qubit. DiVincenzo in 2000 codified a set of criteria which any such hardware must satisfy [10]. The system must firstly be well-understood and capable of scaling up, with the ability to initialise the qubits into a known state. The information stored in the qubits must last much longer than the time to carry out operations on them, and these operations must form a universal set so that any algorithm can be run. This set can be reduced to individual operations on every qubit and entangling operations between arbitrary pairs of qubits. Finally we must be able to measure the state of individual qubits.

Many physical systems have been proposed as candidates for quantum computation. To date the most developed of these include trapped ions, superconducting circuits, quantum dots, crystal defects, neutral atoms and photons; of these the first two have achieved the greatest success to date. While superconducting circuits promise fast operations and easy scalability leveraging decades of microfabrication technology, they currently suffer from short coherence times and poor qubit connectivity, though these shortcomings are rapidly being alleviated.

The appeal of trapped ions as qubits stems from a simple observation: a floating atom isolated from the environment is the quantum physicist's perfect toy. Each atom of a particular species is identical, with the electronic transition frequencies well above any thermal background radiation. Its electric charge along with cooling from laser beams allows it to be held still by electric fields near its ground state of motion. We encode the qubit in two of the electrons many energy levels by labelling them $|0\rangle$ and $|1\rangle$ (or $|\uparrow\rangle$ and $|\downarrow\rangle$), and can manipulate it using lasers and radio waves. Multiple ions trapped together behave like mutually repulsive marbles in a bowl, the confinement such that when cool they lie near stationary in a line, forming an ion 'crystal' [11]. The collective oscillation modes around their equilibrium positions are used to mediate an entangling interaction between them.

This basic scheme was presented in Cirac and Zoller's seminal 1995 paper, and trapped ions have gone on to demonstrate unparalleled qubit coherence times and single and two-qubit operation fidelities, as well as algorithms on a few qubits [12–18]. However significant challenges remain in scaling up and speeding up these systems while retaining this level of control with low errors. While high fidelity operations (or gates) between two ions have recently been carried out close to their natural speed limit, the largest collection of trapped ions with full control numbers only 20 ions [19, 20]. Just adding more ions to a trap does not work, as the spectrum of their oscillation frequencies becomes increasingly crowded. Picking out a single mode on which to carry out an entangling gate then requires increasingly long operation times.

To solve this problem, a number of architectures have been suggested for a scalable ion trap quantum computer. If multiple trapping zones are built into a single device and ions can be moved around between them, groups of ions of any size can be brought together for cooling, storage, readout and entangling operations as needed [21,22]. This monolithic 'Quantum CCD' however looks daunting to build for a large scale quantum computer containing 100s or 1000s of ions, with the complex ion management system and integrated laser and electronic access needing redesign whenever we wish to scale up the system.

An alternative approach is to use a network of many small repeating units, each a self-contained mini-QCCD of tens of ions [23–25]. Each node is as complex as reliability allows and are connected by optical links, with ions in different nodes entangled by interfering and measuring the photons they emit. Local operations within each node, as long as they are of sufficient fidelity, and nearest neighbour links between nodes, even if they are probabilistic and error-prone, are enough to implement computing with a popular error correcting protocol [26]. Scaling up this system involves producing more identical nodes and connecting them to the network, rather than re-engineering the system as a whole.

1.2 Mixed Species Networked Architecture

Realising this networked architecture needs progress on multiple fronts. While remote entanglement of ions in different nodes has recently been achieved with high fidelity at nearly 200 Hz rates, development is needed on the nodes themselves [27]. Using multiple ion species within each node is likely to be an essential ingredient, to split the diverse roles the ions need to play.

Many networked architecture schemes involve one ion species used for logic operations and memory, and another as an ancilla and for remote entanglement generation. This arrangement serves a number of purposes. At a basic level, a crystal of two ions of differing species reduces cross-talk and needs only global, rather than focussed, beams for individual addressing. High power resonant light or radio waves, needed for remote entanglement generation, qubit control or readout, are transparent to the neighbouring ion and leave it undisturbed. While different isotopes of the same element can be used, this effect is substantially better for two different elements. As direct cooling destroys the information on the logic ion, the ancilla can also be used to sympathetically cool it after transport between trap zones or between gate operations [28]. The two ion species can be chosen to reflect their different roles, using an isotope with nuclear spin for the logic ion, as it gives access to long lived 'clock' qubits, and one without for the ancilla, as the simpler cooling and state preparation increases the quality and speed of remote entanglement attempts.

For these reasons we choose $^{43}Ca^+$ and $^{88}Sr^+$ as our mixed species ion pair, as logic and ancilla/interconnect qubit respectively. Their comparable mass and laser frequencies make them well suited together, for efficient sympathetic cooling and similar laser infrastructure. Indeed two $^{88}Sr^+$ and three $^{43}Ca^+$ ions in a two zone trap are sufficient to realise a node in a quantum computing network [29]. This scheme relies on high-fidelity mixed species entangling gates to purify the lower fidelity remote entanglement generated between $^{88}Sr^+$ ions in neighbouring nodes. The focus of this thesis is demonstrating such a mixed species gate, with a simple experimental scheme and low errors.

1.3 Thesis Outline

In this thesis I describe work in which I played a primary part in a team which consisted at various times of Dr. Vera Schäfer, Amy Hughes and I, with guidance provided by Dr. Christopher Ballance and my supervisor Prof. David Lucas. It is split into the following chapters:

Chapter 2 - Ions motivates our choice of ion species and describes how basic operations, such as trapping, cooling, state preparation and readout are carried out on them.

Chapter 3 - Theory introduces the underlying principles that allow us to control ions with radiation, and describes how we carry out operations on individual ions, control their motion, and how we generate entanglement between them.

Chapter 4 - Apparatus describes the design and construction of the experimental system used to trap the ions, as well as apply the control fields needed for the experiments in this thesis.

Chapter 5 - Two-Qubit Gates presents the setup, results and analysis of experiments entangling a $^{43}Ca^+$ and $^{88}Sr^+$ ion, which is currently being prepared for publication, as well as an experiment entangling two $^{88}Sr^+$ ions.

Chapter 6 - Conclusion provides a summary of this thesis, a comparison to related work in the field, and an outlook on future experiments.

2

Ions

To implement our goal of a mixed species ion trap node, we first need to carefully chose our two ion species. They each need to be well suited for trapped ion information processing in general, and also fulfil their particular roles when trapped together. In this chapter we motivate our choice of ions, describe the multiple qubits available in each species, and go on to detail their ionisation, trapping, and cooling. We then explain how we initialise and measure the qubit states, and how we perform non-entangling operations on them.

2.1 Choice of Ions

When choosing which species of ion to use, we must consider the electronic structure of the atomic element in general, and the particular structure of its isotopes.

2.1.1 Elements

We firstly need a strong cycling transition at wavelengths easily produced by commonly available lasers. For this reason we choose group two alkaline earth metals, as when ionised they have a single valence electron which gives a relatively simple level structure with strong optical transitions. The lightest in the group, Be^+ and Mg^+, unfortunately have their lowest energy transitions from the ground $S_{1/2}$ level well into the ultra-violet (see figure 6.1 at the end of the thesis). Producing laser radiation at these wavelengths directly is difficult and so complex frequency multiplication techniques often need to be used. UV radiation is also difficult to guide along optical fibres, can be absorbed in air and can charge trap structures. Though many important results have been achieved with these elements, we choose to avoid this problem entirely and rule them out.

With increased mass further down the group we can take advantage of more convenient transition wavelengths, lower susceptibility to trap heating, and useful low-lying D levels. The latter can be used to form an optical qubit with the ground level, or to shelve electron population out of the Doppler cooling/fluorescence cycle for high fidelity readout schemes. The downside of choosing a heavier ion are the lower motional frequencies one can achieve when trapped, limiting the speed of two-qubit gates.

2.1.2 Isotopes

When making our choice we also need to decide which particular isotope to use, as isotopes with nuclear spin $I \neq 0$ exhibit hyperfine structure due to the interaction between the spin of the nucleus and the total angular momentum of the electron. Though it is difficult to predict nuclear spin for particular isotopes, isotopes with both an even number of protons and an even number of neutrons have no nuclear spin.

Whilst the hyperfine interaction leads to more complex level structures, it importantly allows for so-called 'clock' transitions. Here, for a particular pair of states in the hyperfine split manifold, at a particular magnetic field, the qubit frequency has a turning point with magnetic field (i.e. $\frac{\mathrm{d}\omega_0}{\mathrm{d}B} = 0$), and so is first order insensitive to it. As magnetic field fluctuations are the dominant source of qubit decoherence in ions, using a clock transition greatly increases the qubit's coherence time.

These transitions can be found either at zero magnetic field between states with $m_F = 0$, or in the so-called 'intermediate' field regime where the Zeeman shifts are no longer linear with field but have curvature, and are of similar magnitude to the hyperfine splitting[1]. Though clock qubits at $B = 0$ are used, it is unfortunately necessary to apply a small magnetic field to define a quantisation axis for laser polarisations. This rather defeats the point, as it re-introduces a small linear magnetic field dependence to the qubit frequency and also lifts the state degeneracy creating problematic nearby off-resonant transitions.

Hyperfine structure generally also allows for high fidelity readout schemes involving only dipole rather than quadrupole transitions, using the multiple extra levels available. As the latter have much narrower linewidths, producing lasers to address them is technically much more demanding. This extra level complexity however is a double edged sword, as most operations need multiple extra laser frequencies to clear population out of them. The extra nearby transitions also slow these operations, to avoid off-resonantly exciting them. For probabilistic remote entanglement generation, which involves repeated cooling, preparation, and excitation cycles, this speed restriction is a significant hindrance.

[1]Though in principle hyperfine structure is not necessary to find field insensitive transitions, applying magnetic fields large enough to be in the intermediate regime with respect to the ion's fine structure is impractical.

2.2 Our Ions

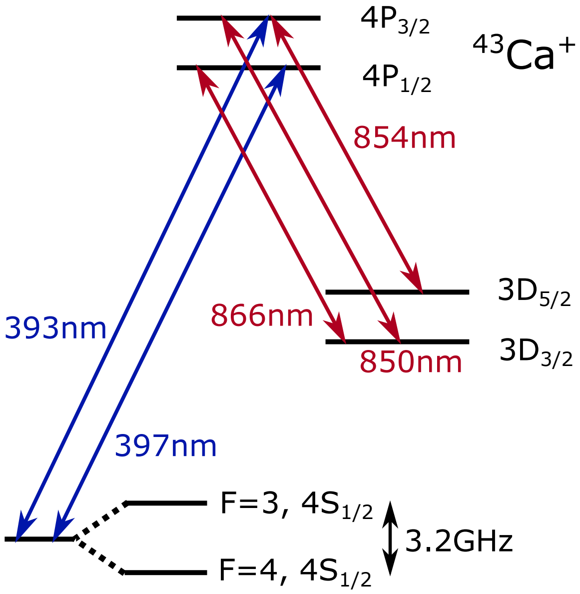

Figure 2.1: Level structure of the $^{43}Ca^{+}$ ion. All the transitions driven by diode lasers in this work are shown, except the Raman lasers (used to bridge the 3.2 GHz qubit transition and perform two-qubit entangling gates) which have been omitted for clarity. Our qubit resides in the hyperfine split ground level.

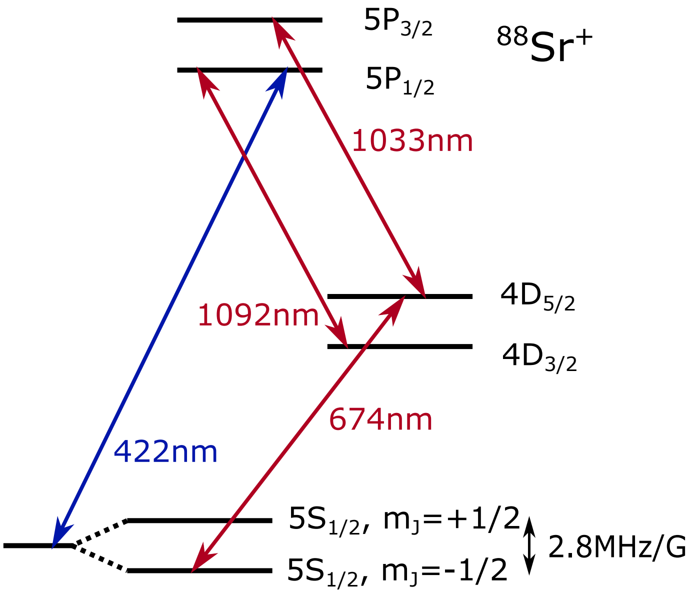

Figure 2.2: Level structure of the $^{88}Sr^{+}$ ion. Note the similarities with the structure of $^{43}Ca^{+}$ in figure 2.1, but with different transition wavelengths. All the transitions driven by diode lasers in this work are shown, along with the 674nm laser which drives a quadrupole transition to the long lived $5D_{5/2}$ 'shelf' level. We use either a qubit in the Zeeman split ground level, or a qubit spread between the ground level and $5D_{5/2}$, for the experiments in this thesis.

With these factors in mind we choose two medium mass elements in the alkaline earth metals, $^{43}Ca^{+}$ and $^{88}Sr^{+}$. They have similar level structures (figures 2.1 and 2.2) with a noble gas core configuration and a single valence electron in the S orbital. The spin-orbit interaction splits the excited P orbital into the fine structure levels $P_{1/2}$ and $P_{3/2}$, and the low-lying D into the $D_{1/2}$ and $D_{5/2}$. Both species have dipole transitions in the blue from S to P and in the infrared from P to D (with the exception of $P_{1/2}$ to $D_{5/2}$ which is forbidden), with a dipole forbidden quadrupole allowed transition from S to D. As this last transition is very weak, the $D_{3/2}$ and $D_{5/2}$ levels have lifetimes of order a second for both species.

The two species also have somewhat similar masses with a ratio of $\sim$2, for efficient sympathetic cooling when trapped together. A larger mass imbalance causes the normal modes of the ion crystal to be more asymmetric (especially radial modes), so modes where the cooling ion has a small eigenvector take longer to cool. In the presence of external heating (for example anomalous trap heating [30]), this can also increase the final Doppler cooling temperature.

The 397 nm $S_{1/2}$ - $P_{1/2}$ transition of $^{43}Ca^{+}$ and the 408 nm $S_{1/2}$ - $P_{3/2}$ transition of $^{88}Sr^{+}$ are also fairly close in frequency, which as we shall see, is important for our entangling gate mechanism.

2.2.1 Calcium

The key difference between the two species is the presence of hyperfine structure in $^{43}Ca^{+}$, as it is an odd isotope[2]. With a nuclear spin of $I = 7/2$, the $S_{1/2}$ ground level is split into $F = 3$ and $F = 4$ hyperfine manifolds separated by 3.226 GHz. We can choose any pair of states between the $F = 3$ and $F = 4$ hyperfine manifolds to form our qubit, but we make use of two particularly useful ones. At the end of the manifold

[2] In fact $^{43}Ca^{+}$ is the only stable odd isotope, and as calcium has an even number of protons it is thus the only one with hyperfine structure.

we form the 'stretch qubit' between $|F=3, m_F=3\rangle$ and $|F=4, m_F=4\rangle$. Though it is magnetically sensitive (2.45 MHz/G), the ability to prepare the qubit state directly via optical pumping makes it very useful. We can also use the 'clock qubit' between $|F=3, m_F=1\rangle$ and $|F=4, m_F=0\rangle$, which has a turning point of qubit frequency with magnetic field at 146.0942 G. This qubit has been used extensively in previous work and exhibits both a long coherence time of $T_2^* = 50\,\mathrm{s}$ [14] and very low memory errors over short timescales ($\epsilon < 10^{-4}$ for $t \lesssim 50\,\mathrm{ms}$) [31]. Due to the usefulness of this qubit in realising a practical quantum computer, all experiments in this thesis are conducted at this magnetic field. At this field the qubit frequency for the stretch qubit is $\omega_0 = 2\pi \times 2.874\,\mathrm{GHz}$ and for the clock qubit is $\omega_0 = 2\pi \times 3.200\,\mathrm{GHz}$.

At this field we are also in the intermediate regime for all the levels shown in figure 2.1, except $D_{5/2}$ where we are in the strong field regime as its hyperfine splitting is small. For convenience we still label the $S_{1/2}$ and $P_{1/2}$ levels with F and m_F, as the states are still divided into two distinct manifolds. For the $P_{3/2}$ and $D_{3/2}$ the states are completely scrambled with respect to their zero field ordering, and the only good quantum number is the total angular momentum projection $m = m_I + m_J = m_F$, which along with the energy can be used to label the states. In $D_{5/2}$ m_I and m_J are good quantum numbers.

2.2.2 Strontium

In the $S_{1/2}$ ground level of $^{88}Sr^+$ we can choose only the two Zeeman states to form a qubit, which has energy $\Delta E = g_J \mu_B m_J B$ given by the Zeeman splitting. It has similar magnetic sensitivity to the $^{43}Ca^+$ stretch qubit (2.8 MHz/G), and at 146 G the qubit frequency is $\omega_0 = 2\pi \times 408.8\,\mathrm{MHz}$. We can also form an optical qubit between $S_{1/2}$ and $D_{5/2}$ which, whilst being metastable with decay lifetime $\tau = 390.8\,\mathrm{ms}$, has the upper qubit level out of the Doppler cooling cycle for easy state readout. We choose the $S_{1/2}\,|m_J=-1/2\rangle$ and $D_{5/2}\,|m_J=-3/2\rangle$ states with $|\Delta m_J| = 1$, as for our beam geometry they minimise the qubit frequency sensitivity to magnetic field (1.12 MHz/G). At 146 G the qubit transition is the −163.5 MHz Zeeman shifted component of the 674 nm quadrupole transition.

2.3 Photo-Ionisation

In the experiment we have two tubes filled with granules of calcium and strontium metal respectively, which can independently be heated by passing a current through them. When heated a hot atomic vapour streams out of a small hole in the side of each 'oven' towards the centre of the trap.

To trap the atoms we need them to have an overall charge, so we remove an electron using a two-stage laser ionisation process. We first excite the outer electron from its ground $^1S_0$ level using a strong transition to the $^1P_1$ level, before further exciting it to the continuum with a second laser. As the first stage is resonant to a transition with linewidth smaller than the relevant isotope shifts, we can use the laser frequency to select particular isotopes without needing an isotopically pure atomic source. The first stage is driven by 423nm and 461nm lasers for calcium and strontium respectively, while the second stage is driven by a common 375nm laser.

2.4 Linear Paul Trap

To trap our ions we make use of their electric charge to confine them in a 3D harmonic potential well. No arrangement of charges can produce a static electric field minima in free space, where all field lines are pointing towards a point (Earnshaw's theorem), as can be seen from Gauss's Law $\nabla \cdot \mathbf{E} = \oiint \mathbf{E} \cdot \mathrm{d}\mathbf{S} = 0$, where $\mathbf{S}$ is a closed surface. If we confine in two dimensions we anti-confine in the third creating a saddle point.

We trap our ions using a linear Paul trap which consists of two 'endcap' electrodes held at V_{DC} applying a static field in the axial (z) direction and 4 'blade' electrodes, two at ground and two oscillating at frequency Ω with amplitude V_{RF} applying a radio frequency (RF) oscillating saddle shaped quadrupole potential in the radial (x-y) plane. Around the centre of the trap the ion sees a confining 3D harmonic potential

$$U = Q_z^{\mathrm{DC}}[z^2 - \frac{1}{2}(x^2 + y^2)] + [Q_z^{\mathrm{RF}} z^2 + Q_x^{\mathrm{RF}} x^2 - Q_y^{\mathrm{RF}} y^2] \cos(\Omega t) \,, \tag{2.1}$$

where

$$Q_z^{\mathrm{DC}} = \frac{\alpha_z V_{\mathrm{DC}}}{z_0^2} \qquad Q_z^{\mathrm{RF}} = \frac{\alpha_z V_{\mathrm{RF}}}{2 z_0^2} \qquad Q_x^{\mathrm{RF}} = \frac{\alpha_x V_{\mathrm{RF}}}{\rho_0^2} \qquad Q_y^{\mathrm{RF}} = \frac{\alpha_y V_{\mathrm{RF}}}{\rho_0^2} \,.$$

The parameters α are electrode geometry factors, where $\alpha_{x,y} = 1/2$ and $\alpha_z = 0$ for infinitely long hyperbolic shaped blades[3], and ρ_0 and z_0 are the distances from the ion to the blades and endcaps respectively. The DC field confines in the axial direction and anti-confines in the radial direction, while the RF field, averaged over one period, confines in the radial and axial directions, the latter due to the asymmetric drive of the blades [32]. The equations of motion for a particle of charge $+Ze$ and mass m are separable and can be written in the standard Mathieu equation form

$$\frac{\mathrm{d}^2 u}{\mathrm{d}\tau^2} + [a_u - 2q_u \cos(2\tau)]u = 0 \,, \tag{2.2}$$

where $u \in \{x, y, z\}$, $\Omega t = 2\tau$, and

$$\begin{aligned} a_x &= -\xi Q_z^{\mathrm{DC}} & \qquad q_x &= -\xi Q_x^{\mathrm{RF}} \\ a_y &= -\xi Q_z^{\mathrm{DC}} & \qquad q_y &= +\xi Q_y^{\mathrm{RF}} \\ a_z &= +2\xi Q_z^{\mathrm{DC}} & \qquad q_z &= -\xi Q_z^{\mathrm{RF}} \end{aligned}$$

[3] $\alpha_{x,y} = 1/2$ rather than 1, as we drive the blades asymmetrically.

$$\xi = \frac{4Ze}{m\Omega^2} .$$

We can find approximate solutions of the form $u = \overline{u} + \delta$ in which the ion undergoes secular motion $\overline{u} = A_u \cos(\omega_u t)$, where the secular frequency is

$$\omega_u = \frac{\Omega}{2}\sqrt{\frac{q_u^2}{2} + a_u} , \tag{2.3}$$

and micromotion

$$\delta = \frac{\overline{u}q_u}{2}\cos(\Omega t) , \tag{2.4}$$

under the assumptions that $\delta \ll \overline{u}$ and $\omega_u \ll \Omega$ [33, 34]. In the axial direction the DC harmonic potential dominates the RF, so $q_z \ll a_z$ and hence the axial secular frequency is approximately

$$\omega_z \approx \frac{\Omega}{2}\sqrt{a_z} = \sqrt{\frac{2ZeQ_z^{\mathrm{DC}}}{m}} \propto \sqrt{\frac{V_{\mathrm{DC}}}{m}} . \tag{2.5}$$

In the radial direction both the RF trapping and DC anti-trapping are relevant, with the secular frequency given by

$$\omega_r = \frac{\Omega}{2}\sqrt{\frac{q_r^2}{2} - \frac{2\omega_z^2}{\Omega^2}} = \sqrt{\left(\frac{\sqrt{2}\,ZeQ_r^{\mathrm{RF}}}{m\Omega}\right)^2 - \frac{ZeQ_z^{\mathrm{DC}}}{m}} \mathrel{\underset{\sim}{\propto}} \frac{V_{\mathrm{RF}}}{m\Omega} \tag{2.6}$$

where $r \in \{x, y\}$. As we drive the RF electrodes asymmetrically, there is a small difference between α_x and α_y, which leads to two radial modes at similar frequency.

The secular confining force in each direction is that of a simple harmonic oscillator $F_u = -m\omega_u^2\overline{u}$. We can see therefore that, approximately, the axial confining force is independent of the ion mass while the radial confining force is proportional to $1/m$.

With multiple cold ions in the trap, their Coulomb repulsion is balanced by the trapping fields, and they form a stable arrangement known as a crystal. We ensure that the ratio ω_r/ω_z is sufficiently large such that the ions lie in a straight line along the trap axis, with the separation between ions increasing away from the centre. The equilibrium positions from the trap centre for N ions are given by

$$\overline{z}_N^i = \left(\frac{k_N^i Z^2 e^2}{4\pi\epsilon_0 m \omega_z^2} \right)^{1/3} \approx \left(\frac{k_N^i Z e z_0^2}{8\pi\epsilon_0 \alpha_z V_{\mathrm{DC}}} \right)^{1/3} \tag{2.7}$$

where k_N^i are scaled positions with $k_1 = \{0\}, k_2 = \{-1/4, 1/4\}, k_3 = \{-5/4, 0, 5/4\}$. For $N \geq 4$ these values have to be found numerically, and are tabulated and approximated in [35] and [36] respectively. Note that the spacing is independent of the ion's mass under the same assumptions of equation 2.5, namely that the axial confinement due to the RF potential is negligible compared to that due to the DC potential [37].

The motion of the ions can be described by a quantum harmonic oscillator for each of the modes, with spacing $\hbar\omega$. The ground state spread of the wavefunction for each ion is $\tilde{z} = \sqrt{\hbar/2m\omega}$. For two $^{43}Ca^+$ ions in a typical 1.9 MHz axial trap the ions are separated by $\Delta\overline{z} \approx 3.5\,\mu$m, orders of magnitude further apart than their axial wavefunction spread of $\tilde{z} \approx 8\,$nm. We can therefore reasonably think of the ions as a string of small wavepackets in a line, centred at their equilibrium positions, and not overlapping with their neighbours [36]. Their separation is also too large for any appreciable interaction between their spins [38].

2.4.1 Normal Modes

There are N normal modes of oscillation in each axis for an ion string of N ions. For two ions these are an in-phase and an out-of-phase mode in each of the axial and two radial directions. For two ions of the same mass the in-phase mode, known as the centre of mass (COM) mode, is at the trap secular frequency for that axis and the other mode frequencies have simple relations to it. The axial 'stretch' mode has frequency $\omega_{\mathrm{str}} = \sqrt{3}\,\omega_z$ and the radial 'rocking' mode has frequency $\omega_{\mathrm{rock}} = \sqrt{\omega_r^2 - \omega_z^2}$. The eigenvectors for these modes, which give each ion's normalised amplitude, are $b_z = b_r = \{1, 1\}/\sqrt{2}$ for the in-phase modes and $b_{\mathrm{str}} = b_{\mathrm{rock}} = \{1, -1\}/\sqrt{2}$ for the out-of-phase modes, which are purely symmetric and purely anti-symmetric respectively. For same-species crystals with $N \geq 4$ the mode frequencies and eigenvectors need to be found numerically, and are again tabulated in [35, 36].

For ion strings of differing species expressions for the mode frequencies and eigenvectors are more complex, but for the two ion case they can be found in [39] and [40]. The modes are no longer purely symmetric or purely anti-symmetric, and they depend principally on the mass ratio between the two ions. Dissimilar ions have modes with asymmetric eigenvectors, where one ion moves with a large amplitude while the other is almost still. This asymmetry is worse for the radial modes compared to the axial modes. For the ions is this thesis the eigenvectors of the axial modes are $b_{\mathrm{ip}} = \{0.453, 0.892\}$ for the in-phase mode and $b_{\mathrm{oop}} = \{-0.892, 0.453\}$ for the out-of-phase mode, where the elements are in the order $\{^{43}\mathrm{Ca}^{+}, ^{88}\mathrm{Sr}^{+}\}$.

2.4.2 Micromotion

For a single ion in the centre of the RF potential (the RF null), the micromotion amplitude is zero. Stray static electric fields, or imperfections in the trap geometry such that the DC and RF potential nulls do not overlap, causes the ion to sit away from the RF null and experience micromotion at the trap RF frequency. We see micromotion when the ion is displaced either radially or axially, the latter because of the small axial confinement from the RF. In general we want to minimise the micromotion amplitude at all times, as excess micromotion depletes the main ‘carrier’ term of all beams on the ion, adds unwanted sidebands to the scanned spectra of transitions, and in extreme cases can cause heating of the ion secular motion.

To return the ion to the RF null, the trap has three compensation electrodes running parallel to the trap axis that can move the ion in any direction in the radial plane. Applying a small asymmetry to the endcap voltages also allows us to move the ions axially. With these controls we can trim out the offending stray electric field at the

RF null.

With multiple ions in the trap we can only minimise rather than eliminate micromotion, as the ions naturally sit displaced axially from the RF null. For a mixed species pair of ions poor compensation creates an additional problem, as a radial displacement of the crystal also causes them to tilt with respect to the trap axis. This occurs as the radial confining force is mass dependent, pushing lighter ions closer to the trap axis.

Compensation

To compensate the trap we measure the micromotion of a single ion using a variety of methods, and adjust the compensation and endcap electrodes to minimise it. The standard method of measurement involves correlating the arrival times of fluorescence photons at our detector, with the phase of the trap RF [34]. As the Doppler shift induced by the micromotion causes the frequency of the Doppler cooling lasers (as seen by the ion) to move towards and away from resonance, we see a modulation of the fluorescence rate at the trap frequency. This needs to be done separately with three non-coplanar beams, to resolve the direction and magnitude of the ion's displacement. The signal can be magnified by reducing the trap RF strength, which increases the stray field displacement of the ions. While minimising this signal in all three directions compensates the trap, in practice it is tricky as neither the compensation electrodes nor probe laser beams form an orthogonal basis set, with no common vector between them.

An alternative method uses the ion imaging camera, to watch the movement of the ions as the trap RF strength is toggled. The ions move in the direction opposite to their displacement from the trap RF null when the trap strength is reduced, until at the null their position is nearly independent of trap strength. This works effectively in the horizontal plane of the trap, but vertical ion displacements are visible only as defocussing of the ion image, so fluorescence correlation compensation is used in this direction instead

Finally, we have a highly accurate method to measure the axial compensation, by looking at the micromotion sidebands of the motionally sensitive Raman beam pair. Their geometry has been set carefully so their difference k-vector lies exactly on the trap axis, so measuring the Rabi frequency of these sidebands allows accurate minimisation of the axial compensation. For two well compensated ions in our trap, with typical trap parameters, the residual axial micromotion sidebands have a Rabi frequency a quarter that of the carrier.

In practice we use a combination of all three methods, setting the rough compensation with the camera, and the final compensation in the radial horizontal direction, radial vertical direction and axial direction with the camera, fluorescence correlations and Raman sideband spectroscopy respectively. At the lower trap strength needed for loading, the compensation voltages required differed, suggesting thermal trap geometry changes with reduced RF power. Thermal cycling of the trap over many weeks, as well as a changing lab environment, required the compensation to be updated periodically.

2.4.3 Recrystallisation

Collisions from background gas molecules can significantly and suddenly heat the ions in the trap, occasionally with enough force to remove them from the trap entirely. Though the latter is rare in our deep trapping potential, we often have to recover from collisions that de-localise the ions from their crystalline arrangement and kick them to high orbits in the trapping potential.

To detect such a decrystallisation event, we check the ion fluorescence is above a certain threshold between each shot of the experiment. If below the threshold, we discard the previous data point and continue. To avoid false positives, we wait until three below threshold shots occur in succession, after which we recrystallise the ions. This is done by weakening the RF and DC trap strengths while detuning the Doppler cooling lasers to increase their effectiveness. A sequence of different trap strengths are used, that were empirically found to reliably recrystallise the ions, with the whole procedure taking approximately 3 seconds.

2.4.4 Ion Order

In principle the ordering of a mixed species crystal does not need to be kept constant in the trap, as long as all beams are aligned to illuminate both ions equally, but in practice experimental drifts necessitate that we do so.

As we do not image the ions with the camera during most experiments, we cannot detect if they have swapped position. This however is unlikely without a preceding decrystallisation event, so we only re-order the ion chain after recrystallisation. We start by weakening the RF confinement so that the ions move from an axial line to a radial line arrangement. We then reduce the voltage on one of the endcaps substantially, so the crystal is pushed axially away from the trap centre. At this location, the axial confinement due to the RF is different enough for each species, due to their differing mass, so the radial crystal tilts to one side, with the lighter ion closer to the trap centre. Returning the RF *then* the endcap voltages back to their usual values leaves the ions in a deterministic order. This procedure takes approximately 0.3 s. Similar schemes have been performed in other experiments, as described in [41].

2.4.5 Finding Mode Frequencies

To calibrate our two-qubit gate, it is important to find the motional mode sideband frequency on which we apply the gate precisely. Using laser spectroscopy unfortunately has a significant drawback, as the applied field causes an unavoidable light-shift on the qubit carrier, which then shifts the measured mode frequency. While the light-shift reduces with reducing laser intensity, it cannot be removed entirely.

Instead we apply a perturbing voltage to one of the trap endcaps, and scan its frequency while looking at the ion fluorescence. This 'tickle' voltage excites the ion motion when on resonance, which causes Doppler shifts that make the laser frequency oscillate from the ion's perspective. Due to the non-linear relationship between laser detuning and fluorescence, this causes a change in fluorescence averaged over a motional cycle, which appears as a peak or trough when the 'tickle' voltage is on resonance. To maximise this signal, we choose to sit at one of the points of maximum curvature on the half-Lorentzian laser spectrum, at either 1/4 or 3/4 of the peak. This allows us to resolve the motion mode frequency to $\sim$100 Hz.

2.5 Doppler Cooling

We now have trapped ions, but we still need to cool and see them. We can do both by shining laser light tuned to the red of the $S_{1/2}$ - $P_{1/2}$ transition, at 397nm for calcium and 422nm for strontium. Photons are preferentially absorbed when the oscillating ion moves in the direction of the laser beam and sees the light frequency Doppler shifted closer to resonance. After absorbing this lower energy photon with its momentum pointing in the opposite direction to the ion's motion, another is emitted in a random direction at a likely higher frequency closer to resonance. A small proportion of these photons are collected by the imaging system in the experiment, so we can detect the ion. As more energy leaves the system than enters, and the average change of momentum is against the motion, the ion is cooled down. We ensure that the laser beam direction, and hence the photon k-vector, has a projection onto all 3 trap axes so we can cool all the normal modes of the system.

In $^{43}Ca^+$ and $^{88}Sr^+$ there is a $\sim$6% probability (branching ratio) that the electron population will decay from the $P_{1/2}$ level to the $D_{3/2}$ rather than to the $S_{1/2}$. To avoid population being optically pumped to the $D_{3/2}$ level and the ion quickly going dark, we repump along this transition also (866nm for calcium and 1092nm for strontium). As we are at 146 G the 422 nm transition in $^{88}Sr^+$ is split into four Zeeman components. We drive the outer pair with two $\hat{\sigma}^{\pm}$ polarised beams separated by 545 MHz. In $^{43}Ca^+$ we use two 397nm beams with $\hat{\sigma}^+$ and $\hat{\pi}$ polarisation, each of which have an EOM induced sideband to hit both hyperfine manifolds of the $S_{1/2}$ level, and need to power broaden the 866nm beam to cover the hyperfine splitting of the $D_{3/2}$ level.

The ion is cooled asymptotically to the Doppler limit, where the cooling rate equals the heating rate due to the recoil momentum of the emitted photons. For a two level system with a very low intensity laser optimally red-detuned by half a natural linewidth Γ, this is roughly [40, 42]

$$k_B T \approx \frac{\hbar\Gamma}{2}\,. \tag{2.8}$$

If we compare this to the motional energy spacing $\hbar\omega$ of the trap, we can see that the thermal mean mode occupation $\overline{n} \approx \Gamma/2\omega$. For typical trap strengths of $\omega_z = 2\pi \times 1.9\,\text{MHz}$ and $\omega_r = 2\pi \times 4.2\,\text{MHz}$, and with $\Gamma = 2\pi \times 23\,\text{MHz}$ for the $P_{1/2}$ level in Ca^+, this gives $\overline{n} \approx 6$ and $\overline{n} \approx 3$ for the axial and radial modes respectively. In reality the situation is more complex, especially in $^{43}Ca^+$ where its hyperfine structure leads to many dark resonances and population spread over many more states, giving a higher Doppler limit and in general lower fluorescence.

2.5.1 Dark Resonance Cooling

We can cool beyond the Doppler limit if we can find regions in the ion's fluorescence spectrum with both low fluorescence and a steep gradient of fluorescence with laser frequency, to minimise heating and maximise cooling. Just such conditions exist adjacent to the many dark resonances in the complex spectrum of isotopes with hyperfine structure, for example $^{43}Ca^+$. Using parameters similar to [43], empirically adjusted for our beam geometry and polarisations, we find the side of a dark resonance feature on which to cool, needing only a lowering of intensity in the 397nm and 866nm beams and a small change in frequency of the latter.

2.6 Sideband Cooling

If the natural linewidth γ of a transition is such that $\gamma \ll \omega_z$, we can resolve the motional sidebands around the transition as long as our probe Rabi frequency is kept low enough to not excessively broaden the spectrum.

Driving the lower frequency ('red') sideband of a qubit transition flips the spin, but also removes a motional quanta of energy. If we can return the electron population back to the starting qubit state and repeat this process, we can cool the ion in principle down to the motional ground state. This can be done in two distinct ways, either continuously applying the red sideband excitation while repumping from the upper state, or applying a repeated cycle of a red sideband π-pulse followed by repump pulses. While continuous sideband cooling achieves a higher cooling rate, important initially to clear population from higher motional states, pulsed sideband cooling achieves a lower final temperature. The final temperature is ultimately limited by off-resonant excitation of the qubit carrier transition, after which the population can decay along the blue sideband adding a motional quanta.

In $^{43}\mathrm{Ca}^{+}$, we use the Raman beams, with difference vector along the trap axis, to excite the red sideband. A weak σ^{+} polarised 397 nm beam, is used to remove population from the $|\uparrow\rangle$ qubit state only, which then decays back to $|\downarrow\rangle$, with a simultaneous 866 nm beam clearing population from the $\mathrm{D}_{3/2}$ as before.

In $^{88}\mathrm{Sr}^{+}$ we use the 674 nm quadrupole laser, which has projection onto all the trap axes, to excite the red sideband transition of the optical qubit. The upper qubit state in the $\mathrm{D}_{5/2}$ level is cleared out with the 1033 nm laser to $\mathrm{P}_{3/2}$, though as this can decay back to both states in the ground level, the lower frequency 422 nm laser and 1092 nm laser are used to pump to $|\downarrow\rangle$, as described in the next section.

2.7 Qubit Preparation and Readout

In $^{88}Sr^+$ we prepare the Zeeman and optical qubits in the same way, using the lower frequency 422 nm beam to optically pump to the common $|{\downarrow}\rangle$ state, $S_{1/2}\,|m_J = -1/2\rangle$. In $^{43}Ca^+$ we optically pump to the $S_{1/2}\,|F = 4, m_F = 4\rangle$ state at the end of the manifold using the $\hat{\sigma}^+$ polarised 397nm beam, to form the $|{\downarrow}\rangle$ state of the stretch qubit. To prepare the clock qubit we first prepare the stretch qubit before applying a sequence of four microwave pulses to transfer the population along the ground level to $|F = 3, m_F = 1\rangle$ (for details see [44]).

When reading out the state of the ion we apply the Doppler cooling beams and look for ion fluorescence, by counting the number of photons that arrive at our detector in a given readout time. If the ion is bright (i.e. our count is above a certain threshold) we know the ion was in the $|{\uparrow}\rangle$ state, and if it is dark (i.e. our count is below the threshold) we know the ion was in the $|{\downarrow}\rangle$ state. For the $^{88}Sr^+$ optical qubit we can do this directly as the upper qubit state in the $D_{5/2}$ level is out of the Doppler cooling cycle[4]. For the ground level qubits however we must first selectively shelve the population from one of the qubit states out of the Doppler cooling cycle before we apply the Doppler cooling beams and count photons. In both species we shelve the $|{\downarrow}\rangle$ state population to the $D_{5/2}$ level. As this level has a finite lifetime, the readout time is limited to prevent errors from shelf decay, but set long enough to cleanly discriminate between bright and dark photon counts.

[4]For this qubit $|{\uparrow}\rangle$ is therefore dark rather than bright.

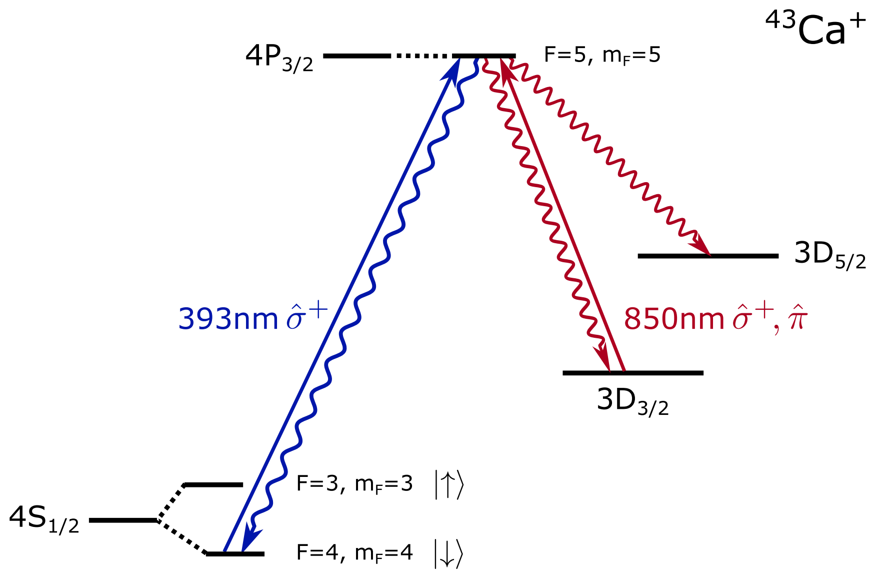

Figure 2.3: Readout scheme for ^{43}Ca$^+$. A 393 nm $\hat{\sigma}^+$ polarised pulse selectively excites the $|{\downarrow}\rangle$ qubit state to $P_{3/2}\,|F=5, m_F=5\rangle$, which cannot decay to any other state in the ground level. An 850 nm pulse split into three frequency components (two $\hat{\sigma}^+$ and one $\hat{\pi}$ polarised), is used to clear the $D_{3/2}$ level. Electron population is pumped to the $D_{5/2}$ shelf, which is out of the Doppler cooling cycle, so if the ion does not fluoresce when cooling light is subsequently applied, it must have been in the $|{\downarrow}\rangle$ state.

To shelve the Zeeman qubit in ^{88}Sr$^+$ we apply a 674 nm π-pulse identical to that used to apply a bit-flip to the optical qubit. With the added hyperfine structure and without the quadrupole laser in ^{43}Ca$^+$ the situation is trickier. Using a frequency resolved $\hat{\sigma}^+$ polarised 393 nm beam, population is transferred from the stretch qubit down state $S_{1/2}\,|F=4, m_F=4\rangle$ to $P_{3/2}\,|F=5, m_F=5\rangle$[5] (figure 2.3). From here selection rules allow decay either to the $D_{5/2}$ shelf, three states in $D_{3/2}$ or back down to $S_{1/2}\,|F=4, m_F=4\rangle$. As there is no decay to other ground level states we avoid mixing population between the qubit states. After the 393 nm pulse, two $\hat{\sigma}^+$ and one $\hat{\pi}$ polarised frequency resolved 850 nm pulses clear out the $D_{3/2}$ states. This set of $\{393, 850\sigma_{1,2}, 850\pi\}$ pulses is repeated 8 times in total followed by a final 393 nm pulse, shelving the population into the $D_{5/2}$. This method, as optimised in [44], is both

[5]Though strictly speaking F is not a good quantum number here, see section 2.2.1, we use $|F, m_F\rangle$ for convenience to label the state one would arrive at if the magnetic field were adiabatically ramped up [44].

efficient and prevents exciting population to $P_{3/2}$ states other than $|F = 5, m_F = 5\rangle$, which would decay to other ground level states. To readout from the clock qubit, we first need to map its state back to the stretch qubit, using the same microwave pulses used to prepare it but in reverse.

2.7.1 Multiple Ion Readout

To readout the state of two distinguishable ions in a crystal, e.g. $^{43}Ca^+$ and $^{88}Sr^+$, we sequentially shelve, illuminate and collect photons from each species. For two ions of the same species however, since we have no individual addressing, we can only shelve and illuminate with global beams, and look for how many ions are fluorescing. We then see three levels of fluorescence from the ions; either both ions are dark, one ion is bright, or both are bright. Setting two thresholds between them therefore allows us to measure these possibilities, but when one ions is bright we cannot determine which of them it is.

2.8 Single Qubit Operations

The $^{88}Sr^+$ optical qubit is directly driven by a 674 nm narrow linewidth laser along its quadrupole transition. The Zeeman qubit in $^{88}Sr^+$ and the stretch and clock qubits in $^{43}Ca^+$ are all driven directly using RF and microwaves respectively, that are applied to the trap blade electrodes. The $^{43}Ca^+$ qubits can also be driven by a Raman transition using a pair of phase locked lasers detuned from the $P_{1/2}$ level.

3

Theory

In this chapter I explain the origin of the light-matter interaction, which underpins how our control fields interact with the electronic structure of the ions. I describe how we control the state of each qubit, with a field resonant to its frequency, or via two fields coupled to an upper level. The coupling of these fields to the ion's motion is described, as well as how operations on its motion are achieved.

I go on to explain the principles behind our two-qubit entangling operations, which use the ions shared motion to couple the qubits together. I briefly describe the details of the two gate mechanisms used in this thesis, along with an extension to the mixed species case. Finally, I survey common error sources for these operations, and methods to mitigate them.

3.1 Multipole Expansion of the Atom-Light Interaction

The Hamiltonian for a single valence electron[1] atom in an electromagnetic field characterised by vector potential $\mathbf{A}$ and scalar potential φ is [45–48]

$$H = \frac{1}{2m}\Big[\boldsymbol{\sigma}.(\mathbf{p} + e\mathbf{A}(\mathbf{r},t))\Big]^2 - e\varphi(\mathbf{r},t) + V(\mathbf{r})\,. \tag{3.1}$$

[1]With charge $-e$, mass m and displacement from the nucleus $\mathbf{r}$.

This is the minimal coupling Hamiltonian for a charged particle interacting with an electromagnetic field, but including its intrinsic magnetic moment $\boldsymbol{\mu_S} = -g_S \mu_B \mathbf{S}/\hbar$ due to its spin $\mathbf{S} = \hbar\boldsymbol{\sigma}/2$, and including $V(\mathbf{r})$ the atomic binding potential[2]. The Bohr magneton is $\mu_B = e\hbar/2m$, we use the approximation $g_S \approx 2$, and $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ is a vector of the Pauli matrices. The potentials $\mathbf{A}$ and φ are related to the more familiar field quantities by the definitions

$$\begin{aligned} \mathbf{E}(\mathbf{r},t) &= -\boldsymbol{\nabla}\varphi(\mathbf{r},t) - \frac{\partial \mathbf{A}(\mathbf{r},t)}{\partial t}\,, \\ \mathbf{B}(\mathbf{r},t) &= \boldsymbol{\nabla} \times \mathbf{A}(\mathbf{r},t)\,. \end{aligned} \tag{3.2}$$

Expanding out the brackets in the Hamiltonian, using the Pauli identity $(\boldsymbol{\sigma}\cdot\mathbf{a})(\boldsymbol{\sigma}\cdot\mathbf{b}) = \mathbf{a}\cdot\mathbf{b} + i\boldsymbol{\sigma}\cdot(\mathbf{a}\times\mathbf{b})$, we can separate it into atomic and interaction terms[3]:

$$H = H_0 + H_I\,, \tag{3.3}$$

$$H_0 = \frac{\mathbf{p}^2}{2m} + V(\mathbf{r})\,, \tag{3.4}$$

$$H_I = \frac{e}{m}\mathbf{p}\cdot\mathbf{A}(\mathbf{r},t) - \boldsymbol{\mu_S}\cdot\mathbf{B}(\mathbf{r},t) + \frac{e^2}{2m}\mathbf{A}(\mathbf{r},t)^2 - e\varphi(\mathbf{r},t)\,. \tag{3.5}$$

We can Taylor expand the field quantities about the origin so that

$$\begin{aligned} \mathbf{E}(\mathbf{r},t) &= \mathbf{E}_0(t) + [\mathbf{r}\cdot\boldsymbol{\nabla}]\mathbf{E}_0(t) + \frac{1}{2!}[\mathbf{r}\cdot\boldsymbol{\nabla}]^2\mathbf{E}_0(t) + \cdots\,, \\ \mathbf{B}(\mathbf{r},t) &= \mathbf{B}_0(t) + [\mathbf{r}\cdot\boldsymbol{\nabla}]\mathbf{B}_0(t) + \frac{1}{2!}[\mathbf{r}\cdot\boldsymbol{\nabla}]^2\mathbf{B}_0(t) + \cdots\,. \end{aligned} \tag{3.6}$$

The gradient of the vector field $\boldsymbol{\nabla}\mathbf{E}_0 = (\boldsymbol{\nabla}\otimes\mathbf{E})\big|_0$ with the subscripts denoting evaluation at $\mathbf{r} = 0$ after any gradients have been taken. We can choose expressions for $\mathbf{A}$ and φ based on this expansion [49], which when inserted back into H_I yield its multipole

[2] Without the atomic binding potential this is the Hamiltonian of the Pauli equation.

[3] We have also used the fact that $[\mathbf{p}, \mathbf{A}] = 0$ in the Coulomb gauge ($\boldsymbol{\nabla}\cdot\mathbf{A} = 0$).

moment expansion

$$H_I = -\mathbf{d} \cdot \mathbf{E}_0 - \boldsymbol{\mu} \cdot \mathbf{B}_0 - \mathbf{Q} \,\colon\, \boldsymbol{\nabla} \mathbf{E}_0 + \cdots \,, \tag{3.7}$$

where we have dropped the time dependence for clarity and omitted the $e^2\mathbf{A}^2/2m$ diamagnetic term which is negligible for small fields[4]. $\mathbf{d} = -e\mathbf{r}$ is the electric dipole operator, $\boldsymbol{\mu} = \boldsymbol{\mu_L} + \boldsymbol{\mu_S}$ is the magnetic dipole operator comprised of both orbital and spin parts of the electron's magnetic moment, and $\mathbf{Q} = -e\mathbf{r} \otimes \mathbf{r}/2$ is the electric quadrupole operator. The orbital magnetic moment $\boldsymbol{\mu_L} = -\mu_B \mathbf{L}/\hbar$ arises from the electron's angular momentum $\mathbf{L} = \mathbf{r} \times \mathbf{p}$ about the nucleus.

Consider a plane wave with polarisation $\hat{\boldsymbol{\epsilon}}$, angular frequency ω, wavevector $\mathbf{k} = 2\pi\hat{\mathbf{k}}/\lambda$ and phase ϕ that interacts with our atom and takes the form

$$\begin{aligned} \mathbf{E} &= E\hat{\boldsymbol{\epsilon}} \cos(\mathbf{k} \cdot \mathbf{r} - \omega t + \phi) \,, \\ \mathbf{B} &= B\hat{\boldsymbol{\nu}} \cos(\mathbf{k} \cdot \mathbf{r} - \omega t + \phi) \,, \end{aligned} \tag{3.8}$$

where the ratio of the electric and magnetic field amplitudes $E/B = c$, and the polarisation vectors and wavevector form an orthonormal basis $\hat{\mathbf{k}} = \hat{\boldsymbol{\epsilon}} \times \hat{\boldsymbol{\nu}}$. As the wavelength is usually much larger than the physical extent of the electron wavefunction, the spatial variation of the field over the atom is negligible; i.e. $\mathbf{k} \cdot \mathbf{r} \approx \mathcal{O}(a_0/\lambda) \ll 1$. Thus gradients of the field and therefore higher order terms in H_I are small. Given also that magnetic interactions are suppressed by a factor of c compared to the same order electric interactions, we need only consider the lowest order moment with an allowed interaction (i.e. one where its term is non-vanishing) as it is vastly more likely than the rest.

The first term in the expansion is the electric dipole (E1) interaction. Where selection rules allow, this interaction is very often the dominant term and so we can make the

[4]The double dot product between two matrices is defined as $\mathbf{A} \,\colon\, \mathbf{B} = \sum_{ij} A_{ij} B_{ij}$, producing a scalar.

'dipole approximation' and neglect higher order terms. If a transition between a pair of states (e.g. $|g\rangle$,$|e\rangle$) is 'dipole forbidden' ($\langle e|-\mathbf{d}\cdot\mathbf{E}_0|g\rangle = 0$), we must consider higher order terms. The second term in H_I is the magnetic dipole (M1) interaction, and if this too is forbidden we consider the third term, the electric quadrupole (E2) interaction.

3.2 Single Qubit Rotations

Taking the lowest order allowed term in H_I for a particular pair of states, we can expand the interaction Hamiltonian in their basis[5] using $H_I = \mathbb{I}H_I\mathbb{I}$, with the identity $\mathbb{I} = |e\rangle\langle e| + |g\rangle\langle g| = \left(\begin{smallmatrix}1 & 0\\ 0 & 1\end{smallmatrix}\right)$ for excited and ground qubit states $|e\rangle = \left(\begin{smallmatrix}1\\ 0\end{smallmatrix}\right)$ and $|g\rangle = \left(\begin{smallmatrix}0\\ 1\end{smallmatrix}\right)$ respectively. As the electric dipole operator has negative parity (changes sign under the transformation $\mathbf{r} \to -\mathbf{r}$), the diagonal matrix elements $\langle e|\mathbf{r}\cdot\hat{\boldsymbol{\epsilon}}|e\rangle$ and $\langle g|\mathbf{r}\cdot\hat{\boldsymbol{\epsilon}}|g\rangle$ vanish. For the magnetic dipole and electric quadrupole interactions, the diagonal matrix elements lead to differential Zeeman and Stark shifts respectively, which we ignore here[6]. Now our Hamiltonian is

$$H_I = \hbar\Omega\cos(-\omega t + \phi)\sigma_x \tag{3.9}$$

with the interaction strength characterised by Ω the Rabi frequency, which for each term up to the electric quadrupole is

$$\begin{aligned}
\mathrm{E1}: \Omega &= \frac{eE}{\hbar}\langle e|\mathbf{r}\cdot\hat{\boldsymbol{\epsilon}}|g\rangle\ , && (3.10)\\
\mathrm{M1}: \Omega &= -\frac{B}{\hbar}\langle e|\boldsymbol{\mu}\cdot\hat{\boldsymbol{\nu}}|g\rangle\ , && (3.11)\\
\mathrm{E2}: \Omega &= \frac{eE}{2\hbar}\langle e|(\mathbf{r}\cdot\hat{\boldsymbol{\epsilon}})(\mathbf{k}\cdot\mathbf{r})|g\rangle\ , && (3.12)
\end{aligned}$$

[5]The similarity transformation $A \to UAU^\dagger$ gives operator A in the new basis if unitary operator U transforms from the old to the new basis.

[6]Electric and magnetic multipoles Eℓ and Mℓ have parity $(-1)^\ell$ and $(-1)^{\ell+1}$ respectively.

ignoring a $\pi/2$ offset to ϕ for the E2 Hamiltonian[7]. If we expand the atomic Hamiltonian (equation 3.4 plus any fine or hyperfine structure terms) in the same basis it becomes

$$H_0 = \frac{\hbar\omega_0}{2}\sigma_z\,. \tag{3.13}$$

To make the effect of the interaction H_I clearer, it is helpful to move into a frame rotating at the qubit frequency ω_0. We can do this by shifting to the 'Interaction picture' with respect to H_0, using the unitary operator $U = e^{iH_0t/\hbar}$ so that $H_I \to UH_IU^\dagger$ and becomes[8]

$$H_I = \frac{\hbar\Omega}{2}\left(e^{-i\delta t}e^{i\phi}\sigma^+ + e^{i\delta t}e^{-i\phi}\sigma^-\right)\,, \tag{3.14}$$

and adding a term $-\hbar\omega_0\sigma_z/2$ to the total Hamiltonian so that $H = H_I$. We have also kept terms rotating at the detuning $\delta = \omega - \omega_0$ but dropped those rotating at $\omega + \omega_0$, as assuming we are close to the atomic resonance ($\delta \ll \omega_0$) these terms rotate very fast and will quickly average to zero; this is the rotating wave approximation. When the light is on resonance ($\delta = 0$) we find that

$$H_I = \frac{\hbar\Omega}{2}\left[\cos(\phi)\sigma_x - \sin(\phi)\sigma_y\right]\,, \tag{3.15}$$

which as it is time-independent has a time evolution operator $U = e^{-iH_It/\hbar}$ which is[9]

$$U = R(\Omega t, \phi) = \cos\left(\frac{\Omega t}{2}\right)\mathbb{I} - i\sin\left(\frac{\Omega t}{2}\right)\left[\cos(\phi)\sigma_x - \sin(\phi)\sigma_y\right]\,. \tag{3.16}$$

This describes Rabi oscillations (or Rabi 'flopping') about an axis perpendicular to z-axis of the Bloch sphere. The state/Bloch vector rotates from $|g\rangle$ to $|e\rangle$ and back at the Rabi frequency Ω, and by choosing a pulse time such that $\Omega t = \pi$ we can do a $\pi-$pulse

[7]This is due to the gradient we took, but as we care about relative rather than absolute phase this is unimportant.

[8]Note this is the inverse of the time evolution operator for H_0.

[9]Using the identity $e^{i\theta(\hat{\mathbf{n}}\cdot\boldsymbol{\sigma})} = \cos(\theta)\mathbb{I} + i\sin(\theta)(\hat{\mathbf{n}}\cdot\boldsymbol{\sigma})$ where $\hat{\mathbf{n}}$ is a unit vector in the basis of the Pauli matrices.

so that $R(\pi, \phi)|g\rangle \to |e\rangle$, flipping the state of the qubit up to a global phase. We can do σ_x rotations with $R_X(\theta) = R(\Omega t, 0)$ and σ_y rotations with $R_Y(\theta) = R(\Omega t, -\pi/2)$, and by adding a constant offset to ϕ for all subsequent pulses we can also implement z-axis rotations, as this phase offset between our local oscillator and the qubit simply rotates the x- and y-axes of the Bloch sphere around the z-axis.

3.2.1 Off-Resonant Effects

To find the general time evolution of H_I for any detuning δ, it is helpful to remove the time dependence by moving into the interaction picture frame rotating at the laser frequency ω. From equation 3.14, we can boost the frame by the detuning frequency and use the same similarity transformation as before but with $U = e^{i\delta t\sigma_z/2}$ so the total Hamiltonian is now

$$H = -\frac{\hbar\delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\left[\cos(\phi)\sigma_x - \sin(\phi)\sigma_y\right] . \tag{3.17}$$

After some algebra one finds that the time evolution operator is

$$U = \begin{pmatrix} \cos\left(\frac{\Omega' t}{2}\right) + i\frac{\delta}{\Omega'}\sin\left(\frac{\Omega' t}{2}\right) & -ie^{i\phi}\frac{\Omega}{\Omega'}\sin\left(\frac{\Omega' t}{2}\right) \\ -ie^{-i\phi}\frac{\Omega}{\Omega'}\sin\left(\frac{\Omega' t}{2}\right) & \cos\left(\frac{\Omega' t}{2}\right) - i\frac{\delta}{\Omega'}\sin\left(\frac{\Omega' t}{2}\right) \end{pmatrix} , \tag{3.18}$$

where the Bloch vector rotates about an axis somewhere between the z-axis and equator of the Bloch sphere at the effective Rabi frequency $\Omega' = \sqrt{\Omega^2 + \delta^2}$.

In the limit where $|\Omega| \gg |\delta|$ we recover equation 3.16. In the off-resonant case when $|\delta| \gg |\Omega|$ however, if we expand $\Omega' = \delta + \Omega^2/2\delta \cdots$ we can see that the excitation probability $|\langle e|U|g\rangle|^2$ has amplitude $\approx \Omega^2/\delta^2$. Dropping these small off-diagonal terms we can also see that the time evolution operator looks like a σ_z rotation at frequency Ω', which is an energy shift to our qubit states. Working backwards the Hamiltonian can be written as

$$H \approx -\frac{\hbar\delta}{2}\sigma_z - \frac{\hbar\Omega^2}{4\delta}\sigma_z\,, \tag{3.19}$$

where the first term is just the atomic splitting[10], with the second term an A.C. Stark shift or light-shift induced by our beam.

3.3 Raman Interactions

We sometimes wish to drive radio/microwave frequency magnetic dipole qubits with optical frequency radiation, for reasons that will become clear in section 3.4. We can achieve this by coupling the qubit states $|g\rangle$ and $|e\rangle$ to an upper level $|u\rangle$, using two laser beams on electric dipole transitions with detunings, Rabi frequencies and phases δ_1,Ω_1,ϕ_1 and δ_2,Ω_2,ϕ_2 respectively.

In the interaction picture after making the rotating wave approximation the interaction Hamiltonian is the sum of that for each beam:

$$H_I = \frac{\hbar\Omega_1}{2}\left(e^{-i\delta_1 t}e^{i\phi_1}\,|u\rangle\langle g| + e^{i\delta_1 t}e^{-i\phi_1}\,|g\rangle\langle u|\right) \\ + \frac{\hbar\Omega_2}{2}\left(e^{-i\delta_2 t}e^{i\phi_2}\,|u\rangle\langle e| + e^{i\delta_2 t}e^{-i\phi_2}\,|e\rangle\langle u|\right)\,. \tag{3.20}$$

We are interested in the systems behaviour when $|\delta_1| \gg |\Omega_1|$ and $|\delta_2| \gg |\Omega_2|$. By choosing an appropriate frame one can find the time evolution operator and adiabatically eliminate the population in $|u\rangle$, but it is easier to find the time averaged dynamics using the James-Jerke approximation [50]. Here we terminate the series expansion of the time evolution operator after two terms, neglecting any terms that oscillate at the

[10] Remember we are in a frame rotating at the laser frequency ω, so the qubit rotates at a frequency δ relative to this.

detunings or faster leaving the effective Hamiltonian as

$$H_I \approx -\frac{\hbar\Omega_1^2}{4\delta_1}\left(|u\rangle\langle u| - |g\rangle\langle g|\right) - \frac{\hbar\Omega_2^2}{4\delta_2}\left(|u\rangle\langle u| - |e\rangle\langle e|\right) + \frac{\hbar\Omega_1\Omega_2}{4\Delta}\left(e^{-i\Delta\delta t}e^{i\Delta\phi}|e\rangle\langle g| + e^{i\Delta\delta t}e^{-i\Delta\phi}|g\rangle\langle e|\right). \quad (3.21)$$

The first two terms are single beam light-shifts from each laser, akin to the second term of equation 3.19, which if not balanced produce a differential light-shift on the qubit[11]. The last term is a two-photon stimulated Raman interaction which looks like equation 3.14, and describes Rabi oscillations of the qubit at effective Rabi frequency $\Omega_1\Omega_2/2\Delta$, where the Raman detuning Δ is defined as $1/\Delta = (1/\delta_1 + 1/\delta_2)/2$, difference detuning $\Delta\delta = \delta_1 - \delta_2$ and relative phase $\Delta\phi = \phi_1 - \phi_2$. We can therefore carry out single qubit operations on ground level qubits using a pair of Raman beams detuned from an optical transition to an upper level.

Often there are multiple upper levels through which we can couple, for example due to a fine structure splitting. As long as $|\delta| \gg |\Omega|$ holds true for each level, we can coherently sum the Rabi frequencies from each path such that

$$\Omega_{\text{Raman}} = \sum_i \frac{\Omega_{g,u_i}\Omega_{e,u_i}}{2\Delta_i}\,, \quad (3.22)$$

where Δ_i is the Raman detuning from, and Ω_{g,u_i} & Ω_{e,u_i} are the Rabi frequencies from the qubit states to, each upper state $|u_i\rangle$.

3.3.1 Scattering

While it may seem advantageous to reduce Δ to increase the effective Rabi frequency of the Raman interaction, this leads to an increase in the otherwise small population in $|u\rangle$ which, as the upper states have a finite lifetime, can spontaneously decay into

[11] Indeed the James-Jerke approximation could have been used in the derivation of the light-shift in the previous section.

an arbitrary final state and scramble the qubit creating an error. This can analogously be seen as a Raman interaction between the initial and final states where one beam is replaced by vacuum fluctuations of the electromagnetic field [51]. If the final state is the same as the initial state this is known as elastic Rayleigh scattering, otherwise it is known as inelastic Raman scattering.

For a single beam with field E, the scattering rate from $|i\rangle$ to $|f\rangle$ via upper state $|u_j\rangle$ is given by the Kramers-Heisenberg formula

$$\Gamma_{i\to f} = \frac{E^2\gamma}{4\hbar^2}\sum_k \left|\sum_{j,q} \frac{\langle f|\mathbf{d}\cdot\hat{\boldsymbol{\epsilon}}_q|u_j\rangle\ \langle u_j|\mathbf{d}\cdot\hat{\boldsymbol{\epsilon}}_k b_k|i\rangle}{\Delta_j}\right|^2, \tag{3.23}$$

where γ is the mean natural linewidth of the upper states, and the laser polarisation $\hat{\boldsymbol{\epsilon}} = \sum_k \hat{\boldsymbol{\epsilon}}_k b_k$ and emitted photon polarisation $\hat{\boldsymbol{\epsilon}}_q$ are given in the basis $\hat{\boldsymbol{\epsilon}}_p = \{\hat{\boldsymbol{\sigma}}^+, \hat{\boldsymbol{\sigma}}^-, \hat{\boldsymbol{\pi}}\}$ [51]. The total scattering rate out of $|i\rangle$ is then $\Gamma_i = \Gamma_{\text{Rayleigh}} + \Gamma_{\text{Raman}}$, with $\Gamma_{\text{Rayleigh}} = \Gamma_{i\to i}$ and $\Gamma_{\text{Raman}} = \Gamma_{i\to f}$ for $i \neq f$.

Expressions for the Raman Rabi frequencies, light-shifts and scattering rates for $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ and similar species can be found in [32, 39, 52].

3.4 Coupling to the Ion Motion

In the preceding sections we have not considered the interaction between the applied light and the ions motion in the trap. In the multipole expansion of the interaction Hamiltonian (section 3.1) we argued that as the spatial variation of the light is negligible over the extent of the electron wavefunction $\mathbf{k}\cdot\mathbf{r} \ll 1$, we can neglect orders in the expansion higher than the first allowed term. This is not necessarily the case when considering the variation over the entire atoms wavefunction.

Considering only the axial motion, an ion of mass m is in a harmonic potential $m\omega_z^2 z^2/2$ and has ground state wavefunction spread $\tilde{z} = \sqrt{\hbar/2m\omega_z}$, which for $^{43}\text{Ca}^+$ in a typical

1.9 MHz trap is ~8 nm. For microwave/RF radiation with wavelengths of a few cm or longer the spatial variation over the atoms wavefunction $\mathbf{k} \cdot \mathbf{r_a}$ is also negligible[12], but this is not true for optical radiation with wavelengths of a few hundred nm. We now need to include the spatial variation of the field over the atoms co-ordinates in the interaction Hamiltonian of equation 3.9, so that

$$H_I = \hbar\Omega \cos(\mathbf{k} \cdot \mathbf{r_a} - \omega t + \phi)\sigma_x \,. \tag{3.24}$$

For a pair of Raman beams we replace the motional phase with $\Delta\mathbf{k} \cdot \mathbf{r_a}$, using the difference wavevector $\Delta\mathbf{k} = \mathbf{k_1} - \mathbf{k_2}$. As the magnitudes of the wavevectors for each beam are almost identical $k = |\mathbf{k_1}| = |\mathbf{k_2}|$, we can write $|\Delta\mathbf{k}| = 2k\sin(\theta/2)$ where θ is the angle between the beams. In the following we can always replace $\mathbf{k} \to \Delta\mathbf{k}$ to get the result for a pair of Raman beams.

3.4.1 Quantisation of Motion

To analyse the interaction of light with the ions motion we first need to quantise the motion. Considering the one dimensional case for N ions of differing species in a line along the z-axis, we can write the Hamiltonian for the atoms in the harmonic potential as

$$H_m = \sum_{j=1}^{N} \hbar\omega_j \left(a_j^\dagger a_j + \frac{1}{2} \right) , \tag{3.25}$$

where the raising and lowering operators $a_j^\dagger$ and a_j add and subtract a quanta of motion (a phonon) from normal mode j with frequency ω_j respectively. The displacement operator for ion i from its equilibrium position, in the direction of unit vector $\hat{\mathbf{z}}$, is related to the sum of the raising and lowering operators for each of the N modes [41]:

$$\mathbf{r}_i = \mathbf{z}_i = \hat{\mathbf{z}} \sum_{j=1}^{N} \tilde{z}_{ij} \left(a_j + a_j^\dagger \right) . \tag{3.26}$$

[12] If the intensity of the radiation is made to vary over short length scales it can become significant, e.g. near-field microwaves.

The ground state wavefunction spread $\tilde{z}_{ij}$ now depends on b_{ij}, the component of the j^{th} normal mode eigenvector for the i^{th} ion[13] (as described in section 2.4.1), so

$$\tilde{z}_{ij} = b_{ij}\sqrt{\frac{\hbar}{2m_i\omega_j}} \,. \tag{3.27}$$

The motional phase is now

$$\mathbf{k}\cdot\mathbf{r}_i = \sum_{j=1}^{N}\eta_{ij}\left(a_j + a_j^\dagger\right) \,, \tag{3.28}$$

where we have defined the Lamb-Dicke parameter as

$$\eta_{ij} = \tilde{z}_{ij}\mathbf{k}\cdot\hat{\mathbf{z}} \,, \tag{3.29}$$

which quantifies the coupling strength of the light to the ion's motion. The Lamb-Dicke parameter is proportional to the ratio of the motional ground state wavefunction spread to the wavelength, and its square is the ratio of the photon recoil energy transferred into mode j, $(\hbar k_z b_{ij})^2/2m_i$, to the energy spacing of that mode, $\hbar\omega_j$. The Lamb-Dicke parameter depends on the beam geometry, where $\mathbf{k}\cdot\hat{\mathbf{z}}$ quantifies the projection of the wavevector onto the mode direction. We can easily generalise the above to all three dimensions and quantise all $3N$ modes, but the results for the radial modes are analogous to the axial modes.

We wish to consider the interaction Hamiltonian for a single beam (or single pair of Raman beams) equally illuminating the whole ion string, tuned near to the qubit frequency ω_0 of a particular species. As the qubit frequencies of different species are usually very different, true in this thesis, the beam is so far detuned from the other species as to be impotent and their terms can be dropped from the Hamiltonian. In the interaction picture with respect to $H_0 + H_m$, where $H_0 = \sum_{i=1}^{N}\hbar\omega_{0,i}\sigma_{z,i}/2$ is the sum of

[13] For a single ion $b = 1$.

the individual atomic Hamiltonians, and after taking the rotating wave approximation, it is therefore

$$H_I = \sum_{i=1}^{N} \frac{\hbar\Omega}{2} \Big(e^{i \sum_{j=1}^{N} \eta_{ij} \left(a_j e^{-i\omega_j t} + a_j^\dagger e^{i\omega_j t}\right)} e^{-i\delta t} e^{i\phi} \sigma_i^+ \\ + e^{-i \sum_{j=1}^{N} \eta_{ij} \left(a_j e^{-i\omega_j t} + a_j^\dagger e^{i\omega_j t}\right)} e^{i\delta t} e^{-i\phi} \sigma_i^- \Big) \delta_{\omega_0 - \omega_{0,i}} \,, \quad (3.30)$$

where $\delta_{\omega_0 - \omega_{0,i}}$ is the Kronecker delta function selecting ions of the same species.

3.4.2 Lamb-Dicke Regime

In order see the effect of the motional phase factor it is useful to Taylor expand it. Looking at a particular ion in a particular mode it is

$$e^{i\mathbf{k}\cdot\mathbf{r_a}} = e^{i\eta\left(a+a^\dagger\right)} = 1 + i\eta\left(a + a^\dagger\right) - \frac{\eta^2\left(a+a^\dagger\right)^2}{2} + \cdots \,. \quad (3.31)$$

As the Lamb-Dicke parameter η is small, we often choose to neglect the $\mathcal{O}(\eta^2)$ and higher terms and work in the Lamb-Dicke regime, where the amplitude of the ions motion in the beam direction is small compared to $\lambda/2\pi$; i.e. $\langle\Psi_{\text{motion}}|(\mathbf{k}\cdot\mathbf{r_a})^2|\Psi_{\text{motion}}\rangle^{1/2} = \eta\sqrt{2n+1} \ll 1$, known as the Lamb-Dicke criterion. For hot ions with a beam strongly coupled to the motion this criterion is not necessarily satisfied.

For a single ion in the Lamb-Dicke regime equation 3.30 simplifies to

$$H_I = \frac{\hbar\Omega}{2}\left(\left[1 + i\eta\left(ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t}\right)\right] e^{-i\delta t} e^{i\phi} \sigma^+ + \left[1 - i\eta\left(ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t}\right)\right] e^{i\delta t} e^{-i\phi} \sigma^-\right) \\ = H_{\text{Carrier}} + H_{\text{RSB}} + H_{\text{BSB}} \,, \quad (3.32)$$

where we can identify three separate interactions.

$$H_{\text{Carrier}} = \frac{\hbar\Omega}{2}\left(e^{-i\delta t}e^{i\phi}\sigma^+ + e^{i\delta t}e^{-i\phi}\sigma^-\right) \tag{3.33}$$

is the carrier interaction which does not change the motional state, and has a Hamiltonian identical to equation 3.14. It causes Rabi oscillations as before between $|g,n\rangle \leftrightarrow |e,n\rangle$.

The second is the red sideband interaction

$$\begin{aligned} H_{\text{RSB}} &= \frac{i\hbar\Omega\eta}{2}\left(ae^{-i\omega_z t}e^{-i\delta t}e^{i\phi}\sigma^+ - a^\dagger e^{i\omega_z t}e^{i\delta t}e^{-i\phi}\sigma^-\right) \\ &= \sum_n \frac{i\hbar\Omega\eta\sqrt{n}}{2}\left(|n-1\rangle\langle n|\, e^{-i\delta_r t}e^{i\phi}\sigma^+ - |n\rangle\langle n-1|\, e^{i\delta_r t}e^{-i\phi}\sigma^-\right)\,, \end{aligned} \tag{3.34}$$

which is analogous to the Jaynes-Cummings Hamiltonian from cavity QED [53, 54], where we picture a ladder of states of increasing n, with states $|g,n\rangle \leftrightarrow |e,n-1\rangle$ connected by subtracting a phonon. By expanding in the basis of motional states $\mathbb{I} = \sum_n |n\rangle\langle n|$, using $a|n\rangle = \sqrt{n}\,|n-1\rangle$, and defining the detuning from the red sideband $\delta_r = \delta + \omega_z$, we can see that this describes Rabi oscillations between the aforementioned states with effective Rabi frequency $\Omega_{\text{RSB}} = i\Omega\eta\sqrt{n}$. The interaction is on resonance when the laser is tuned to the first red sideband of the carrier; $\delta = -\omega_z$.

Finally we have the blue sideband interaction

$$\begin{aligned} H_{\text{BSB}} &= \frac{i\hbar\Omega\eta}{2}\left(a^\dagger e^{i\omega_z t}e^{-i\delta t}e^{i\phi}\sigma^+ - ae^{-i\omega_z t}e^{i\delta t}e^{-i\phi}\sigma^-\right) \\ &= \sum_n \frac{i\hbar\Omega\eta\sqrt{n+1}}{2}\left(|n+1\rangle\langle n|\, e^{-i\delta_b t}e^{i\phi}\sigma^+ - |n\rangle\langle n+1|\, e^{i\delta_b t}e^{-i\phi}\sigma^-\right)\,, \end{aligned} \tag{3.35}$$

where we have defined $\delta_b = \delta - \omega_z$ and used $a^\dagger|n\rangle = \sqrt{n+1}\,|n+1\rangle$. This describes Rabi oscillations which add a phonon between $|g,n\rangle \leftrightarrow |e,n+1\rangle$ with effective Rabi frequency $\Omega_{\text{BSB}} = i\Omega\eta\sqrt{n+1}$, on resonance when the laser is tuned to the first blue sideband of the carrier; $\delta = +\omega_z$.

Scanning the light frequency these three interactions produce a spectrum of a central peak with smaller sidebands at $\pm\omega_z$. As long as $\omega_z \gg \Omega$ these resonances are well resolved and we can tune our beam to add to, subtract from, or ignore the phonons in the mode while flipping $|g\rangle \leftrightarrow |e\rangle$.

We still need to be mindful of off-resonant excitation however, especially of the carrier when tuned to a sideband as its strength is enhanced by a factor $1/\eta$ with respect to the sideband. When scanning with higher intensity this produces a light-shift which reduces the carrier-sideband spacing below ω_z.

Outside the Lamb-Dicke Regime

With increasing ion temperature higher order sidebands of diminishing amplitude at multiples of ω_z appear, which add or subtract multiple phonons from the mode. We also see a suppression in the strength of the carrier, with the effective Rabi frequency at the next order of approximation $\Omega_{\text{Carrier}} = \Omega(1 - \eta^2 n)$. If not in a pure motional eigenstate (i.e. in a superposition of multiple $|n\rangle$ states), the differing carrier Rabi frequencies of each state increasingly dephase the longer we apply our light pulse, decohering our qubit. Even in fairly cold (low $\overline{n}$) thermal states, this suppresses the amplitude of the carrier population inversion after several Rabi cycles. This problem can be avoided by using radiation with a very small Lamb-Dicke parameter (e.g. RF/microwave) or by using Raman beams in a co-propagating geometry where $\Delta\mathbf{k} = 0$ giving only a motionally insensitive interaction.

The full expression for the Rabi frequency for an arbitrary Δn transition is somewhat complex and can be found in [21], but it shows that with increasing n the carrier and then the sidebands are eventually suppressed to zero before reviving. This can be understood as a decrease in the light amplitude as seen by the ion, as it is averaged over the ion's growing motional wavepacket. Eventually the $\mathrm{n^{th}}$ order gradient of the light when averaged over the ion reduces also, suppressing the $\mathrm{n^{th}}$ order sideband.

We see similar sideband resonances for each mode the beam has a non-zero projection onto, at multiples of the mode frequency. We also see mixing products between modes where one or multiple phonons are added or subtracted from different modes. For example for a single ion we would expect to see mixing products at detunings $\delta = \pm p\omega_z \pm q\omega_r \pm q'\omega_{r'}$, where $\{p, q, q'\} \in \mathbb{N}_0$, with Rabi frequencies suppressed by a factor $(\eta_z)^p(\eta_r)^q(\eta_{r'})^{q'}$ with respect to the carrier.

3.5 Two-Qubit Gates

To create entanglement between the spins of two ions in a trap, one might first suppose we could use their direct interaction, like neighbouring bar magnets. This interaction however falls off sharply with distance, and with atomic separations four orders of magnitude further apart than in solids, it becomes much too weak to be useful [38]. Instead we turn to the strong Coulomb interaction between the ions, and the shared motion it creates, which can be used as a bus through which we can mediate an interaction between the spins. The key step in achieving this is finding a method to link the ions internal spin degree of freedom with its external motional degree of freedom.

The first proposal implementing this idea came from Cirac and Zoller in their 1995 paper laying out the foundations of the trapped ion quantum information processor [12]. Their scheme involves a series of sideband pulses focussed on each ion, which conditionally excites the motion depending on the spin state of the 'control' ion. If the motion is excited, the 'target' ion picks up a phase, with the motion then de-excited at the end of the process. While this scheme implements a two-qubit gate, it requires the ions to begin and end in the motional ground state, a significant limitation.

Since then a number of gate mechanisms have been proposed, which require only that the ions are in the Lamb-Dicke regime. The most popular of these share a common feature; a spin dependent force that excites then de-excites the motion, with the spin states acquiring a geometric phase in the process [55–57].

3.5.1 The Geometric Phase Gate

The basic principle of a geometric phase gate is as follows. Consider driving the harmonic motion of one or more ions, which start off near still, with a sinusoidal force that is slightly off-resonant by detuning δ from the natural oscillation frequency of a particular motional mode. Just as with a classical oscillator, the motion starts out in phase with this force and the amplitude of the oscillations grow. Over time t a phase difference δt develops between the motion and the force, which eventually exceeds π and causes the force to 'reverse' and damp the motion. When $\delta t = 2\pi$ the force has removed the excitation it added, and the ions return to their initial state.

In a position-momentum phase space this plots a trajectory that spirals out from near the origin and then spirals back in again, as the amplitude of the oscillations grow and fall. In a reference frame rotating at the motional mode frequency, this draws a circle starting and ending near the origin, where a resonant force would push the motional state out in a straight line.

This loop drawn in phase space is consequential, as the ions acquires a path dependent phase from it. This phase has two components; the dynamic phase acquired by spending time at a higher energy state, and a geometric 'Berry' phase equal to the phase space area enclosed by the loop [58,59]. These partially cancel, since the dynamic phase is -2 times the geometric phase, so the total phase is proportional to the geometric phase. A stronger driving force moves the ions through a larger loop in phase space, acquiring a larger geometric phase.

Ordinarily, this just applies a global phase to the ions' wavefunction, which does nothing. If however the force depends on the ions' spin, and acts with different strength on different spin states, each spin state will traverse a loop of different size and acquire a different geometric phase. If the ions' start out in a superposition of multiple spin states, these acquire phases relative to each other, making a material change to the overall wavefunction.

For a single ion, relative phase changes in a superposition are nothing more than single qubit operations, which can be done much more easily by other means. For multiple ions however, this process can transform separable input states to maximally entangled output states.

3.5.2 Spin Dependent Force Hamiltonian

To understand the action of the gate, we can construct the Hamiltonian for this off-resonant force on the ions [32,60–62]. In the interaction picture with respect to H_0+H_m, a spatially uniform force $F(t) = F_0\cos(\omega t - \phi)$ acting on a single motional mode gives

$$\begin{aligned} H_I &= zF_0\cos(\omega t - \phi) \\ &= \tilde{z}F_0\left(ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t}\right)\cos(\omega t - \phi) \\ &= \frac{\tilde{z}F_0}{2}\left(ae^{i\delta t}e^{-i\phi} + a^\dagger e^{-i\delta t}e^{i\phi}\right) , \end{aligned} \tag{3.36}$$

where $\tilde{z}$ is the ground state wavefunction spread, as in equation 3.26, and in the last line we have taken the rotating wave approximation with the motional frequencies, keeping terms rotating at the force detuning $\delta = \omega - \omega_z$.

To encapsulate the spin dependence of the force, we can describe its effect with a Rabi frequency for each orthogonal spin state $|s\rangle$, so $\tilde{z}F_0 = \hbar\sum\Omega_s|s\rangle\langle s|$. Alternatively, we can use a single Rabi frequency acting on a general normalised[14] spin operator Λ, so $\tilde{z}F_0 = \hbar\Omega_F\Lambda$. Making this latter choice, the interaction Hamiltonian is now

$$H_I = \frac{\hbar\Omega_F}{2}\Lambda\left(a^\dagger e^{-i\delta t}e^{i\phi} + ae^{i\delta t}e^{-i\phi}\right) . \tag{3.37}$$

[14]The Hilbert-Schmidt norm $\mathrm{tr}\left[\Lambda\Lambda^\dagger\right] = 2^N$, for N qubits.

Unfortunately, we cannot simply find the time evolution operator of this Hamiltonian with $U = e^{-\frac{i}{\hbar}\int_0^t H_I(t')\mathrm{d}t'}$, as H_I does not commute with itself at different times. The general Magnus expansion expression for the propagator in this case is [63, 64]

$$
\begin{aligned}
U = \exp\Big\{ & -\frac{i}{\hbar}\int_0^t H_I(t')\,\mathrm{d}t' \\
& -\frac{1}{2\hbar^2}\int_0^t \mathrm{d}t' \int_0^{t'} [H_I(t'), H_I(t'')]\,\mathrm{d}t'' \\
& +\frac{i}{6\hbar^3}\int_0^t \mathrm{d}t' \int_0^{t'} \mathrm{d}t'' \int_0^{t''} \left([H_I(t'),[H_I(t''),H_I(t''')]] + [H_I(t'''),[H_I(t''),H_I(t')]]\right)\mathrm{d}t''' \\
& + \cdots \Big\}.
\end{aligned}
\tag{3.38}
$$

Thankfully, the commutator in the second term in the exponent evaluates to a scalar, so higher terms vanish. We then find the propagator of our spin dependent force to be

$$U(t) = D\left(\alpha(t)\Lambda\right) e^{-i\Phi(t)\Lambda^2} \tag{3.39}$$

$$\alpha(t) = \frac{\Omega_F}{2\delta} e^{i\phi}\left(e^{-i\delta t} - 1\right) \tag{3.40}$$

$$\Phi(t) = \frac{\Omega_F^2}{4\delta^2}\left(\delta t - \sin(\delta t)\right). \tag{3.41}$$

The propagator consists of two parts, corresponding to the first two terms in the Magnus expansion; a spin dependent displacement operator $D(\alpha') = e^{\alpha' a^\dagger - \alpha'^* a}$, which moves a state through a trajectory $\alpha(t)$ in phase space, and a geometric phase which corresponds to the area enclosed.

Phase space is now the complex plane, with $\alpha(t)$ drawing a circle with radius $\Omega_F/2\delta$. The initial motional state need not lie at the phase space origin, and the initial direction of the path is set by the phase ϕ, with the sense of the rotation (i.e. clockwise or anti-clockwise) by the sign of the detuning. The origin of the geometric phase can be seen in the non-commutativity of the displacement operator, which when acting sequentially on a state produces an extra phase factor, $D(\alpha)D(\beta) = D(\alpha+\beta)e^{i\,\mathrm{Im}(\alpha\beta^*)}$. Another expression for the geometric phase is therefore $\Phi(t) = -\operatorname{Im}\left(\int_0^t \alpha^*(t')\,\mathrm{d}\alpha(t')\right)$.

The spin dependence can be understood by mapping the action of the propagator on the 2^N eigenstates $\{|\lambda\rangle\}$ of the spin operator Λ, where $\Lambda|\lambda\rangle = \lambda|\lambda\rangle$ and N is the number of qubits. Now

$$U(t) = \sum_{\lambda}^{2^N} D\left(\alpha(t)\lambda\right) e^{-i\Phi(t)\lambda^2} |\lambda\rangle\langle\lambda| \; . \tag{3.42}$$

We can therefore see that the spin dependent force pushes each of the eigenstates around in phase space, with the size and initial direction of the path given by the eigenvalue λ's magnitude and sign. Consequently, the area this encloses is proportional to λ^2.

If the ion starts out in a superposition of these eigenstates, the spin-dependent force therefore entangles the spin state with the motional state. After a time $t = 2\pi/\delta$ has elapsed however, each spin state has completed a loop in phases space, so $\alpha(t) = 0$ and the two disentangle. This process can continue, making multiple loops in phase space, which close at times $t_g = 2\pi K/\delta$ with K the number of loops. The propagator is now

$$U(t_g) = \sum_{\lambda}^{2^N} e^{-i\Phi(t_g)\lambda^2} |\lambda\rangle\langle\lambda| \tag{3.43}$$

$$\Phi(t_g) = \frac{\pi K \Omega_F^2}{2\delta^2} \tag{3.44}$$

To see how this can generate entanglement, consider the spin operator[15] $\Sigma_z = (\sigma_z \otimes I + I \otimes \sigma_z)/\sqrt{2}$, in place of Λ. As $\Sigma_z = \operatorname{diag}\left(\sqrt{2}, 0, 0, -\sqrt{2}\right)$, its eigenstates are the four spin states of the usual Z-basis, $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$, with the diagonal elements their eigenvalues. Since the eigenvalues of the odd spin states are zero, the force magnitudes on them are zero, so they do not move in phase space. The odd states make loops in opposite directions, and after a time t_g the propagator becomes

[15] In a Hamiltonian $H_I = \hbar\Omega\Sigma_z/\sqrt{2}$, this produces when $\Omega t = \pi$ the propagator $U = -i\sigma_z \otimes \sigma_z$, a global σ_z operator on both qubits.

$U(t_g) = \mathrm{diag}\left(e^{-i2\Phi(t_g)}, 1, 1, e^{-i2\Phi(t_g)}\right)$. Choosing the force strength and detuning such that $2\Phi(t_g) = \pi/2$, we implement up to a global phase the gate operation $U(t_g) = -iG$

$$G = \begin{pmatrix} 1 & & & \\ & i & & \\ & & i & \\ & & & 1 \end{pmatrix}. \tag{3.45}$$

Applying this to a (separable) equal superposition of the Z-basis states, we get the maximally entangled output $|\Psi\rangle = (|\uparrow\uparrow\rangle + i\,|\uparrow\downarrow\rangle + i\,|\downarrow\uparrow\rangle + |\downarrow\downarrow\rangle)/2$.

The question now arises, how do we generate a spin dependent force on the ions in practice? Firstly, we need control fields which interact with the ions' motion as described in section 3.4, which will have their Rabi frequencies suppressed by a factor η. We then need to create a Hamiltonian with the same functional form as equation 3.37, with the same spin operator acting on both a and $a^\dagger$.

3.5.3 Mølmer Sørensen Force

Looking at the sideband interaction Hamiltonians of section 3.4.2, one might suppose some combination of sidebands could do the job. Indeed this is the idea behind the Mølmer Sørensen interaction; two equal strength beams arranged symmetrically around the carrier, detuned from the red and blue sidebands respectively [55, 65]. This gives the spin dependent force Hamiltonian with spin operator $\Sigma_\phi = (\sigma_\phi \otimes I + I \otimes \sigma_\phi)/\sqrt{2}$, where $\sigma_\phi = \cos(\phi)\sigma_x - \sin(\phi)\sigma_y$. Its eigenstates are the four spin states of the ϕ-basis $\{|\rightarrow\rightarrow\rangle, |\rightarrow\leftarrow\rangle, |\leftarrow\rightarrow\rangle, |\leftarrow\leftarrow\rangle\}$, with eigenvalues as before. Described in this basis, the gate unitary matrix is G, but in the usual Z-basis for $\phi = 0$ it is, up to a global phase,

$$G_{\mathrm{MS}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & & & -i \\ & 1 & -i & \\ & -i & 1 & \\ -i & & & 1 \end{pmatrix}. \tag{3.46}$$

Conveniently, our usual starting state $|\downarrow\downarrow\rangle$ is an equal superposition of these basis states, so applying the gate and writing the output state in the Z-basis, we directly get $G_{\mathrm{MS}} |\downarrow\downarrow\rangle = \frac{1}{\sqrt{2}} (|\uparrow\uparrow\rangle - i\,|\downarrow\downarrow\rangle)$, a Bell state.

3.5.4 Light-Shift Force

Another route to creating a spin-dependent force is with a pair of Raman beams (section 3.3) with their polarisations arranged so they apply different light-shifts to $|\uparrow\rangle$ and $|\downarrow\rangle$, but don't drive spin flips between them. The differential light-shift between them is important, as it changes the qubit's energy giving a σ_z term to the Hamiltonian. When at the same frequency, the overlapping beams create a standing wave interference pattern, where the magnitude of the light-shift varies over the bright and dark regions. This spatially varying energy is of-course a force.

If we set a frequency difference between the beams, this standing wave pattern starts to move at this frequency, applying a periodic force to the ions. Setting the difference frequency of this 'walking wave' slightly detuned from a motional mode frequency then implements the off-resonant spin dependent force we need.

This is the mechanism behind the light-shift, or 'wobble', gate [32, 39, 56]. We get the Hamiltonian of equation 3.37, with spin operator Σ_z as previously described. We need to map into a superposition of the Z-basis states before the gate with a global $\pi/2$

rotation, and undo it after. The gate unitary is therefore

$$G_{\mathrm{LS}} = R_{XX}\left(\frac{\pi}{2}\right) \cdot G \cdot R_{XX}\left(-\frac{\pi}{2}\right) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & & & i \\ & 1 & -i & \\ & -i & 1 & \\ i & & & 1 \end{pmatrix}, \tag{3.47}$$

with $R_{XX}(\theta) = R_X(\theta) \otimes R_X(\theta)$.

A limitation of this mechanism is that, at the large Raman detunings needed to avoid photon scattering, the beams can only produce differential light-shifts on qubit states which are sensitive to the magnetic field [60]. Therefore this mechanism cannot be used on first order field insensitive 'clock' qubits.

Ion Spacing and Force Asymmetry

A complication with the light-shift gate is that it is affected by the spacing of the ions, as the relative positions of the ions in the intensity standing wave pattern affects the relative phase (but not magnitude) of the forces on each ion. For each collective spin state, the forces on each ion sum together to act on the motional mode, and can be thought of as the sum of two complex vectors.

For a same species pair we can describe the force on each spin state with a separate Rabi frequency ($\tilde{z}F_0 = \hbar \sum \Omega_s |s\rangle\langle s|$ as suggested earlier):

$$\begin{aligned} \Omega_{\uparrow\uparrow} &= (1 \pm e^{i\phi_m})\Omega_{\uparrow} \\ \Omega_{\uparrow\downarrow} &= \Omega_{\uparrow} \pm e^{i\phi_m}\Omega_{\downarrow} \\ \Omega_{\downarrow\uparrow} &= \Omega_{\downarrow} \pm e^{i\phi_m}\Omega_{\uparrow} \\ \Omega_{\downarrow\downarrow} &= (1 \pm e^{i\phi_m})\Omega_{\downarrow}\,. \end{aligned} \tag{3.48}$$

$\Omega_{\uparrow}$ and $\Omega_{\downarrow}$ are the force Rabi frequencies for a single ion, $\phi_m = \Delta\mathbf{k} \cdot \mathbf{d}$ is the standing wave phase difference between the ions separated by $\mathbf{d}$, and $\Delta\mathbf{k}$ is the wavevector difference between the Raman beams, or equivalently the wavevector of the standing wave. $|\Delta\mathbf{k}| = 2\pi/\lambda_z$ with λ_z the standing wave wavelength, and the positive and negative signs apply for the in-phase (COM) and out-of-phase (Stretch) modes respectively. The polarisations of the Raman beams are chosen to give a π phase difference between the force on the up and down states, $\mathrm{sgn}(\Omega_\downarrow) = -\mathrm{sgn}(\Omega_\uparrow)$, to maximise the differential light-shift[16].

The geometric phase acquired by each spin state $\Phi_s \propto \Omega_s^2$. For a maximally entangling gate we want

$$\begin{aligned} \Phi_{\uparrow\uparrow} &= \Phi_{\downarrow\downarrow} = \Phi_{\text{even}} \\ \Phi_{\uparrow\downarrow} &= \Phi_{\downarrow\uparrow} = \Phi_{\text{odd}} \\ |\Phi_{\text{even}} - \Phi_{\text{odd}}| &= \pi/2\,. \end{aligned} \tag{3.49}$$

When the force on the two single ion spin states have equal magnitude (in atoms without hyperfine structure, e.g. $^{40}Ca^+$ or $^{88}Sr^+$), $\Omega_\downarrow = -\Omega_\uparrow$, we can see that a phase difference of $\phi_m = n\pi$ is optimal, with n an integer, as then either the even or odd states have zero force and acquire no geometric phase. This corresponds to an ion spacing of $|\mathbf{d}| = d = n\lambda_z/2$; both integer and half integer multiples of the standing wave wavelength. If $\phi_m = (n+1/2)\pi$ or $d = (n/2+1/4)\lambda_z$ however, $\Phi_{\text{even}} = \Phi_{\text{odd}}$, so no phase difference can be accrued and no entanglement created. In between these values the gate is inefficient, as some of the geometric phase created becomes an unneeded global phase, so we need larger force magnitudes to amplify the two-qubit phase until it reaches $\pi/2$. We can quantify this efficiency as

$$\Theta = \frac{|\Phi_{\text{even}} - \Phi_{\text{odd}}|}{\Phi_{\text{even}} + \Phi_{\text{odd}}}\,. \tag{3.50}$$

For dissimilar force magnitudes $|\Omega_\uparrow| \neq |\Omega_\downarrow|$ however (in atoms with hyperfine structure,

[16] The sign or signum function is defined $\mathrm{sgn}(z) = z/|z|$.

e.g. $^{43}Ca^+$), while the two odd spin states acquire equal geometric phase, the phases acquired by the even states differ, causing the ions to accumulate an additional unneeded single qubit phase during the gate. This problem can be solved by setting the ion spacing to cancel the force on the even spin states, either integer spacing ($\phi_m = 2n\pi$, $d = n\lambda_z$) in the case of the out-of-phase mode *or* half integer spacing ($\phi_m = (2n+1)\pi$, $d = (n+1/2)\lambda_z$) in the case of the in-phase mode. Only the odd spin states acquire geometric phase and the gate has unit efficiency.

Mixed Species Gate

For gates between ions of different species the situation is now more complex. Each spin state of each ion has a different force magnitude in general, due to their different Raman coupling strengths and different Lamb-Dicke parameters, including different mode eigenvector magnitudes. Now the state Rabi frequencies are

$$\begin{aligned}
\Omega_{\uparrow\uparrow} &= \Omega_{\uparrow,1} \pm e^{i\phi_m}\Omega_{\uparrow,2} \\
\Omega_{\uparrow\downarrow} &= \Omega_{\uparrow,1} \pm e^{i\phi_m}\Omega_{\downarrow,2} \\
\Omega_{\downarrow\uparrow} &= \Omega_{\downarrow,1} \pm e^{i\phi_m}\Omega_{\uparrow,2} \\
\Omega_{\downarrow\downarrow} &= \Omega_{\downarrow,1} \pm e^{i\phi_m}\Omega_{\downarrow,2} \, ,
\end{aligned} \tag{3.51}$$

with the positive and negative signs for in-phase and out-of-phase modes respectively. Now however, no ion spacing can completely cancel the force on any collective spin state, so unit gate efficiency Θ is not possible. The best we can do is maximise the force difference between the even and odd spin states.

If neither species has hyperfine structure, we retain some of the symmetries of the same species case, and both integer and half integer spacings, $\phi_m = n\pi$ or $d = n\lambda_z/2$, are optimal and maximise Θ.

If either or both species has hyperfine structure however, all four spin states acquire a different geometric phase and we again get an unwanted single qubit phase. To deal with this complication we symmetrise the gate by driving two complete loops in phase space, each acquiring half the necessary geometric phase, with a global π-pulse in the middle. Flipping the spins halfway through then ensures that the even and odd geometric phases separately match:

$$\begin{aligned} \Phi_{\uparrow\uparrow} &= \Phi_{\downarrow\downarrow} \propto \Omega_{\uparrow\uparrow}^2 + \Omega_{\downarrow\downarrow}^2 \\ \Phi_{\uparrow\downarrow} &= \Phi_{\downarrow\uparrow} \propto \Omega_{\uparrow\downarrow}^2 + \Omega_{\downarrow\uparrow}^2 \,. \end{aligned} \tag{3.52}$$

This two loop arrangement also suppresses several technical error sources shared in common with same species gates, such as unequal illumination of the ions, qubit de-phasing and residual single beam light-shifts.

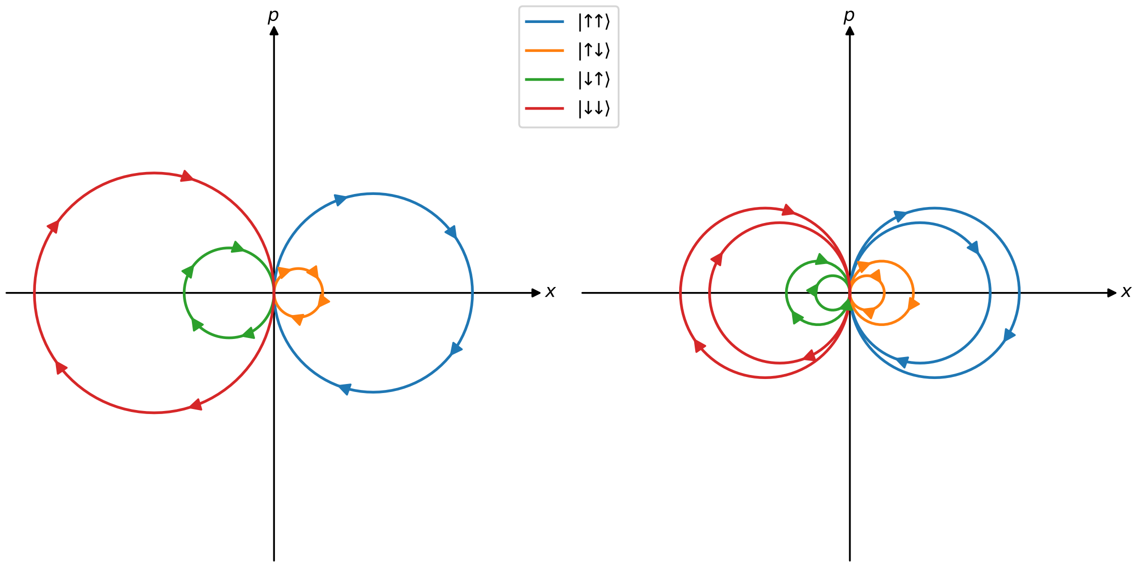

Figure 3.1: Motional phase space trajectories in a rotating frame during a mixed species gate. In a single loop arrangement (Left), hyperfine structure in one species causes all four spin states to acquire a different geometric phase (proportional to the loop area). Splitting into two loops and flipping the spins midway (Right), symmetrises the phases acquired by the even and odd states. Shown at an optimal ion spacing and without Walsh modulation for clarity (see section 5.1.4).

If only one species has hyperfine structure the integer and half integer spacing is again optimal. If both species have hyperfine structure, either the integer *or* half integer spacing is optimal, depending on the sign and magnitude of the force on each ion's spin state and the parity of the motional mode.

3.6 Two-Qubit Gate Error Sources and Mitigation

There are many sources of error which can reduce the fidelity of two-qubit gates. They can be broadly split into technical errors, which arise from imprecise control (i.e. noise or drift) of the trapping fields, laser pulses, etc. and non-technical errors which arise even with perfect control of experimental parameters. These errors are described in more detail in [32,39], but the physical origin of non-technical errors significant to these experiments are discussed below.

3.6.1 Thermal Errors

Though ideally we would like to cool all the motional modes of an ion string to the ground state, we are always left with a non-zero temperature in some modes before we apply the gate operation. We also unavoidably see heating during the gate operation, with both situations causing thermal errors.

Heating Errors

A cooled ion sitting in a trap experiences heating due to noise and fluctuations in the trapping field it sees. This heating has been investigated many times and cannot be explained by Johnson noise or noise in the voltages applied to the trap. It seems instead to be associated with the structure and contamination of the electrode surfaces, which could have fluctuating patch potentials [21, 30, 66]. Whilst this 'anomalous heating' can be reduced with ion-beam or laser cleaning, increasing the ion-electrode distance or operation at cryogenic temperatures, the error caused is still significant.

Heating acts to linearly increase the ions thermal occupation of a mode over time with gradient $\dot{\bar{n}}$, the heating rate for that mode. The error in the gate arises due to the random displacements in phase space the heating causes, which varies shot to shot. When recombining the spin states at the end of the gate, this causes a reduction in contrast exponential in both the square of the phase space separation of the wavepackets during the gate and the size of the heating displacements [67]. Therefore gates where wavepackets make smaller excursions in phase space, and gates of shorter duration, have lower errors. The former can be achieved by maximising the gate efficiency Θ, or using gates with multiple loops.

Carrier Suppression

Another effect of non-zero temperature in some modes is suppression and dephasing of the carrier. As described in section 3.4.2, the carrier Rabi frequency is dependent to second order on the mode occupation n, and is therefore different for each phonon number state (or Fock state) in a thermally occupied mode. This leads to an incoherent spread of force magnitudes during the gate which leads to error. As shown in [32], the error caused by a warm spectator mode is similar to that caused by a warm gate-mode, though the dynamics of the latter are more complex.

Kerr Cross Coupling

The oscillation frequencies of the modes in an ion crystal are not completely independent. Their Coulomb repulsion gives rise to a non-linearity between pairs of modes, which causes the frequency of one to be shifted due to the occupation of another, in a cross-Kerr interaction. This occurs between all modes which are not purely symmetric, as their mode frequencies depend on the mean separation between ions.

In a same-species two ion crystal where we conduct the gate on one of the axial modes, we are most concerned with the Kerr cross-coupling between the axial stretch mode and the two radial rocking modes, which each add a term $H_{\mathrm{Kerr}} = \hbar\chi N_{\mathrm{str}} N_{\mathrm{rock}}$ to the Hamiltonian. N_{str} and N_{rock} are the phonon number operators for the axial stretch and a radial rocking mode respectively, with $N_j = a_j^\dagger a_j$. Physically this occurs because, as the occupation (and so amplitude) of the rocking modes increase, the mean separation between the ions increases, reducing the stretch mode frequency. For gates conducted on the stretch mode this appears as a gate detuning error. If the rocking modes were in pure Fock states a corrected detuning could be found, but in their usual thermal states the incoherent spread of gate detuning leads to error. An expression for χ for a same-species crystal can be found in [68] which shows that, apart from cooling the radial modes better, avoiding the resonance condition $\omega_{\mathrm{str}} = 2\omega_{\mathrm{rock}}$ lowers the error from Kerr cross-coupling.

For a mixed species crystal, we no longer have purely symmetric motional modes. Where previously for two ions the in-phase axial and in-phase radial modes did not participate, all modes of the system now have a cross-Kerr interaction with each other. We can therefore no longer rely on in-phase axial mode gates being free of Kerr cross-coupling errors, as the mode will couple to the axial out-of-phase and all radial modes. For our case of a $^{43}\mathrm{Ca}^{+}$-$^{88}\mathrm{Sr}^{+}$ crystal with a mass ratio close to 2, we are still somewhat close to having a symmetric and an anti-symmetric mode (see section 2.4.1). The out-of-phase mode therefore is still significantly more susceptible to Kerr cross-coupling errors as compared to the in-phase mode, principally from warm radial modes.

3.6.2 Off-Resonant Excitation Errors

Any pulse of radiation of finite duration from a monochromatic source has a non-zero spread of frequencies. Whilst during our gate operation we aim to excite only a single motional mode, errors can arise when either the carrier or other nearby motional mode transitions are off-resonantly excited. Excitation of the carrier will cause unwanted spin flips, while on spectator modes it will drive loops in phase space that in general are not closed at the gate time, causing qubit decoherence.

As shown in section 3.2.1, we expect the excitation probability of nearby transitions to be proportional to Ω'^2/δ'^2, where Ω' and δ' are the effective Rabi frequency of, and the laser detuning to, the off-resonant transition. We therefore only need worry about couplings to the nearest of these transitions. Compared to the Rabi frequency on the gate mode or other nearby motional modes, Ω' on the carrier is enhanced by a factor $1/\eta$, so its off-resonant errors dominate and we can ignore those on spectator modes.

Carrier Excitation

For the Mølmer Sørensen gate where each tone is symmetric about the carrier, the light-shifts from each cancel but the off-resonant excitations sum. For the light-shift gate, one might initially assume that since there is no qubit carrier coupling, we cannot see off-resonant errors of this form. However, over the course of the gate a time-dependent light-shift accumulates on the ions as the standing wave pattern travels over their position. In the frequency domain, this looks like a light-shift on the carrier, but at zero frequency. This phase shift depends on the initial optical phase of the standing wave pattern, which we do not control, so the output of each gate has a random single qubit phase applied, which can drastically reduce the fidelity.

Pulse Shaping

The solution to this problem is to slowly turn on and off our gate force, by shaping the otherwise sharp edges of the square laser pulses. This reduces the spread of frequencies seen by the ions, and in the time domain 'averages' over the initial and final optical phases. The exact shape of the pulse edges is not important, as long as the pulse falls off with a timescale $t_{\text{shape}} > 1/\delta'$, where in our case of coupling to the carrier $\delta' = \omega_z$, the gate mode frequency [32].

3.6.3 Photon Scattering Errors

The Raman coupling used for the light-shift or Mølmer Sørensen gate can cause errors from photons scattered off the upper states, as population in these states decay spontaneously. As described in section 3.3.1, this comes in the form of Rayleigh scattering, where the initial state (one of the qubit states) is the same as the final state scattered into, and Raman scattering, where the two are different.

The latter process emits a photon whose polarisation and frequency are entangled with the final state of the ion, which is then measured by the environment collapsing any superposition and destroying the coherence. The Rayleigh scattered photon is not entangled with the internal state of the ion, but if the scattering rate for each qubit state is different it still leads to a dephasing of the qubit. Both types of scattering also lead to a dephasing of the ion motion due to the recoil momentum imparted by the emitted photon, but the error this creates is negligible in our case.

Reducing Scattering

Increasing the Raman detuning Δ reduces the scattering rate faster than it reduces the Rabi frequency, as $\Gamma_i \propto 1/\Delta^2$ while $\Omega_{\mathrm{Raman}} \propto 1/\Delta$. Their ratio, the error rate, therefore also reduces with detuning [32]. In practice for a particular Rabi frequency we can trade-off Raman detuning with laser power until we run out of the latter.

However even with very large powers we cannot suppress the scattering rate entirely. Ions commonly used for quantum information often have ground $S_{1/2}$ level qubits, which are driven by Raman transitions coupled to both the $P_{1/2}$ and $P_{3/2}$ levels, separated by fine structure splitting ω_f. In this situation the Rayleigh scattering rate has an asymptotic limit with detuning of $\sim\Gamma/\omega_f$, where Γ is the spontaneous decay rate/natural linewidth of the P levels [69]. Whilst this scattering rate from both qubit states can be balanced to prevent dephasing, in species with low lying D levels (e.g. $^{43}\mathrm{Ca}^+$ and $^{88}\mathrm{Sr}^+$ used in this thesis) there is a small but unavoidable decay to them from the P levels.

Scattering to the D levels therefore represents the only fundamental error source, which cannot be removed by technical means. The rate is proportional to the total scattering rate and causes an error for all Raman operations in these ions. For single qubit gates this puts a lower bound of $\sim 10^{-6}$ on the error rate, while for two-qubit gates it is $\sim 10^{-4}$.

4

Apparatus

The apparatus used to trap, control and measure the ions for the experiments in this thesis is described in this chapter. At its heart it consists of the trap, an electrode structure which confines the ions, housed inside a vacuum chamber with windows on all sides. Lasers to cool, manipulate and fluoresce the ions are shined in and are focussed on the ions position, with microwave and RF tones to drive their qubit frequencies applied to the trap electrodes. Water cooled coils around the vacuum chamber apply a magnetic field, which sets an axis for laser polarisations and allows access to a long-lived 'clock' qubit in $^{43}Ca^{+}$. An imaging system above the vacuum chamber collects the ion fluorescence to readout the state of the ions. The lasers are housed elsewhere in the laboratory where they are frequency stabilised, switched on and off, and split into multiple frequency components as necessary before delivery to the vacuum system down optical fibres.

This hardware is essentially the same as that described in [39], but with some further improvements and changes, notably the removal of the 408 nm laser and introduction of a 674 nm $^{88}Sr^{+}$ quadrupole laser. It was designed for gate experiments on $^{43}Ca^{+}$ and $^{88}Sr^{+}$, and also ion-photon entanglement experiments requiring high efficiency photon collection [70].

4.1 Ion Trap

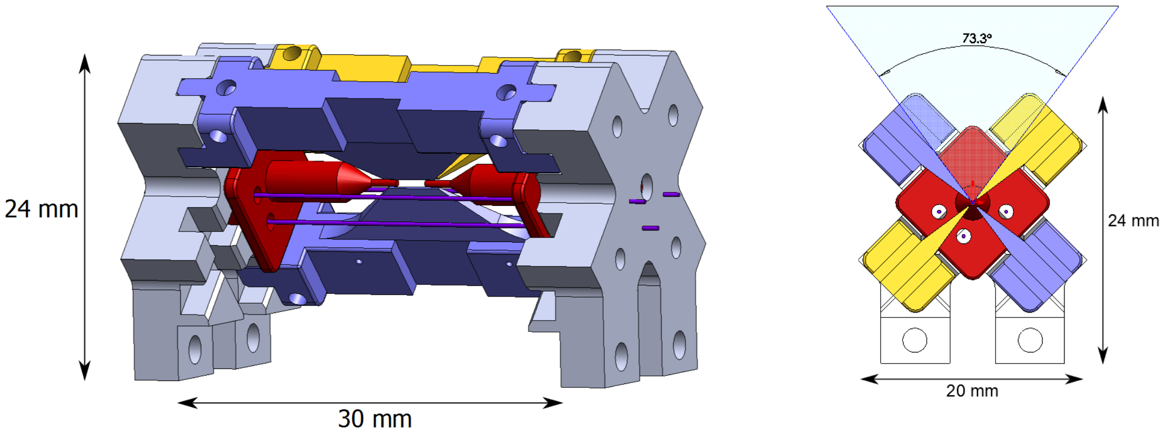

Figure 4.1: Rendered CAD model of the ion trap. Ceramic 'X' shaped supports hold the endcap (red) and blade (blue & yellow) electrodes. On the left the foreground blade has been hidden to show the compensation electrodes (purple). The ion state is read out by collecting fluorescence over a large angle above the trap. Figure from [71].

The trap used in this system, shown in the figure above, is an improvement on the 'Blade' linear Paul trap used in our previous experiments [32,72]. Four blade electrodes are arranged around the trap axis with one pair driven by a radio frequency (RF) potential and one pair grounded. This produces an oscillating quadrupole field in the radial plane which gives a radially confining pseudo-potential, with endcap electrodes on the trap axis at a DC potential applying a static axial harmonic field. Ions whilst ideally static also unavoidably experience oscillation at the RF driving frequency both radially and axially, known as micromotion. Thin wires running parallel to the trap axis act as compensation electrodes to modify the DC trapping potential so its minimum coincides with that of the RF quadrupole potential.

4.1.1 Design

To collect more ion fluorescence for separate ion-photon entanglement experiments, the upper window (viewport) of the vacuum system was designed to be recessed into the vacuum chamber, to bring the main imaging lens closer to the ions. To make room for this re-entrant viewport, the trap has been shrunk by 50% as compared to our previous designs, reducing the ion-window distance to 11 mm. Thinner blades with cut out sections allow an imaging cone full angle of 73.3°, equivalent to a numerical aperture (N.A.) of 0.6. The trap also has a slightly increased endcap spacing, as a compromise allowing both micromotion sideband addressing (not used in this thesis) and a fairly constant carrier amplitude across an ion string [72–74]. The trap dimensions and geometry factors α are presented in table 4.1, where the latter scale the strength of the trapping field created by the electrodes (see section 2.4).

z_0	1.15 mm
ρ_0	0.5 mm
α_x	0.486
α_y	0.487
α_z	0.207

Table 4.1: Characteristic trap dimensions and geometry factors α. z_0 and ρ_0 are the ion-endcap and ion-blade electrode spacings respectively. $\alpha_x \neq \alpha_y$ as the blades are driven asymmetrically.

4.1.2 Construction

The four trap blades are aligned and located in slots in Macor® machinable glass-ceramic holders, which are rigid enough and can be machined precisely enough to minimise trap geometry distortion. The blades and endcaps are made from 316LN grade non-magnetic austenitic stainless steel, which resists work-hardening during machining, which would otherwise transform machined regions into the magnetic martensite phase [75]. This choice of material ensures trap electrodes are not magnetised, and the ions see a uniform magnetic field environment from only the external coils. Electrical discharge machining was used to shape the electrodes instead of conventional cutting techniques, as this also reduces the possibility of work-hardening. We measured the magnetisation of the completed blades and found a maximum field of 0.02(5) G, consistent with zero given the Hall effect sensor's minimum measurable field of 0.04 G.

To smooth the $\sim$20 μm sized bumps and pits on the blade and endcap electrode tips left after machining, micron-size diamond lapping compound was used to polish them. As these surfaces are the closest to the ion, rough or oxidised areas left from the machining process could cause imperfections or fluctuations in the trapping potential, leading to an increased heating rate from ground state of motion [21, 30]. Sequentially finer compounds, containing $^1/_4$ micron diamonds in the final stage, were used to polish the electrodes by hand, to produce a mirror finish (figure 4.2).

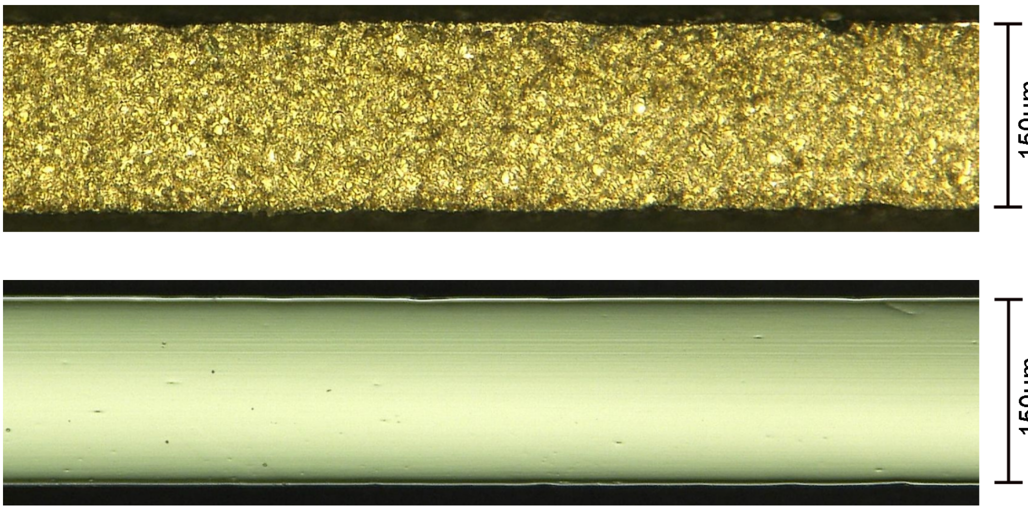

Figure 4.2: Image of a trap blade electrode tip (x400 optical zoom) before (Top) and after (Bottom) polishing with abrasive diamond lapping compound. The final stage of polishing with $^1/_4$ micron diamond paste produced a mirror finish ensuring a uniform potential in the trap.

Unfortunately this step proved futile, as after construction of the trap we measured a heating rate two orders of magnitude higher than the heating rate of the previous trap, which has similar ion-electrode separations and on which no treatment of the electrode surfaces was carried out [32]. While trap heating is known to have a strong inverse dependence on ion-electrode spacing, its origin and reliable mitigation strategies (except cryogenic trap operation) are poorly understood [66]. We therefore cannot give a reason for this higher heating rate in an ostensibly similar trap.

For a single $^{43}Ca^{+}$ ion trapped with a fairly typical 1.92 MHz axial motional frequency, we measured a heating rate of $\dot{\bar{n}} = 75(10)$ quanta/s. No noise sources near the motional frequencies were detected in the trap RF or DC drive, with the former being filtered by the helical resonator and the latter by a filter board before the trap regardless. For two $^{43}Ca^{+}$ ions at the same trap strength we measured $\dot{\bar{n}} = 145(10)$ quanta/s and $\dot{\bar{n}} = 0.5(1)$ quanta/s for the axial COM and stretch modes respectively, with the low value of the latter expected from the spatially uniformity of the heating noise. These results, along with finding an anomalous scaling of the axial heating rate with radial trap strength, are reported in more detail in [39].

4.1.3 Drive

The trap needs stable RF voltages on the blades, and DC voltages on the endcaps and compensation electrodes, which can be adjusted to change the trap depth for loading, recrystallisation, re-ordering and normal operations.

The production and distribution of the RF and DC voltages needed to drive the trap and compensation electrodes are identical to the description in [39] and only slightly changed from [32]; a brief description follows here.

RF

A synthesiser[1] connected to a variable attenuator[2] sends RF to a fixed +46 dB power amplifier[3], before the voltage is stepped up to just under 1 kV using a copper helical resonator RF transformer [76]. A small coil induces a voltage in a larger coaxial helix all inside a cylindrical shield, grounded to the vacuum chamber and to the undriven trap blades [77]. When one side of the helix is connected to the driven blades it applies

[1]Hewlett Packard, 8656B
[2]Mini-circuits ZX76-31R5-PP-S+, 0 to 31.5 dB attenuation in 0.5 dB steps
[3]Frankonia FLL-25, 1 dB compression point at 20 W output power

a $\times 35$ stepped up voltage at its 28 MHz resonance frequency.

DC

A custom digital-to-analogue converter (DAC) supplies DC voltages for the endcap and compensation electrodes between -240 V and $+240$ V in 7 mV steps. They are applied via a filter board connected to the electrical feed-throughs on the underside of the vacuum system, which suppresses noise at the trap motional frequencies which would cause heating. A separate input on the board, which bypasses the filters, allows a 'tickle' signal to be applied to an endcap or compensation electrode, to excite and find the motional mode frequencies. With 194 V on the endcaps and $\sim$776 V (zero-peak) on the blades, we obtain 1.9 MHz axial and 4.2 MHz radial trap frequencies with $^{43}Ca^{+}$.

4.2 Vacuum System

The trap is mounted inside a 6" 'spherical octagon' vacuum chamber[4] oriented horizontally, with 8 small CF40 ports on the sides and 2 large CF100 ports on the top and bottom for optical, electrical and vacuum pumping access. We attach 6 broadband anti-reflection (BBAR) coated quartz viewports on to the sides[5], along with a 4-pin 5kV electrical feed-through[6] for the trap RF. A custom base flange, with four[7] and eight[8] pin feed-throughs for the trap DC, microwaves and Ca/Sr ovens, supports a cradle holding the trap, with its centre coincident with the centre of the vacuum chamber. A combined non-evaporable getter and ion pump[9], with a nominal pumping speed of 200 l/s for hydrogen are attached to the remaining side port of the vacuum

[4] Kimball MCF600M-SphOct-F2C8-A
[5] Allectra 110-QZ-UV-C40-316LN
[6] LewVac FHP5-50C4-40CF, 5 kV, 50 A
[7] LewVac FHP1-C4-W, 1 kV, 25 A
[8] LewVac FHP1-C8-W, 1 kV, 25 A
[9] SAES NEXTorr D 200-5

chamber through a custom CF63 internal diameter T-piece, onto which are connected an ion gauge[10] and valves[11]. On top a BBAR coated custom re-entrant viewport is mounted, with the glass 11 mm away from the ion to increase the light collection angle. Most parts, including the spherical octagon, viewports and T-pieces, along with the trap cradle and fasteners are made from various grades of 316 non-magnetic austenitic stainless steel, similarly to the trap electrodes.

[10]Varian UHV-24P
[11]VG Scienta ZCR40R

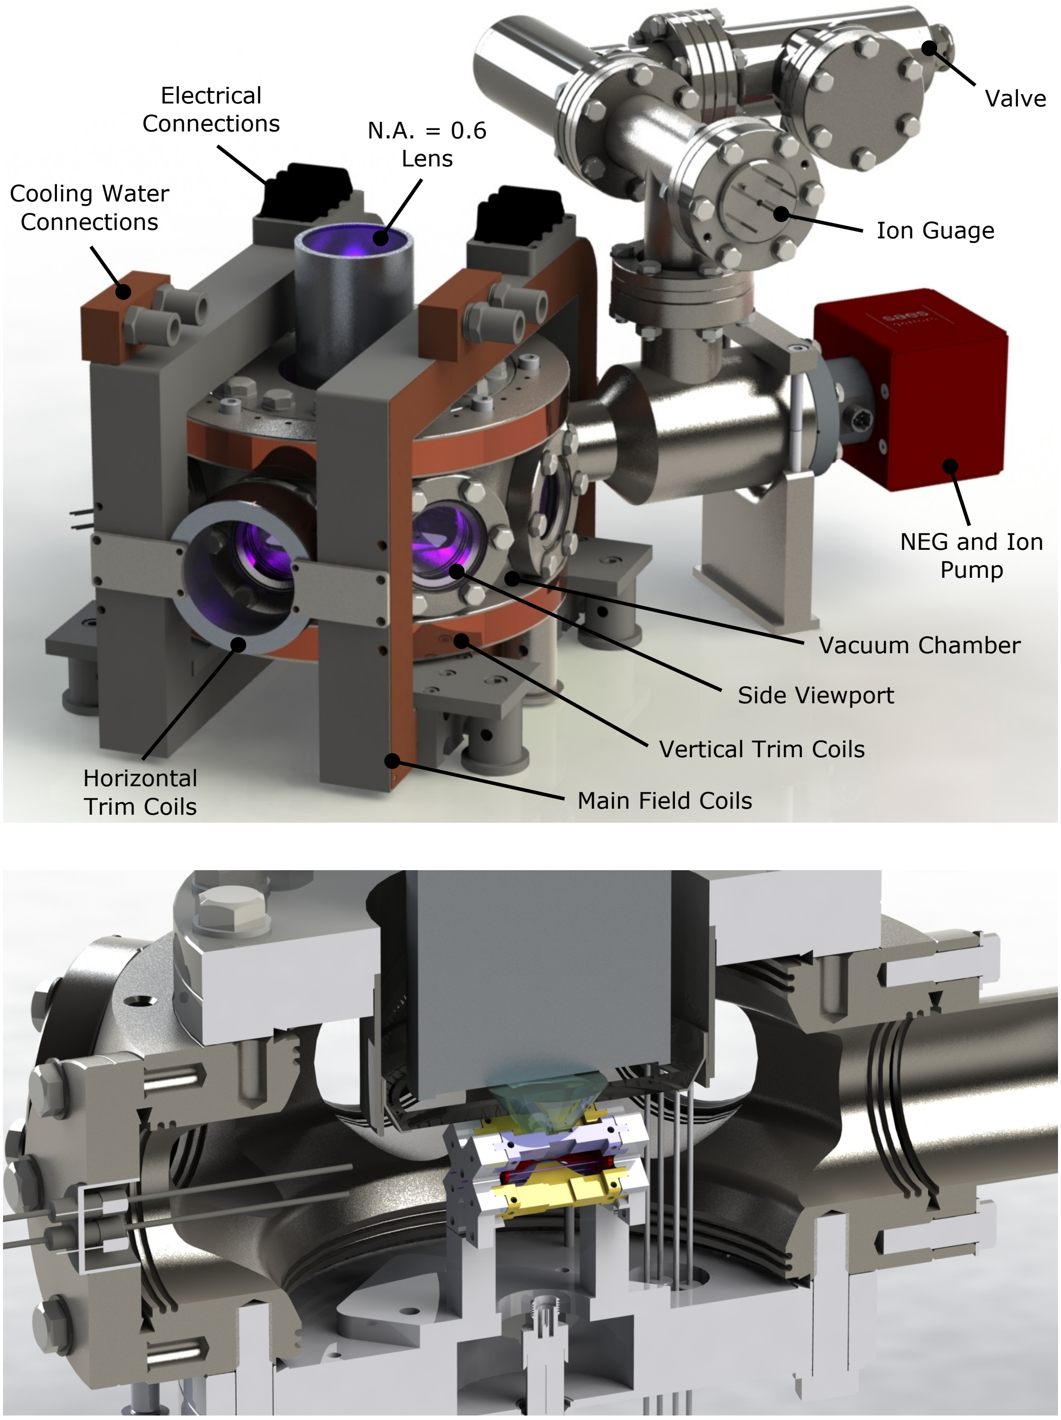

Figure 4.3: Rendered CAD model of the vacuum system, exterior and cross-section. Custom-wound foil electromagnets closely surround a vacuum chamber with the ion trap at its centre. Electrical feed-throughs and windows for laser and imaging access are attached, along with a combined getter and ion pump, ion gauge, and valves. A large N.A. = 0.6 lens drops down from above into a re-entrant viewport to image the ions.

The atomic source ovens are 2 mm diameter thin-walled stainless steel tubes, made to a standard design [78]. Mounted in the corner of the vacuum system, they have a 0.3 mm diameter hole pointing towards the trap centre, through which a flux of atomic vapour is emitted when heated by a current. The ovens were filled with ∼30 mg of either isotopically-enriched calcium[12] (12% $^{43}Ca^{+}$, 88% $^{40}Ca^{+}$), or natural-abundance strontium (82.6% $^{88}Sr^{+}$) granules, before the ends were crimped closed. The strontium oven was filled and sealed in an inert atmosphere due to its rapid oxidation rate, and both were installed as the final step before the system was sealed and pumped down, to minimise their exposure to air.

4.2.1 Assembly

Before assembly, to remove oils and other residues left on vacuum system parts (after machining or handling) that would vaporise and limit the final pressure, the following cleaning procedure was used:

1. Neutracon[13] ∼3% solution in tap water | 50°C for 15 mins
   - Rinse in tap water
   - Dry in air and with argon
2. Deoxidine[14] 10% in deionised water | 35°C for 5 mins
   - Rinse in deionised water
   - Dry with argon
3. Acetone | 35°C for 10 mins
   - Rinse in deionised water
4. Deionised water | 35°C for 10 mins
   - Dry with argon

For each stage, large parts suspended with copper wire, or small parts loose or in

[12] Oak Ridge National Laboratory, TN, USA
[13] Decon Neutracon
[14] Trimite Deoxidine 624

stainless steel mesh balls, were placed in vessels of the solvent or detergent in an ultrasonic bath. The bath was heated and the ultrasonic frequency swept, to avoid areas at nodes in the standing wave pattern remaining uncleaned. Any obviously dirty areas were cleaned manually first with detergent or isopropanol. Parts were dried by blowing off excess solvent using argon (as it is supplied cleaner than compressed air), to prevent dissolved residues re-adhering as the solvent evaporated. Wearing nitrile gloves, a face mask and hairnet throughout, the first step uses diluted Neutracon, a mild neutral detergent, to remove the heaviest greases and oils. Deoxidine, an acidic detergent, was then used to remove remaining deposits, with parts left in the solution for only 5 mins and thoroughly rinsed due to its aggressive nature. After this step parts appeared noticeably cleaner with water wetting the surface uniformly rather than 'beading off'. The final steps of acetone and deionised water were used to remove remaining traces of non-polar and polar residues respectively, with the acetone preceding as it evaporates too fast to be dried. Parts were then double-wrapped in clean UHV aluminium foil[15] to prevent contamination before assembly, making sure they were dry to avoid corrosion.

[15] All-foils UHV foil. Food grade aluminium foil is often coated in oil to aid lubrication and hence cannot be used.

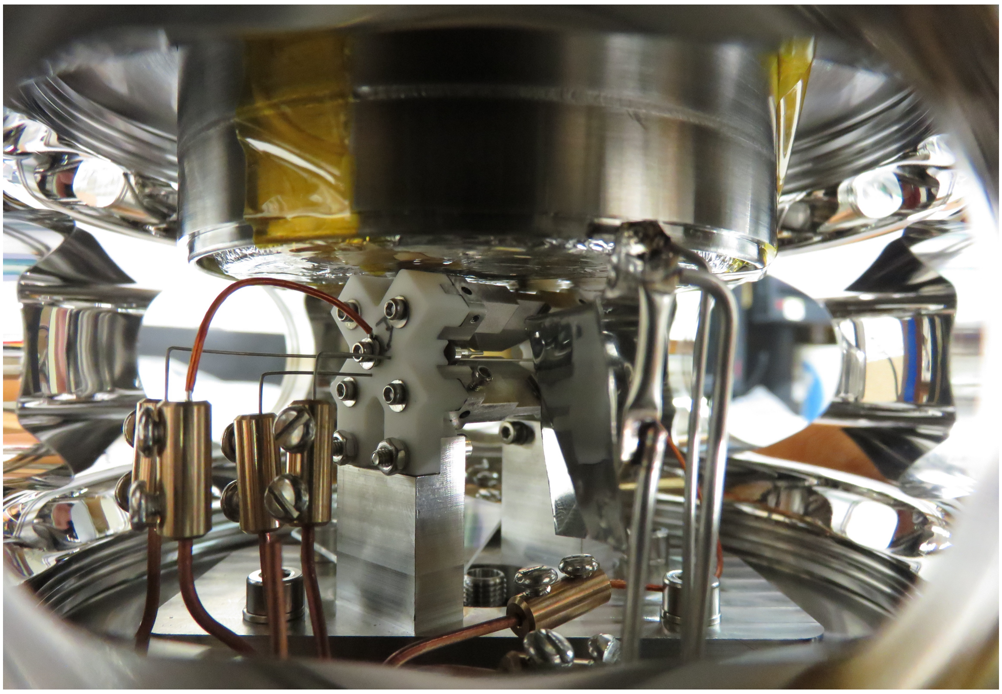

Figure 4.4: Photograph of the assembled trap inside the vacuum chamber, with the Ca oven in the foreground. The trap rests on a cradle bolted to the bottom feed-through, with three compensation electrodes connected on the left. The re-entrant viewport above rests very close to maximise the collection of ion fluorescence light, covered with foil at this stage in the assembly to prevent accidental damage.

During assembly we baked the system to high temperature three times, to reduce the gas desorption rate from the inside surfaces of the vacuum system when in use. We first baked large metallic parts (unassembled) in air at 350°C for one day, building a thick straw-coloured oxide layer that blocks hydrogen outgassing from the steel bulk. Next, with blank flanges used in place of delicate viewports and feed-throughs, and loose trap parts inside, the assembled system was then ‘hard’ baked under vacuum for 2 weeks at 350°C, the temperature limit being set by the ion pump. With the getter element activated a pressure of 6.3×10^{-11} Torr was reached. The final ‘soft’ bake on the fully assembled system was for 2 weeks at 190°C, with both the temperature and ramp rate of 1°C/min just below the limit set by the viewports. The atomic source ovens along with the system ion gauge were heated hotter in order to degas them. Maintained by the system pumps, the final pressure attained is below 5×10^{-12} Torr, the X-ray limited measurement floor of the ion gauge. With the typical trap depths mentioned in 4.1.3, we can hold ions indefinitely and see decrystallisation events caused by background gas collisions every 5-10 mins with 2-ion strings.

4.2.2 Magnetic Field Coils

To achieve 'clock' qubits in $^{43}Ca^{+}$, where the qubit transition frequency is insensitive to magnetic field noise to first order [14], we apply a quantisation field of 146 G using a pair of rectangular water-cooled foil-wound coils. Built by Stangenes Industries, Inc. to a custom design for this system, they are made from layered sheets of copper separated by thin insulating material, and have a filling factor much higher than traditional wire-wound coils. A copper water-cooling plate is attached to the outer face of each coil, electrically insulated by a thin thermally conductive layer, with the whole construction made rigid by a baked epoxy surface. The coils themselves have inner dimensions of 226 mm $\times$ 134 mm, cross section approx. 27 mm $\times$ 27 mm, and the two coils are separated by 76 mm, to fit as closely as possible around the vacuum system leaving clearances of a few millimetres.

To produce 146 G the coils are supplied in series with 56.8A, and dissipate roughly 216W of electrical power, including that in the large gauge cables from the power supply[16], a significant improvement on the 1.35 kW that was dissipated in the home-made wire-wound coils of a previous experiment [44]. When water-cooled at 1 litre/minute they are specified to run 20°C above the water inlet temperature, which we supply at 10.5°C. The coils are connected in parallel with several meters of 10 mm piping to a chiller unit outside the lab operating at a pressure of 8 bar, through a flow rate sensor connected to a safety interlock on the coil power supply.

To trim out the earth's and any stray lab fields, and to align the quantisation axis

[16] Agilent 6671A, capable of supplying the coils with sufficient current to reach 288G, where the next 'clock' qubit transition lies.

onto the optical pumping beam, we use two further orthogonal wire-wound trim coils. A pair of coils providing a field in the vertical direction are bolted onto the top and bottom of the vacuum chamber, closely mounted around the CF100 flanges, with the latter supporting the vacuum chamber on feet. Each vertical trim coil has 70 turns, with an internal diameter of 172.7 mm, cross section approx. 16 mm × 16 mm and are separated by 73.8 mm. Horizontal coils provide a field parallel to the optical table plane and perpendicular to the field of the main coils, and are bolted to them with brackets. These each have 117 turns, an internal diameter of 65 mm, cross section approx. 4.5 mm × 25 mm and are separated by 230 mm.

Field Stabilisation and Servo

To stabilise the current being provided to the main coils, a custom active feedback circuit was used. This was designed and constructed by Benjamin Merkel, with further details in [79]. The coil current is measured[17], and the noise fed back with opposite phase. The ambient magnetic field noise, primarily 50 Hz and harmonics from the mains, can be measured with the ion and actively fed back. With this circuit, magnetic field noise is reduced to 45(1) μG RMS, representing only 0.31(1)ppm on our 146 G field.

Shifts in the ambient magnetic field, due to a nearby goods lift and superconducting magnet experiments, cannot be directly stabilised by this circuit. Instead, we regularly measure the $^{43}Ca^{+}$ stretch qubit frequency by measuring either side of its expected resonance, and shift an offset in the stabilisation circuit to correct shifts as necessary. This B-field servo is run regularly between experiments to catch slow drifts, but cannot compensate if the ambient field changes rapidly.

[17]LEM IT 405-S ULTRASTAB

4.2.3 Imaging System

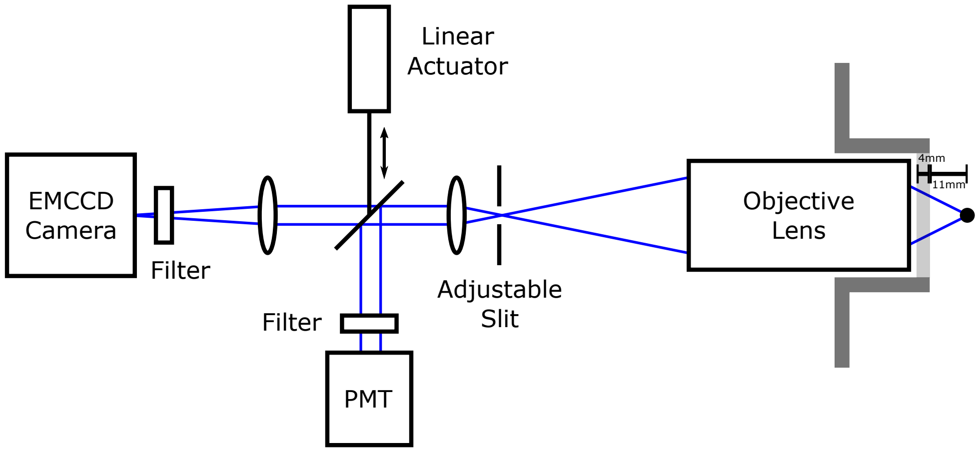

Figure 4.5: Schematic of the ion fluorescence collection system with folding mirrors omitted (not to scale). A large N.A. 0.6 lens close to a re-entrant window images the ions in the trap. An x-y adjustable slit along with shortpass filters reduce background scatter, and a moveable mirror sends light to either a PMT or EMCCD camera.

To readout the ions' state, we collect their fluorescence using a large numerical aperture lens in an imaging system mounted above the vacuum chamber. A few millimetres from the window of the re-entrant top viewport, the custom N.A. 0.6 lens[18] mounted on xyz translation and xy goniometric stages, images the ions onto either a photo-multiplier tube[19] (PMT) or electron multiplying charged coupled device (EMCCD) camera[20]. An adjustable 2-dimensional slit at the focal plane of the lens removes stray laser beam scatter from the trap electrodes and room lights, before an aspheric lens[21] collimates the beam. A mirror on a linear actuator reflects the beam through a filter[22] which blocks light with wavelengths above 440 nm from reaching the PMT, which has a specified counting efficiency of 30%. With the mirror moved aside, the beam can continue instead to the camera through a 125 mm lens and a filter of the same type. The total magnification of the ion image on the camera is 40.2, so two ions separated by 3.5 μm

[18] Photon Gear Inc.
[19] Hamamatsu H10682-02
[20] Andor Technology iXon^{EM}+ DU-897E
[21] Asphericon AFL12-20-S-U
[22] Semrock FF01440/SP-25

in a typical 1.9 MHz axial trap appear 8.8 pixels apart. We primarily use the PMT in this work, with the camera used for loading and diagnostics. The imaging system is housed in a light-tight box, and shutters around the experiment table keep out room light. The PMT dark count rate is $<100\,\mathrm{s}^{-1}$, and we typically see ion fluorescence $\approx 35,000\,\mathrm{s}^{-1}$ on a scattered light background $\approx 300\,\mathrm{s}^{-1}$.

4.3 Microwaves and RF

The microwaves needed to drive the $^{43}\mathrm{Ca}^{+}$ qubit and the RF need for the $^{88}\mathrm{Sr}^{+}$ Zeeman qubit are applied to electrodes inside the vacuum system near the ions. In previous experiments this was done by applying the fields to the grounded trap blades, but in an attempt to increase their Rabi frequency we built a small loop antenna mounted on the base flange directly below the ions. Unfortunately, due to the difficulty of controlling the output polarisation of the radiation, it proved to produce lower Rabi frequencies on the qubit transitions than simply using the blades. We therefore use a grounded blade each to apply the microwaves and RF.

The 2.87 GHz microwaves are produced by summing the frequency from a 2.77 GHz synthesiser with the output of a direct digital synthesiser (DDS) board, whose frequency can easily be varied. Further details on this scheme can be found in [32, 39]. The 408 MHz RF can directly be produced by a DDS. Both are separately amplified to $\sim$10 W, before being delivered to the trap.

4.4 Diode Laser Systems

For Doppler cooling, fluorescence and state preparation of both calcium and strontium we use DL Pro External Cavity Diode Lasers (ECDLs), in blue and IR wavelengths, manufactured by Toptica Photonics. The external cavity is formed by a blazed diffraction grating in the Littrow configuration that reflects the 1st order diffracted light back into the laser diode, and by the rear facet of the diode itself. The laser frequency can be scanned over a large range (~20 GHz) by moving the grating with a piezo-electric actuator, or by changing the diode current. The diode temperature is kept constant using a thermo-electric cooling (Peltier) chip and PID loop, with an optical isolator used to prevent feedback from back reflections in the beam path.

To stabilise each of the laser frequencies, we lock them to individual 1.5 GHz free spectral range Fabry-Perot reference cavities. Part of the laser output is sent along a fibre to the cavities, with sidebands produced by modulation either by an electro-optic modulator (EOM) or directly on the laser diode current[23]. The cavity reflection is used to produce a Pound-Drever-Hall (PDH) error signal [80, 81] to which the laser is locked by feeding back to the grating piezo for slow drifts, and to the laser current for fast fluctuations. Frequency scanning could be achieved by adjusting the reference cavity length with an integrated piezo, but as piezo relaxation after voltage changes causes frequency drifts over a period of several hours, the RF drive frequency of acousto-optic modulators (AOMs) further down the laser's beam path are often adjusted instead.

A sample of each laser output is also sent via fibre to a laser diagnostic system. A 16-way switcher[24] sends light to one of two scanning confocal cavities (which have a frequency spectrum independent of the transverse mode), in either the blue[25] or IR[26], to analyse the optical spectrum, while a wavemeter[27] measures the absolute frequency.

4.4.1 Calcium Lasers

Only a basic outline of the calcium laser system and beam paths set up by Vera Schäfer are given here, more details being found in her thesis [39]. Two blue, 397nm

[23] The latter method is typically used for the IR lasers where we are not sensitive to the sidebands on the ion.

[24] Leoni FibreSwitch mol 1x16

[25] Thorlabs SA200-3B, 1.5 GHz FSR, finesse ~250

[26] Toptica FPI 100, 1 GHz FSR, finesse ~100

[27] High Finesse WS7, 10 MHz resolution

and 393nm, and three IR, 866nm, 854nm & 850nm ECDLs are need to move population around all the relevant levels of $^{43}Ca^+$ (see figure 2.1). To remove unwanted amplified spontaneous emission the light is first reflected from a blazed diffraction grating external to the ECDL. Two pick-offs sample $\sim$100 μW of the beam for the cavity lock and the diagnostic system. For fast switching of the beams with a high extinction ratio, AOMs are used in a double-pass configuration. EOMs are used to apply sidebands to the optical frequency where necessary. At various points, half-wave ($\lambda/2$) plates and polarising beam splitter (PBS) cubes act as variable power splitters, before the light is finally sent along single-mode optical fibre to the experiment table.

4.4.2 Strontium Lasers

Figure 4.6: Photograph of the 19” rack containing the strontium lasers. Three lasers and their controllers are housed below a laser frequency and spectral diagnostic system. In each drawer, light is split into multiple channels for frequency locking, diagnostics and for the experiment. A neighbouring rack houses AOMs for pulse and frequency control, with reference cavities nearby.

To save lab space and attempt to miniaturise the system, all the strontium lasers and the diagnostic system are housed in a 42U 19” rack (figure 4.6). Initially developed by Rustin Nourshargh [82], 422nm, 1092nm & 1033nm lasers on 10 mm thick aluminium breadboards are mounted inside individual 4U drawers on rubber feet. Each breadboard contains ∼1/2 inch sized optical components which, after filtering the laser's amplified spontaneous emission with a grating, split its power into multiple outputs for the experiment, diagnostics and locking. The power of each output can be varied by a $(\lambda/2)$ plate and PBS cube, and the light is coupled into polarisation-maintaining single-mode fibres for the lock and experiment.

The 422nm and 1092nm laser are PDH-locked to reference cavities[28] on a nearby optical table, with the lock found to be robust to sizeable bumps of the (relatively flexible) rack. Two fibre coupled AOMs[29] control the power of the 1092nm and 1033nm light before they are combined using a fibre 50:50 splitter[30] into a single-mode fibre to the experiment table.

To drive the $\hat{\sigma}^{\pm}$ transitions between the $5S_{1/2}$ and $5P_{3/2}$ levels of $^{88}Sr^{+}$ (see figure 2.2) at 146G for optical pumping, we split the 422nm light into two components separated by 545 MHz. Two AOMs[31] at half this frequency, one in +1 order producing the 'upper' component, and the other in -1 order producing the 'lower' component, each shift the frequency of half the light before the beam is recombined and coupled into a fibre to the experiment. A preceding AOM switches both beams on/off, and its frequency tracks slow drifts of the reference cavity to keep the light near resonance. This AOM network is also mounted in a drawer in a separate 19" rack, with the breadboard custom designed to pre-align much of the setup.

4.4.3 Photo-Ionisation Lasers

A two step photo-ionisation (PI) process is used to generate ions from the neutral atomic vapour emitted by atomic ovens [83]. 423nm and 461nm lasers drive the reso-

[28]Stable Laser Systems, 1.5 GHz FSR. Based on a design by the National Physical Laboratory (NPL) UK.
[29]Gooch & Housego FibreQ
[30]Thorlabs PMC1060-50B-APC
[31]Isomet 1250C-829A

nance transitions of neutral Ca and Sr respectively, from their ground levels to upper states which are excited to the continuum by a 375nm laser[32]. A bare diode is used for the latter, as its frequency stability is unimportant, and none of the three PI lasers are locked. Since loading is required relatively infrequently, their free-running stability, along with manual adjustment, is sufficient to stay on resonance with the particular isotope-shifted component needed, e.g. $^{43}Ca^+$, $^{40}Ca^+$, which have resonances split by 620 MHz. The PI beams enter the vacuum chamber nearly perpendicular to the atomic source ovens (see figure 4.7), reducing the Doppler width of the resonances for better isotope selectivity and more efficient ionisation. The PI lasers are housed in a separate lab and controlled remotely, along with the mechanical shutters which switch them. All 3 beams are superimposed using dichroic mirrors into a single-mode fibre to the experiment.

[32]Thorlabs iBeam smart

4.4.4 Beam Delivery

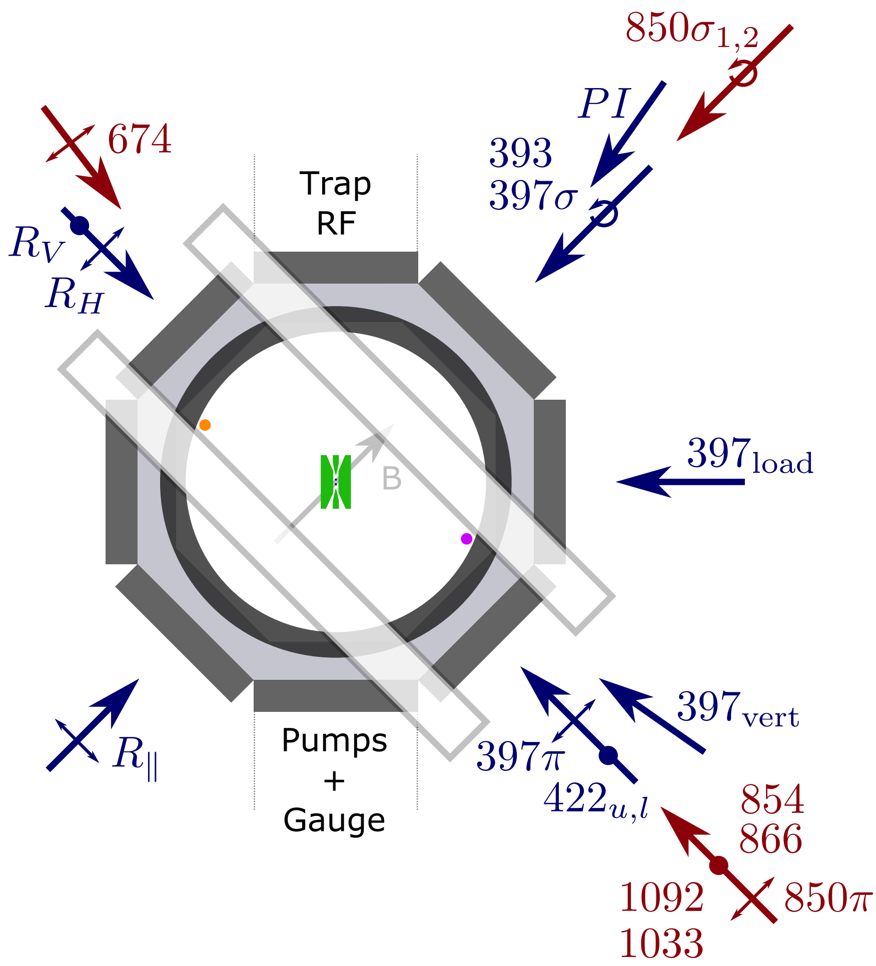

Figure 4.7: Laser beam and polarisation arrangement around the vacuum chamber and trap (green), modified from [39]. The main $^{43}Ca^{+}$ & $^{88}Sr^{+}$ Doppler cooling beams enter superimposed from the bottom right, nearly counter-propagating with the 674 quadrupole beam. The 397_{vert} enters $\sim 10°$ above the horizontal for micromotion compensation. The quantisation axis B is aligned along the 397σ, at a small angle to the $850\sigma_{1,2}$ beams. The calcium and strontium ovens are shown in orange and purple respectively, with atomic emission to the trap centre travelling almost perpendicular to the photo-ionisation (PI) beam. The R_H and R_V co-propagating beams form a Raman pair with the $R_{\|}$ beam, with their k-vector difference lying along the trap axis.

To reduce our sensitivity to beam-pointing drifts, light from each laser's AOM network is sent to the vacuum system along lengths of polarisation-maintaining single-mode fibre, which decouples the alignment of the laser and trap side of each beam. After the fibre a PBS cube turns polarisation variations into power variations, with care taken to align one of the stress-induced axes of the fibre with the polarisation vector of the light launched into it.

A sample of the beam is reflected onto a photodiode used to implement an FPGA-based intensity stabilisation system, with the RF drive power sent to AOMs before the fibre controlled to keep the amount of light illuminating the ions constant [32]. If a pair of beams are sent down the same fibre, intensity stabilisation is available only when one beam is on. After a pair of mirrors for alignment onto the ion, the beam is focussed with a final lens, whose focal length along with that of the fibre out-coupler chosen to set the desired beam waist at the trap centre.

Figure 4.7 shows the laser beam geometry on the experiment table, relative to the quantisation axis. 393 and 397σ light share a fibre to the experiment table, with the magnetic field direction aligned onto the latter using the trim coils described in section 4.2.2. The PI and $850\sigma_{1,2}$ enter from the same direction at a small angle. Through an unused window an extra 397_{load} beam detuned 400 MHz to the red with a large spot size is used to cool hot $^{43}\mathrm{Ca}^{+}$ ions in large orbits during loading, or after collisions from background gas molecules. The 397π, 866, 854 and 850π calcium beams are superimposed with the $422_{u,l}$, 1092 and 1033 strontium beams on a stack of long- and short-pass dichroic mirrors[33], and all pass through a 250 mm achromatic doublet lens[34] to the trap. The 1092 and 1033 beams share a fibre, as do the 866 and 854. The 422 beam consists of upper and lower frequency components derived from the same laser, described in section 4.4.2. To measure and compensate the ions' micromotion in the vertical direction, a flip mirror sends the 397π to a periscope which gives a small vertical projection to the 397_{vert} beam $\sim 10°$ above the horizontal.

[33] Thorlabs DMSP950 (950nm short-pass) and DMLP650 (650nm long-pass), Edmund Optics 86-323 (409nm long-pass)

[34] Thorlabs AC254-250-A. The lens is anti-reflection coated for the blue wavelengths, as the power requirements and problems caused by scatter are much less critical for the red and IR beams.

The spot size at the ion is a compromise between intensity, beam-pointing stability and scatter off the trap electrodes for each beam. Since we filter out light above 440nm wavelength in the imaging system, scatter is less of an issue for the red and IR beams. The blue beams (except the 397_{load}) have $1/e^2$ intensity waist diameters at the trap centre of $\sim$60 μm; the red beams have $\sim$100 μm.

4.5 Raman Lasers

To produce the pairs of beams needed for Raman transitions we use two phase-locked frequency-doubled SolsTis Ti:Sapphire lasers from M Squared Lasers Ltd. Their details and the associated AOM network and beam delivery are detailed in Vera Schäfer's thesis [39], but a brief description follows here. Each laser produces $\sim$1–2 W of blue light detuned from atomic resonances near 400nm, which after passing through an AOM-switching network and short lengths of fibre produce 3 beams incident into the trap (see figure 4.7). A beam from each laser co-propagating perpendicular to the magnetic field, the R_V with vertical and the R_H with horizontal polarisation, are used to drive single qubit gates. To bridge the 3.2 GHz qubit transition in $^{43}\text{Ca}^+$, a sample of the red Ti:Sapphire outputs before the doubling stage from each laser interfere on a photodiode. A 1.6 GHz local oscillator mixes down this beat note to produce an error signal which is fed back onto one of the lasers, known as the 'Slave'. Feedback on a piezo and EOM internal to the laser cavity suppress phase noise below and above $\sim$50 kHz respectively, to lock its frequency relative to the 'Master'. A second horizontally polarised beam from the 'Master' along the magnetic field direction, the $R_{\|}$, along with the R_V from the same laser have a k-vector difference along the trap axis. This beam pair is used to perform two-qubit entangling gates (with AOMs providing the frequency differences needed), with the $R_{\|}$ and R_H pair driving axial motional mode sidebands. The $R_{\|}$ polarisation is set very carefully using a $\lambda/4$ waveplate, so it produces no single beam light-shift on $^{43}\text{Ca}^+$, counter-acting the small birefringence of the vacuum system windows. All three beams are carefully focussed to a 50 μm diameter spot at the ion.

4.6 Quadrupole Laser

To drive the $^{88}Sr^{+}$ optical qubit, readout the state of the $^{88}Sr^{+}$ Zeeman qubit and conduct gate operations between two $^{88}Sr^{+}$ ions, we need a laser to drive the quadrupole $S_{1/2}$-$D_{5/2}$ transition. As this has a narrow natural linewidth of 2.5 Hz, and we hope to use it for coherent operations on the qubit, we need a laser with good frequency stability and good short term linewidth, to avoid decreasing the coherence time that can be achieved. A high power output is also necessary, as while tens of microwatts are needed to drive the dipole transitions, tens of milliwatts are required to achieve Rabi frequencies of more than a MHz on the quadrupole transition.

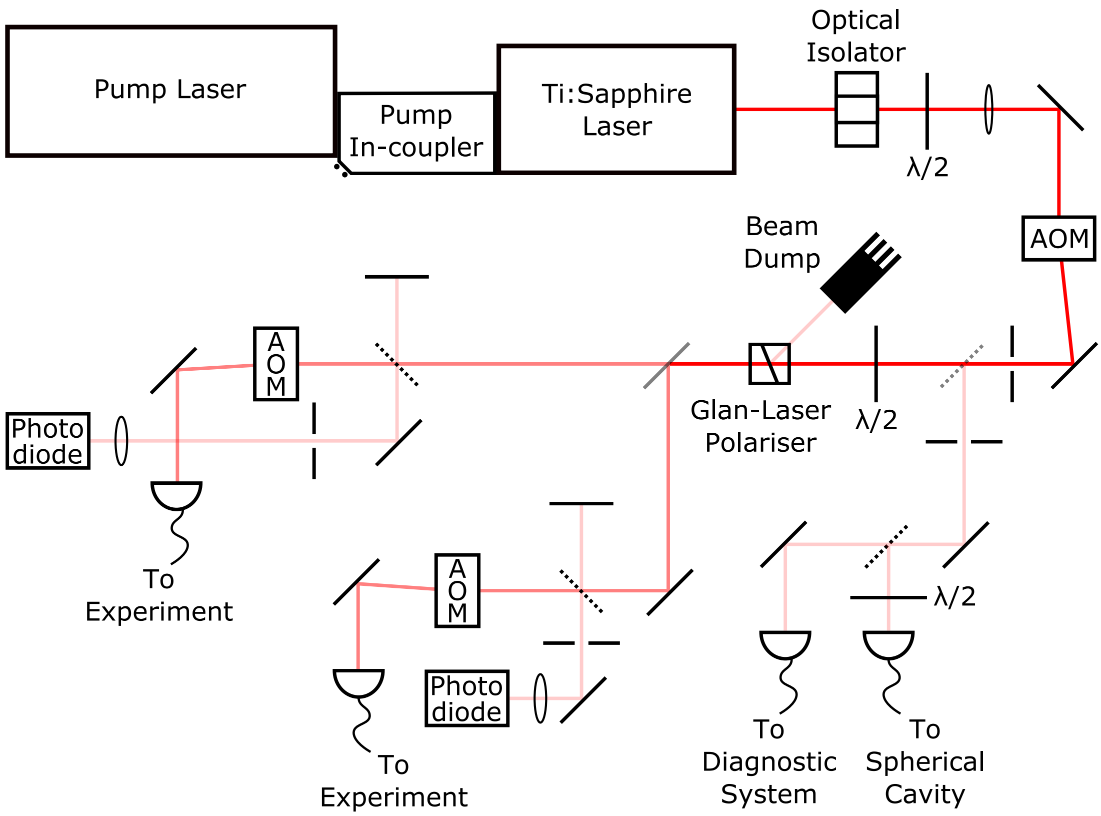

Figure 4.8: Block diagram of the quadrupole laser system. A Ti:Sapphire laser outputting >1 W at 674nm is locked to a spherical reference cavity. The beam is split into two arms, each coupled into a fibre to an experiment. A fibre noise cancellation system beats a sample of the beam with a retro-reflection from the end of the fibre onto a photodiode, feeding back on an AOM before the fibre to suppress phase noise to 5 mrad RMS.

For this purpose we use another SolsTis Ti:Sapphire laser, manufactured by M Squared Lasers Ltd., locked to a high-finesse spherical reference cavity system provided by Stable Laser Systems. The laser, pumped by 15 W of green light at 532nm by a diode pumped solid state (DPSS) laser[35], outputs over 1 W at 674nm. A sample of the light is sent to the diagnostic system, and to a Ultra-Low Expansion (ULE®) glass sphere, with mirror coatings of specified finesse 300, 000 on either side forming a cavity ∼50 mm long. When locked to the cavity the laser drifts are almost exclusively linear, at 17.2 kHz per day (figure 4.9), with the lock robust to laboratory noise (such as slamming doors).

[35] Lighthouse Photonics Sprout-G series

The latter is achieved due to the cavity spacers' spherical design, where the optical axis is at a 'squeeze insensitive angle' from the support axis [84]. The cavity is temperature controlled in two stages, with the inner temperature at the cavity expansion zero-crossing temperature of 29.67°C and the outer close to room temperature. The whole assembly is inside a vacuum chamber at 3×10^{-8} Torr mounted on a breadboard on an active vibration isolation stage. A 5 MHz EOM applies sidebands to the light before it enters the cavity, and is used to form a PDH lock which feeds back on the Ti:Sapphire internal cavity piezos, and on a free space 80 MHz AOM after the laser, for slow and fast corrections respectively. Since the cavity and atomic resonance are in general different, an offset-locking scheme is used to bridge the 3 GHz free spectral range between cavity modes. A synthesiser[36] generates sidebands on the PDH signal which are used for the lock, and can be adjusted to tune the laser frequency by ±1.5 GHz. The linewidth of the system as measured in the factory is 1.5 Hz over 1s.

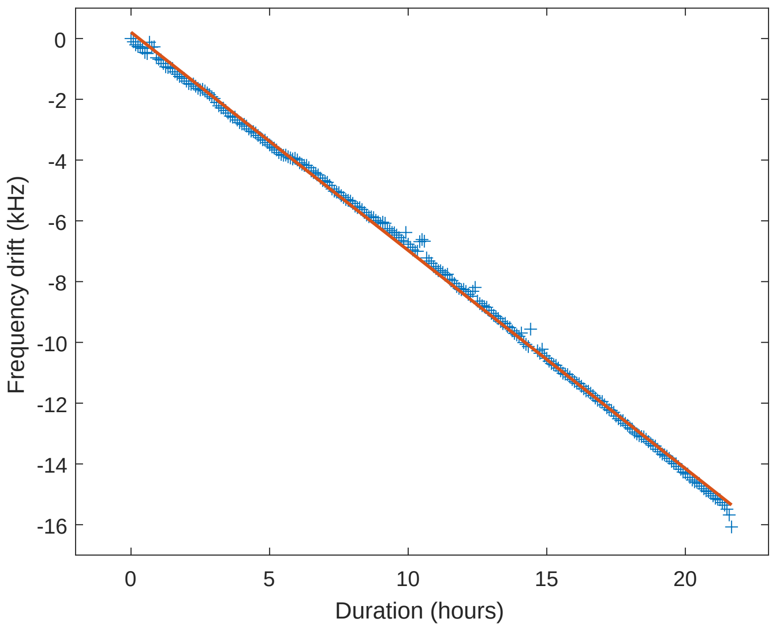

Figure 4.9: Quadrupole laser frequency drift as measured by Ramsey interferometry on a $^{88}Sr^{+}$ ion. Relaxation of machining stresses in the ULE® glass locking cavity produce an almost linear drift of 720 Hz per hour (red line).

[36]Holzworth HSM2001A

The laser is housed in a neighbouring lab to the experiment table, with the light sent down a 15m polarisation-maintaining single-mode fibre routed along a ceiling mounted cable tray. To ensure acoustic noise coupled into the fibre does not unacceptably reduce the beam's temporal coherence, an active fibre phase noise cancellation system (FNCS) is used. Light passing through an 80 MHz AOM before the fibre is retro-reflected by its far flat polished end facet and superimposed with a sample of the incoming beam onto a photodiode. The 160 MHz beat frequency between the two is used to generate an error signal which is fed back to the AOM, reducing the phase noise to 5 mrad RMS. We use an optical isolator[37] on the experiment table directly after the fibre to prevent light scattered back from components further down the path interfering with the FNCS signal.

An 80 MHz AOM[38] in double-pass configuration is used for switching and frequency scanning of the beam. This is followed by a single pass 200 MHz AOM[39], driven either with two frequencies for the Mølmer-Sørensen bichromatic tone, or at its centre frequency for other operations, which is then imaged into a short fibre to the trap.

As shown in figure 4.7, the 674nm beam enters the vacuum chamber almost perpendicular to the magnetic field, with a good projection onto the trap axis to drive axial motional sidebands. In this geometry transitions between Zeeman states of the $5S_{1/2}$ and $4D_{5/2}$ levels that have the least magnetic field sensitivity, that do not change the magnetic quantum number ($\Delta m_J = 0$), are suppressed. Instead a ($\Delta m_J = -1$) transition between $S_{1/2}\,|m_J = -1/2\rangle$ and $D_{5/2}\,|m_J = -3/2\rangle$ is driven. Figure 4.10 shows a frequency scan around this transition, where the axial and radial modes sidebands can be clearly seen. The beam is focussed to a $\sim$60 μm waist diameter at the ion.

[37] Thorlabs IO-5-670-VLP
[38] Isomet 1205C-2-804B
[39] Isomet 1250C-848

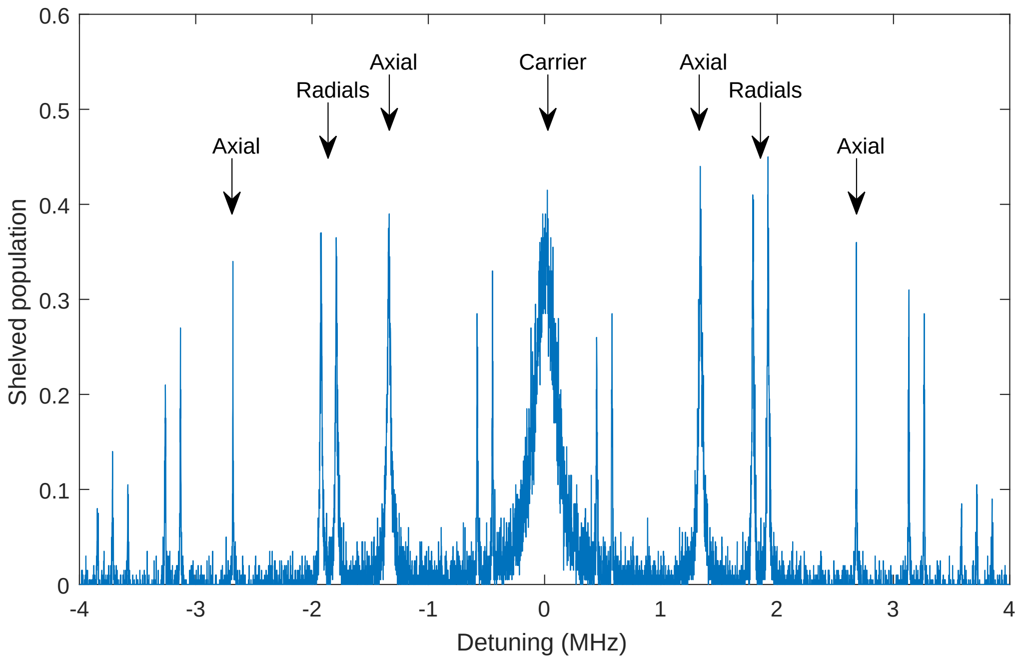

Figure 4.10: Frequency scan around a component of the quadrupole transition (carrier) showing multiple sidebands at the axial and two radial normal mode oscillation frequencies, and their mixing products, for a hot $^{88}Sr^{+}$ ion in the trap. The axial and mean radial trap frequencies are 1.34 MHz and 1.86 MHz respectively.

5

Two-Qubit Gates

In this chapter I discuss experiments performing two-qubit entangling gates on mixed and same species ion pairs. Using a novel technique developed here, we use the similar atomic structure and proximity of transitions in $^{43}\mathrm{Ca}^+$ and $^{88}\mathrm{Sr}^+$ to carry out a light-shift gate on the two ions, using beams derived from a single laser, in a scheme technically simpler than other methods. This work is currently being prepared for publication. In a precursor to this experiment, and as a test of our new 674 nm laser, I also describe the preliminary results of a Mølmer-Sørensen gate carried out on two $^{88}\mathrm{Sr}^+$ ions.

5.1 $^{43}Ca^+$– $^{88}Sr^+$ Light-Shift Gate

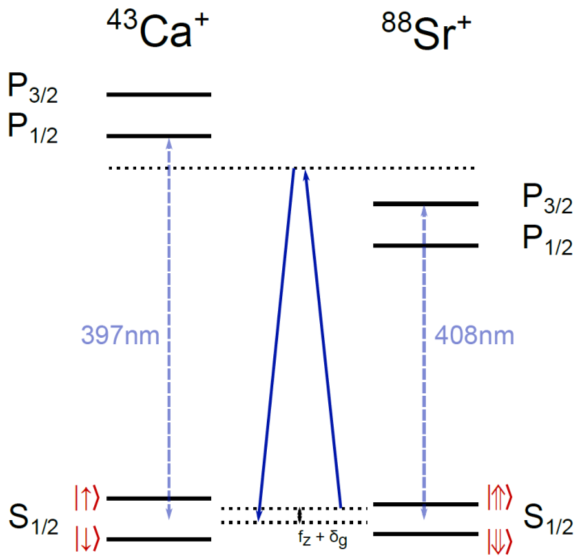

Figure 5.1: Frequency configuration of the gate laser beams for the mixed species gate. A single pair of beams derived from the same laser has a Raman coupling to the S-P transitions in both $^{43}Ca^+$ and $^{88}Sr^+$. Tuned between the 397 nm and 408 nm transitions (which are separated by 20 THz) the beams create a light-shift force off-resonant to the ions motional frequency, thereby driving a geometric phase gate.

In this section I describe the implementation, optimisation and results from an entangling gate between a $^{43}Ca^+$ and $^{88}Sr^+$ ion. Though the two species have different transition wavelengths, and hence for other operations require different sets of lasers, we carry out a gate operation using a pair of Raman beams derived from a single laser. This is possible because both species have transitions to an upper level separated by only 20.23 THz, namely the 397 nm $S_{1/2}$ - $P_{1/2}$ transition in $^{43}Ca^+$ and the 408 nm $S_{1/2}$ - $P_{3/2}$ transition in $^{88}Sr^+$ (figure 5.1); with the Raman beam pair tuned between these two transitions we can create an appreciable two-photon coupling to both. The difference frequency of the beams is set slightly detuned from the ions' motional frequency, off-resonantly driving the mode and causing the collective spin states of the ions to make loops in position-momentum phase space and acquire a geometric phase. With the correct geometric phase difference between the even and odd spin states we implement a $\sigma_z \otimes \sigma_z$ light-shift gate. The gate mechanism, which is independent of the frequency of either qubit, is described in more detail in section 3.5.4.

Needing only a single laser for both species this mechanism is technically simpler than alternatives, such as a Mølmer-Sørensen gate, which need individual beams near-resonant with each species' qubit frequency. This involves a significant increase in complexity and phase control of beams from different lasers. A disadvantage of our mechanism is that it cannot practically be applied on 'clock' qubits. This is because the differential light-shift we rely on for our gate tends to zero in the limit of large Raman detunings for magnetically insensitive pairs of qubit states [60].

5.1.1 Raman Beam Geometry

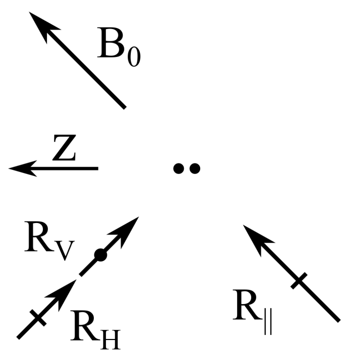

Figure 5.2: Arrangement of Raman beams, reproduced from [32]. The R_V and $R_\parallel$ beams are derived from the same laser and drive the gate in a 'lin perp lin' configuration, with the latter parallel to the static magnetic field B_0. Their difference k-vector is parallel to the trap axis Z. The R_H (from a laser 3.2 GHz higher in frequency) along with the $R_\parallel$, is used for sideband cooling.

To implement our gate, we require the Raman beam pair to have a significant projection onto the mode we wish to drive phase space loops in. Additionally, it is advantageous to choose a geometry which limits the number of modes 'seen' by the beam pair, to limit off-resonant excitation errors (section 3.6.2). The geometry and polarisations we use are summarised in figure 5.2, with more context found in figure 4.7. The R_V and $R_\parallel$ beams, derived from the 'Master' Raman laser, are used to drive the gate. They have perpendicular polarisations and their difference k-vector points along the axial mode direction of the trap, in a 'lin perp lin' configuration. The beams interfere forming a spatially varying polarisation pattern, creating a spatially varying light-shift, dependent on the collective spin state of the ions, i.e. a spin dependent force. The polarisations are chosen to maximise this gate force and minimise single beam light-shifts on both qubits[1]. The R_H, derived from the 'Slave' Raman laser phase locked 3.2 GHz away from the 'Master', along with the $R_\parallel$ drives motional sidebands around the $^{43}\mathrm{Ca}^+$ qubit resonance for sideband cooling and temperature diagnostics. In either case the radial motional modes are orthogonal to the difference k-vector and therefore have no coupling to the beams.

5.1.2 Raman Detuning

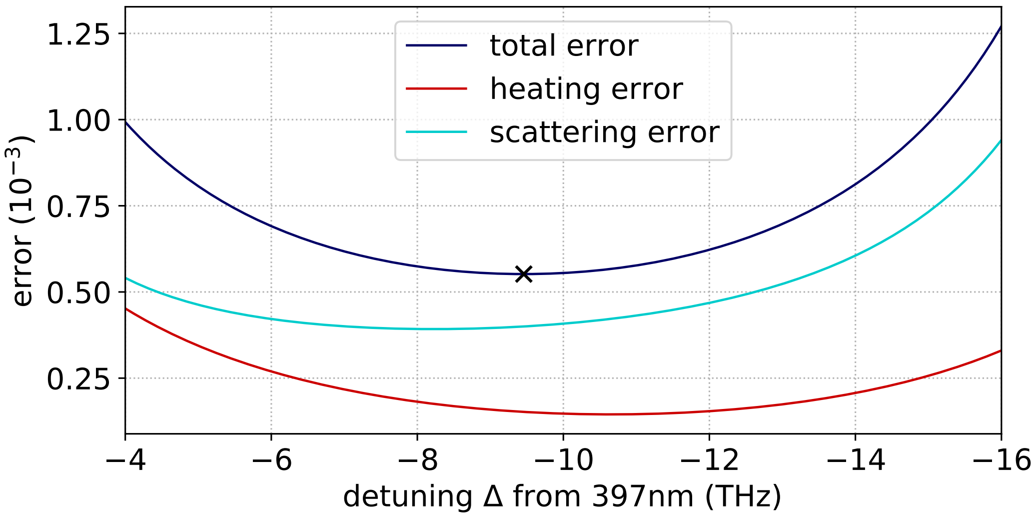

Figure 5.3: Calculated error sources for the gate on the out-of-phase mode versus Raman detuning, between the resonances of $^{43}\mathrm{Ca}^+$ and $^{88}\mathrm{Sr}^+$ at 0 THz and −20.23 THz respectively. Raman detunings near the resonance of either species allow faster gates, but at the expense of higher scattering and heating errors. The point of minimum total error (black x) is close to the $\Delta = -9$ THz detuning used for the gate. Reproduced from [39], which assumes slightly different experimental parameters to this work, though errors should agree within a factor ∼2.

[1]The added hyperfine splittings in $^{43}\mathrm{Ca}^+$ mean the polarisations required to null its light-shift are slightly different to those for $^{88}\mathrm{Sr}^+$. This effect is small and is corrected by the spin-echo sequence, so we set the polarisations to null the light-shift on $^{43}\mathrm{Ca}^+$.

The choice of Raman detuning between the transitions in $^{43}Ca^+$ and $^{88}Sr^+$ (figure 5.1), to minimise the gate error, is a trade-off between Rabi frequency and photon scattering from each species. Tuning close into one species applies a large force allowing for a shorter gate time, but gives large photon scattering errors and only a small force on the other species. The size of the phase space loops for each spin state are now similar, so to acquire the required $\pi/2$ geometric phase difference between the odd and even states we need to drive larger loops. These larger excursions in phase space exponentially increase the error due to mode heating, as explained in section 3.6.1.

A calculation of these error sources versus Raman beam detuning for a fixed laser power is shown in figure 5.3. While tuning close into one species reduces the gate time, and so the duration over which the crystal heats, the increase in heating error due to the large phase space loops required is much larger. Along with photon scattering error, the total error is minimised by tuning somewhere halfway between the resonances of each species, but is fairly insensitive to the exact value.

We sit at $\Delta = -9\,\mathrm{THz}$ from the 397 nm transition in $^{43}Ca^+$, at a wavelength of approximately 402 nm. At this detuning, with the Lamb-Dicke parameters listed in table 5.1, the ratio of the geometric phase acquired by the odd states to that acquired by the even states is roughly $9:1$, for gates on either mode. This corresponds to a gate efficiency of $\Theta = 80\%$ (see equation 3.50).

5.1.3 Experimental Sequence

Each experiment consists of a basic sequence of ground-state cooling, state preparation, optional pre-gate pulses, the spin-echo gate sequence, optional post-gate pulses, state-selective shelving and readout. We use the Zeeman qubit in $^{88}Sr^+$ and the stretch qubit in $^{43}Ca^+$, and simultaneously state-prepare to the $|{\downarrow}\rangle$ state of both. The pre- and post-gate pulses are simultaneous single qubit pulses on each ion used for tomography, to change the input state before the gate and the readout basis after, respectively. All single qubit pulses are applied using RF/microwaves resonant with $^{88}Sr^+$ and $^{43}Ca^+$ respectively.

For each setting of the experiment (i.e. particular pre- and post-gate pulses) this sequence is run N times, and we record whether the ions were bright or dark each time. Combining these shots we find the probabilities $\{p_{\uparrow\uparrow}, p_{\uparrow\downarrow}, p_{\downarrow\uparrow}, p_{\downarrow\downarrow}\}$, where the first index denotes whether $^{43}Ca^+$, and the second whether $^{88}Sr^+$, was bright ($\uparrow$) or dark ($\downarrow$). We check the fluorescence of the ions to see if they have decrystallised after each shot of the experiment, and recrystallise and re-order if necessary. We run the B field servo every 10 minutes, interleaving them between long data runs.

Cooling

Each experiment begins with Doppler cooling of the two ion crystal. We Doppler cool for 4 ms in total, with the first 2/3 of the time spent simultaneously Doppler cooling on $^{88}Sr^+$ and dark resonance cooling on $^{43}Ca^+$, to cool the modes on which they respectively have a large eigenvector. This is the only direct cooling on $^{88}Sr^+$; everything that follows cools the $^{43}Ca^+$ ion which then sympathetically cools the $^{88}Sr^+$. Since their masses are fairly similar, so too are their eigenvectors in each mode, so this works well. The last 1/3 of the Doppler cooling time is spent only dark resonance cooling on $^{43}Ca^+$, as this achieves a better final Doppler temperature. This is then followed by 2 ms continuous sideband cooling and 20 pulses of pulsed sideband cooling on both axial modes, using the R_H and $R_\parallel$ Raman beams. These parameters were set empirically by optimising on the mode temperature and then the gate fidelity. We reach a temperature of $\overline{n} = 0.075(5)$ on the in-phase mode and $\overline{n} = 0.030(5)$ on the out-of-phase mode.

5.1.4 Gate Sequence

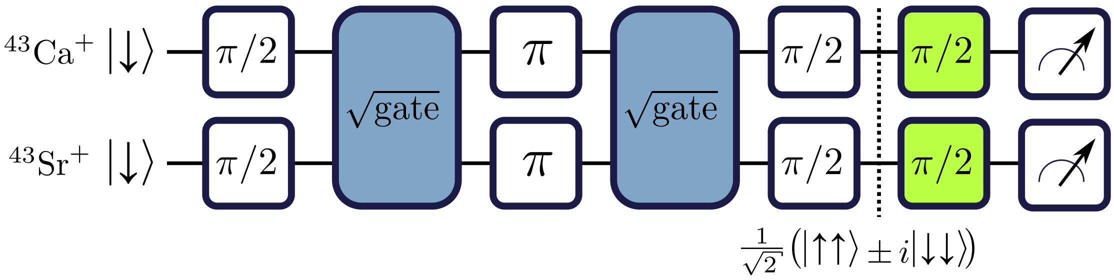

Figure 5.4: Sequence of pulses which make up a gate experiment. The light-shift gate pulses are applied inside both arms of a spin-echo sequence, each moving the spin states through complete loops in phase space acquiring half the total geometric phase. The single qubit pulses are driven by microwaves and RF resonant with each species, with the optional analysis pulses (green) mapping out a parity fringe as their phase is scanned.

We apply the gate inside a spin-echo sequence as shown in figure 5.4, so that in each arm we move the ions through a complete loop in phase space accumulating half the required total geometric phase. This arrangement serves a number of purposes. Firstly, the initial $\pi/2$ pulse rotates the ions into a superposition of all four z-basis two-qubit states ($|\Psi\rangle = \left(|\uparrow\uparrow\rangle + |\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle + |\downarrow\downarrow\rangle\right)/2$), a necessary pre-requisite before applying our z-basis gate operation. The spin-echo also cancels any residual single beam light-shifts from the Raman beams, reduces errors due to dephasing of our magnetically sensitive qubits, and symmetrises the gate operation against the different force magnitudes on ^{43}Ca$^+$ and ^{88}Sr$^+$, due to their dissimilar Lamb-Dicke parameters or even unequal illumination.

Force Phase

While the absolute phase of the gate force relative to the ions' motion does not affect the gate dynamics, their phase difference in each arm of the spin-echo sequence sets the relative angle between the paths they push the spin states through in phase space. In principle this phase difference can be non-zero and does not need to be controlled, as long as each phase space loop is perfectly closed. In practice this is often not the case, for example when scanning the gate pulse duration (as in figure 5.7) where the dynamics can become unintelligible as the spin states take uncontrolled trajectories.

To fix this, we first note that an oscillator at the Raman difference frequency (for example driving the AOM) will deviate from the motional mode frequency by phase δt if the gate pulse starts at $t = 0$. To make the gate force continuous over both pulses we therefore need to shift the phase of the second pulse to remove the phase accumulated during the time the gate beams are off. Accounting for the time taken for the single qubit π-pulse with padding, along with the π phase shift to the force added by the pulse itself, this is $\phi \to \phi - \delta(t_\pi + t_{\mathrm{pad}}) - \pi$.

Making the force phase-continuous however will cause the spin states to return to the phase space origin at half the gate pulse time $t_{\mathrm{p}}/2$, as the two half completed loops sum. Since during optimisation of the gate it is convenient for the dynamics to appear the same for any number of loops, we instead ‘reset’ the phase by updating it by $\phi \to \phi - \delta(t_{\mathrm{p}} + t_\pi + t_{\mathrm{pad}}) - \pi$. Now both pulses start at the same phase with respect to the motion, but the end of the first pulse is not at the same phase as the start of the second, i.e. they are discontinuous.

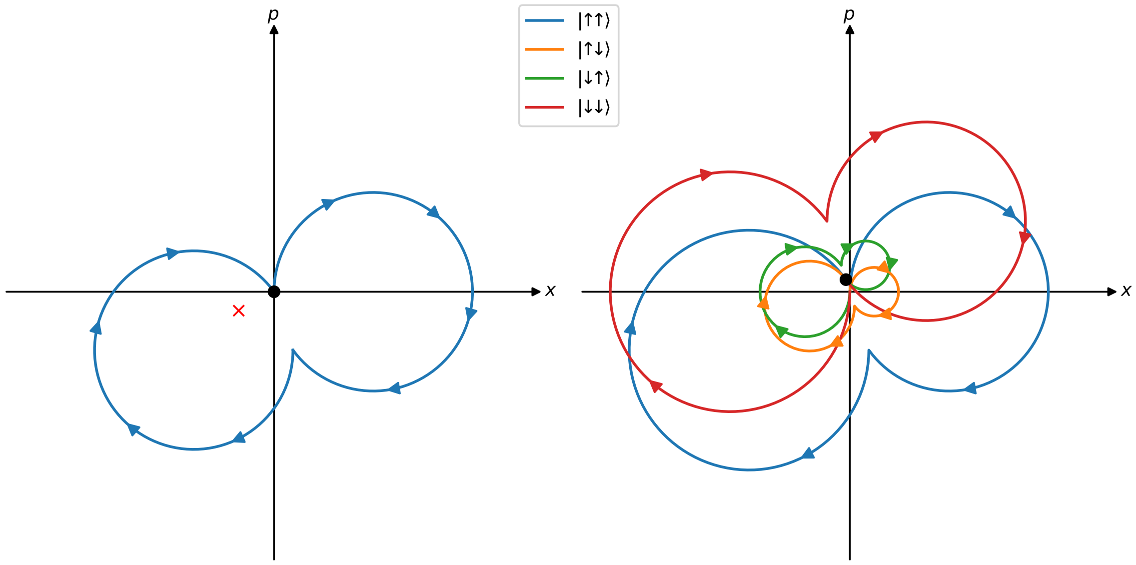

Figure 5.5: Phase space trajectories during a two-loop gate in a Walsh modulated 'reset phase' arrangement (see text), with detuning or timing errors. The black dots mark the end of the second loop. When all four spin states have equal force magnitudes (Left), the trajectories close for all gate times (only one state trajectory shown). In the general mixed species case (Right), the different force magnitudes cause the trajectories to meet, but not close. For comparison, in a Walsh modulated 'continuous phase' arrangement, the trajectories do not meet or close (red x).

Lastly, it is advantageous to add a further π phase shift, or equivalently not correct for the shift caused by the single qubit π-pulse, to reduce our sensitivity to imperfect loop closure caused by detuning or gate duration errors. This reversal in gate phase is a first order Walsh modulation, bringing the spin states closer to the origin at the end of the gate. Whilst in [85] this is carried out on otherwise phase-continuous loops, we apply it to discontinuous 'reset phase' loops as described above. As can be seen in figure 5.5, we now gain an added advantage: if all four spin states have equal force magnitude (for example in a same species gate without hyperfine structure, see section 3.5.4) the total path is now closed for all pulse durations, making the gate insensitive to loop closure errors. This is equivalent to the identity $(e^{-i\delta t} - 1) - (e^{-i\delta t} - 1) = 0$.

If the spin states have different force magnitudes, as is the case for mixed species gates, the trajectories do not individually close, but all meet at a point offset from the phase space origin. By the end of the gate each state has been displaced from the origin by $(\Omega_{\uparrow,1} + \Omega_{\downarrow,1} + \Omega_{\uparrow,2} + \Omega_{\downarrow,2})\,(e^{-i\delta t} - 1)$, which while non-zero in the presence of detuning or timing errors, is the same for all four states (c.f. section 3.5.4). As reduced phase space overlap between the spin states, rather than displacement from the origin, causes errors, the gate is still robust against detuning or timing errors.

Pulse Shaping

We pulse shape the gate pulses to reduce off-resonant excitation errors, as described in section 3.6.2. We apply a pulse shape with shape function $\sin^2(\pi t/2t_{\text{shape}})$ to the R_V beam choosing $t_{\text{shape}} = 2\,\mu$s, which as calculated in [32] is more than sufficient to suppress errors for similar gate parameters. The R_V pulse is bracketed by the unshaped $R_{\|}$ pulse (i.e. the $R_{\|}$ is switched on before and switched off after the R_V), so any timing jitter does not cause pulse area errors.

5.1.5 Setting Ion Spacing

As described in section 3.5.4, the ions need to be spaced close to $d = n\lambda_z/2$ apart in the moving standing wave, or ‘walking wave’, created by the interfering Raman beams for an optimal gate, where n is an integer and λ_z is the instantaneous wavelength of the walking wave. At other spacings the gate efficiency reduces, meaning we have to increase the total geometric phase by driving larger loops in phase space and so increase errors due to heating.

To set the ion spacing we make use of the observation that, for each of the 4 collective spin states in the z-basis, there are spacings where the forces applied to the mode cancel, as long as the forces on each ion are equal in magnitude. At these spacings, which are the same as those optimal for the gate, applying even a resonant gate force does not move the spin state from the phase space origin. We can therefore resonantly excite the mode and look for a minimum in motional excitation while scanning the ion spacing.

This procedure requires equal force magnitudes on each ion, so at the correct ion spacing they are equal and opposite and so cancel. Due to the unequal magnitudes of the Lamb-Dicke parameters (including normal mode vectors) and generally unequal Rabi frequencies on each ion in a mixed species pair, this is not straightforward. Equal force magnitudes can be arranged by adjusting the Rabi frequency on each ion, in our case by a careful choice of Raman detuning between the $^{43}Ca^{+}$ and $^{88}Sr^{+}$ resonances. This however is technically unfeasible, as scanning the Raman detuning would require repeated time-consuming re-alignments of the 'Master' Raman laser doubling cavity.

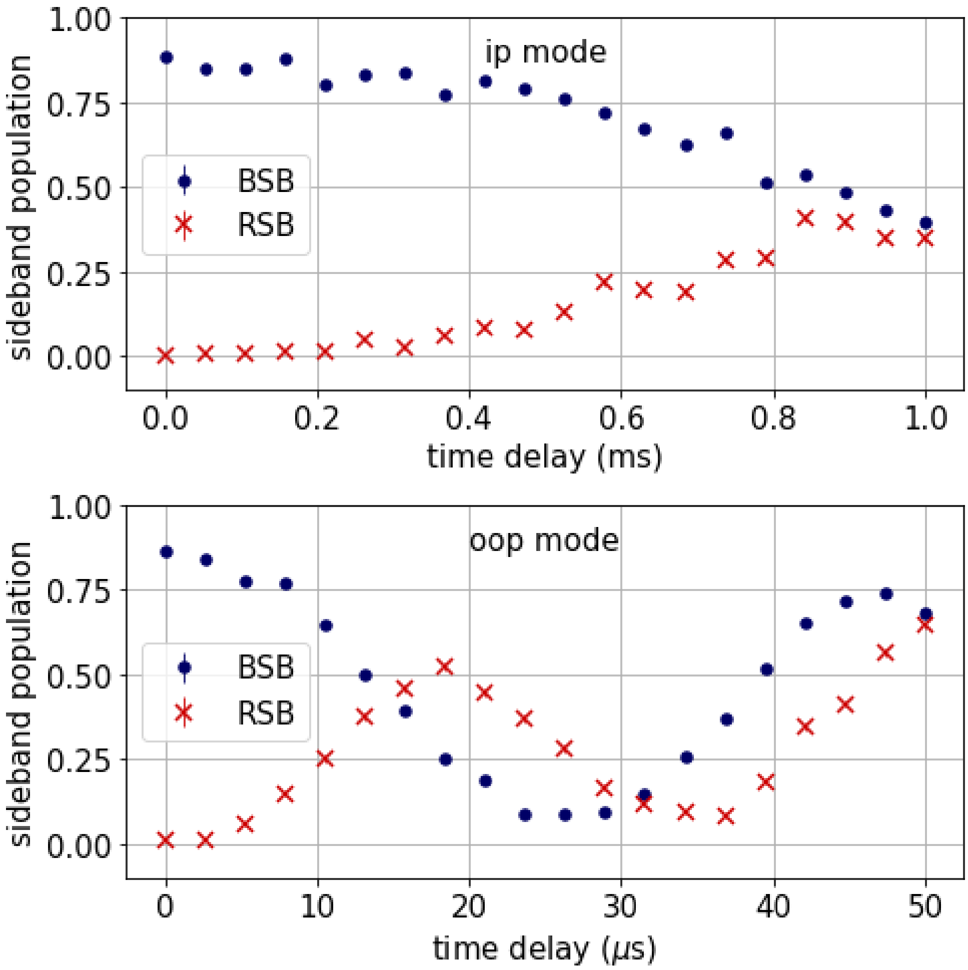

Figure 5.6: Resonant mode excitation of a near ground state cooled $^{43}Ca^+$-$^{43}Ca^+$ crystal at the ion spacing used for the gate; note the difference in scale of the excitation duration of the two plots. The ions are spaced 12.5 standing wave periods apart so that the force on the in-phase mode cancels while that on the out-of-phase mode is maximised. Due to the near invariance of the ion spacing with mass, the same spacing is achieved for our $^{43}Ca^+$-$^{88}Sr^+$ crystal.

We instead rely on the fact that the ion spacing is almost mass independent, as shown in equation 2.7, and set it using a same species ion pair, which trivially gives us equal force magnitudes. Loading a pair of $^{43}Ca^+$ ions and preparing them in the $|\downarrow\downarrow\rangle$ state, we apply gate beams with difference frequency resonant with either axial mode and then look for motional excitation. Scanning the DC endcap voltages to adjust the spacing we can find a point with zero (or very little) excitation in one of the modes.

We find this point to be at $d = 12.5\lambda_z$, with a scan of the sideband excitation versus the force duration shown in figure 5.6.

5.1.6 Axial Tilt

During experiments preparing this gate we noticed that we could excite the radial modes with the Raman lasers. This was an issue as the radial rocking modes and their harmonics were very close to the axial modes on which we drive the gate, and so would cause large off-resonant excitation errors. As the Raman lasers have been aligned carefully so their difference k-vector lies parallel to the trap axis, this should not be the case. Indeed with two $^{43}Ca^{+}$ ions we do not see the radial modes. We therefore established that the $^{43}Ca^{+}$-$^{88}Sr^{+}$ crystal was tilted with respect to the trap axis.

This was due to imperfect compensation of the trap in the radial direction. As the radial confining force from the oscillating RF field is mass dependent (see section 2.4.2), but the force from stray fields are not, residual stray fields will cause ions of differing mass to have differing radial equilibrium positions. In our case the heavier $^{88}Sr^{+}$ ion is more weakly confined by the RF than $^{43}Ca^{+}$, and so sat further away from the RF null line (trap axis), causing a tilt in the crystal. This was remedied by carefully compensating in the radial direction to cancel the stray fields and return the ions to the RF null line, after which we could no longer see the radial modes with the Raman lasers [41].

5.1.7 Implementation

	In-phase mode	Out-of-phase mode
Mode frequency	1.494 MHz	2.906 MHz
Gate detuning	80 kHz	−40 kHz
Gate duration	25 μs	50 μs
Raman beam power	84 mW	60 mW
$^{43}Ca^{+}$ Lamb-Dicke parameter	0.090	0.127
$^{88}Sr^{+}$ Lamb-Dicke parameter	0.124	0.045

Table 5.1: Approximate parameters for the mixed species gates, actual parameters were fined tuned on gate dynamics to improve fidelity. The frequency difference between the two Raman beams used to apply the gate pulse is the sum of the mode frequency and gate detuning. The power given is per beam, each focussed to a 50 μm diameter spot at the ion. Gate duration is the sum for both loops (excluding single qubit pulses). The Lamb-Dicke parameters are calculated from experiment parameters, and include the mode eigenvector magnitudes (given in section 2.4.1), i.e. they are the absolute value of the quantity defined in equation 3.29.

We performed gates on both the in-phase and out-of-phase axial modes of the $^{43}Ca^{+}$-$^{88}Sr^{+}$ crystal aiming for the Bell states $|\Phi\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\uparrow\rangle \pm i\,|\downarrow\downarrow\rangle\right)$ respectively, with the sign of the Bell state matching the sign of the gate detuning. Approximate parameters of the gate on each mode are given in table 5.1. To arrive at these parameters it was often useful to look at the gate dynamics, by scanning the length of the gate beam pulses while looking at the two-ion populations. An example for gates on both modes is shown in figure 5.7. We calibrate the gate duration by looking for the population minimum of the odd states. Here, the correct gate power gives equal population for the even states at 50%. At odd and even multiples of the gate time respectively, we can see the ions entangle and disentangle, where any errors in the parameters appear amplified as systematic offsets. The dynamics show reasonably good agreement with the simulation, showing we have good understanding and control of the mixed species gate mechanism.

Readout error normalisation

We are interested in the quality of our gate operation over its many instances during a quantum computation sequence. It is therefore useful to separate intrinsic gate errors, which compound with sequence length, from state preparation and readout (SPAM) errors, which only occur at the beginning and end. When assessing our gate we therefore wish to calculate the true spin state populations from the measured bright and dark fluorescence probabilities. These are connected for a single ion by the confusion matrix M, in the relation

$$\begin{pmatrix} p_{\mathrm{bright}} \\ p_{\mathrm{dark}} \end{pmatrix} = \begin{pmatrix} 1-\epsilon_\uparrow & \epsilon_\downarrow \\ \epsilon_\uparrow & 1-\epsilon_\downarrow \end{pmatrix} \begin{pmatrix} p_\uparrow \\ p_\downarrow \end{pmatrix}, \tag{5.1}$$

where $\epsilon_\uparrow$ is the probability an ion we intend after state preparation to be in $|\uparrow\rangle$ is measured as bright, and vice versa for $\epsilon_\downarrow$. These SPAM errors are measured with independent experiments close in time to the gate experiment, and are fairly stable. To calculate the true state populations we simply apply M^{-1} to the vector of observed fluorescence probabilities. For two ions where we can distinguish which fluoresces we use the tensor product of the confusion matrix for each species, and so in our case apply $(M_{^{43}\mathrm{Ca}^+} \otimes M_{^{88}\mathrm{Sr}^+})^{-1}$ to correct for SPAM errors.

	$^{43}Ca^+$	$^{88}Sr^+$
$\epsilon_\uparrow$	0.22(4)%	0.24(4)%
$\epsilon_\downarrow$	0.14(3)%	0.58(6)%

Table 5.2: Representative state-preparation and readout error measurement. Each qubit state of each ion was prepared and readout 20,000 times, with $\epsilon_\uparrow$ the percentage of times an ion we attempted to prepare in $|\uparrow\rangle$ was erroneously measured as dark, and conversely for $\epsilon_\downarrow$. The average error for all four states is $\bar{\epsilon} = 0.29(2)\%$.

Our SPAM errors are given in table 5.2. The largest contribution, from the $^{88}Sr^+$ $|\downarrow\rangle$ state, is mostly due to imperfect shelving on this species with the 674 nm laser. As we do not cool the radial modes of the crystal their thermal distribution incoherently suppress the carrier (see section 3.6.1), lowering the quality of the π shelving pulse.

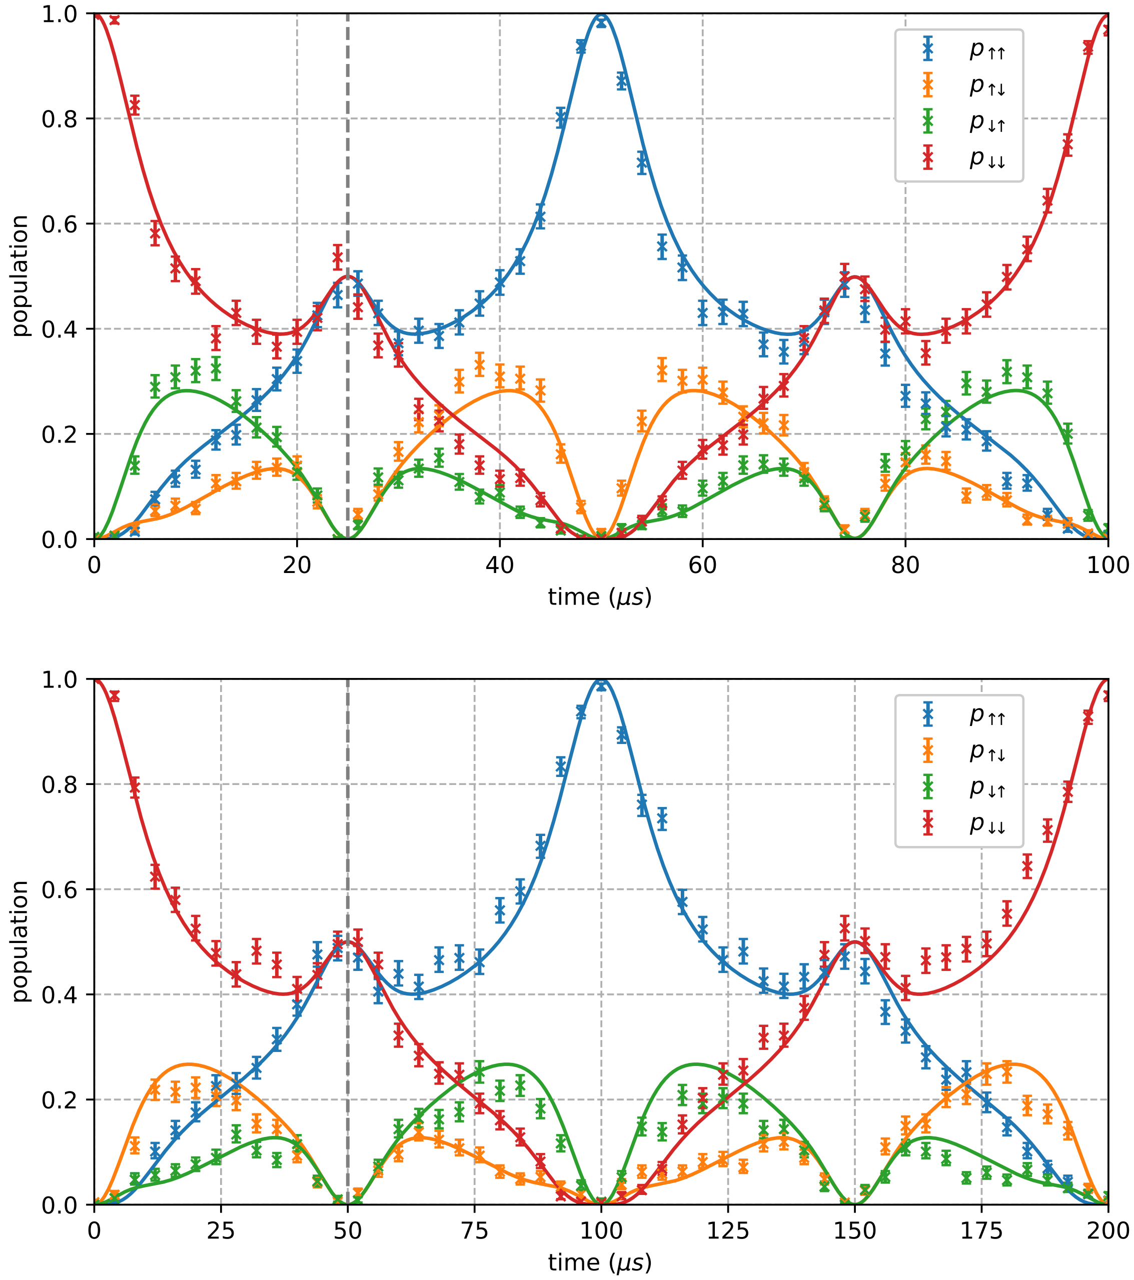

Figure 5.7: Gate dynamics on the in-phase (Top) and out-of-phase (Bottom) modes, versus total gate pulse duration, after readout error normalisation. Starting in $|\downarrow\downarrow\rangle$ the gate force initially excites the motion of the ions, entangling it with their collective spin state. By the gate time (grey dashed line) the motion is disentangled, and the ions are left in the maximally entangled states $|\Psi\rangle = (|\uparrow\uparrow\rangle \pm i\,|\downarrow\downarrow\rangle)\,/\sqrt{2}$ respectively. At twice the gate time the ions are disentangled again but have had their spin flipped. The length of each spin-echo arm was kept constant and Walsh modulation was not used. Solid lines are simulations of the gate including anomalous heating, written by Vera Schäfer, and are not fitted.

5.1.8 Bell State Tomography

The standard and simplest method for measuring the fidelity of a gate is to find the overlap of the output state with the target state, when applied to a separable input state [21,56]. For gates on the in-phase and out-of-phase modes we intend to create the specific states $|\Phi\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\uparrow\rangle \pm i\,|\downarrow\downarrow\rangle\right)$ respectively, but it is easier and more meaningful to quantify the fidelity with respect to the general Bell state with arbitrary phase θ

$$|\Phi_\theta\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\uparrow\rangle + e^{i\theta}\,|\downarrow\downarrow\rangle\right)\,, \tag{5.2}$$

as we can trivially rotate between them with a global single qubit rotation [62]. For an output state with density matrix ρ we can quantify the fidelity as[2]

$$\begin{aligned} F = \langle\Phi_\theta|\rho|\Phi_\theta\rangle &= \frac{1}{2}\left(\rho_{\uparrow\uparrow,\uparrow\uparrow} + \rho_{\downarrow\downarrow,\downarrow\downarrow}\right) + \frac{1}{2}\left(e^{i\theta}\rho_{\uparrow\uparrow,\downarrow\downarrow} + e^{-i\theta}\rho_{\downarrow\downarrow,\uparrow\uparrow}\right) \\ &= \frac{1}{2}\left(\rho_{\uparrow\uparrow,\uparrow\uparrow} + \rho_{\downarrow\downarrow,\downarrow\downarrow}\right) + |\rho_{\uparrow\uparrow,\downarrow\downarrow}|\,. \end{aligned} \tag{5.3}$$

If the output state is pure i.e. $\rho = |\Psi\rangle\langle\Psi|$ the fidelity $F = |\langle\Phi_\theta|\Psi\rangle|^2$. In the second line we have set θ to maximise the fidelity, and so best fit the phase of the real state.

To measure the fidelity we perform Bell state tomography, by first measuring the spin state populations after the gate, with the sum of the symmetric populations $p_{\uparrow\uparrow} + p_{\downarrow\downarrow} = \rho_{\uparrow\uparrow,\uparrow\uparrow} + \rho_{\downarrow\downarrow,\downarrow\downarrow}$. To find the off-diagonal element $|\rho_{\uparrow\uparrow,\downarrow\downarrow}|$ we then need to apply a rotation $R(\pi/2, \phi)$ to both ions after the gate, called the 'analysis pulse'. If we scan its phase ϕ we can obtain from the populations the parity signal

$$\begin{aligned} P(\phi) &= p_{\uparrow\uparrow} + p_{\downarrow\downarrow} - (p_{\uparrow\downarrow} + p_{\downarrow\uparrow}) \\ &= \rho_{\uparrow\downarrow,\downarrow\uparrow} + \rho_{\downarrow\uparrow,\uparrow\downarrow} + 2|\rho_{\uparrow\uparrow,\downarrow\downarrow}|\cos(2\phi + \phi_0)\,, \end{aligned} \tag{5.4}$$

where ϕ_0 is an offset which depends (among other things) on the output state phase.

[2] The square root of this quantity is also sometimes confusingly called the fidelity [86].

The contrast of these 'parity fringes' with periodicity 2ϕ can then be measured to find $|\rho_{\uparrow\uparrow,\downarrow\downarrow}|$.

Examples of such a fidelity measurement are shown in figure 5.8. While the probability of either ion in isolation being bright is 50%, they are strongly correlated and almost always appear bright together, the hallmark of an entangled state.

To find the parity contrast the most obvious method is to fit a sinusoid to the data with fixed period but floating offset phase and contrast. The usual least-squares fit however is not a good choice as it assumes a normal distribution, when in our case the probabilities are binomially distributed. This systematic bias is especially bad for points near $p = 0$ or $p = 1$, exactly those that the fitted contrast is most sensitive to for high fidelity gates. Whilst a maximum likelihood fitting method assuming binomial statistics can solve this problem [32], there is still a danger of insufficiently sampled data causing the fit to slightly over-estimate the fidelity by trading-off between the contrast and offset phase.

Instead, the parity contrast can be directly calculated from four measurements of the parity signal at analysis phases separated by $\pi/4$. With some basic properties of sinusoids it is straightforward to show that

$$\sqrt{[P(0) - P(\pi/2)]^2 + [P(\pi/4) - P(3\pi/4)]^2} = 4|\rho_{\uparrow\uparrow,\downarrow\downarrow}| \tag{5.5}$$

which does not depend on the offset phase ϕ_0. If the offset phase is stable we can do one better. Finding it beforehand with a detailed scan, we can measure at only two points, the peak and trough at $\phi = \{-\phi_0/2, \pi/2 - \phi_0/2\}$, their difference being twice the contrast. This is equivalent to finding the overlap fidelity with a particular, rather than general, Bell state, which is in some sense better. These methods have the advantage of focussing data collection on points which have the largest impact on the measured contrast, as well as preventing over-estimation of the fidelity. Examples of fidelity measurements are shown in figure 5.8.

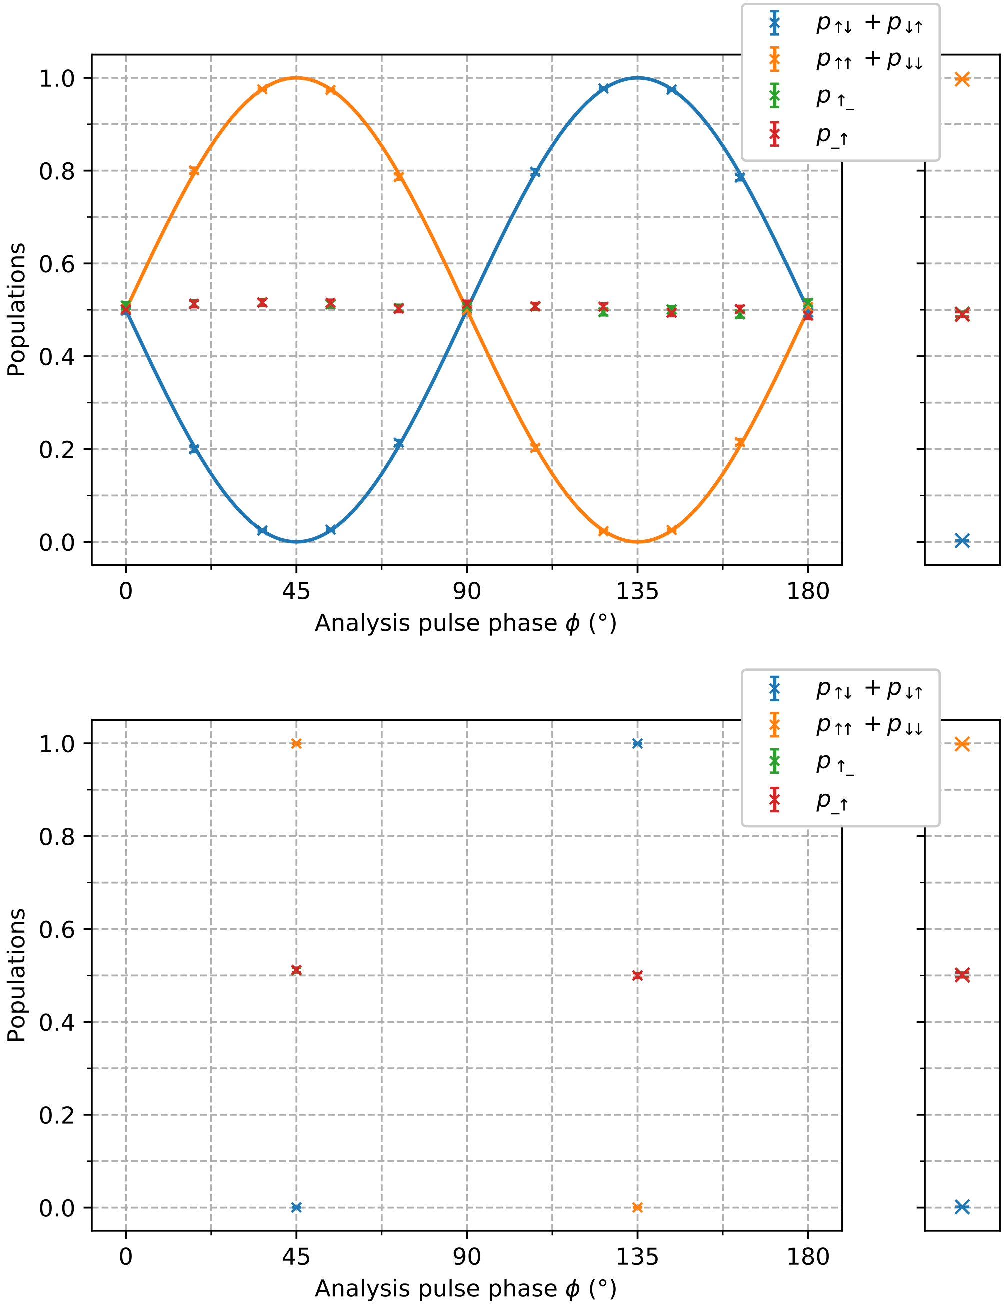

Figure 5.8: Bell state fidelity measurements after readout error normalisation. After the gate a $\pi/2$ analysis pulse is applied with a scanned phase ϕ (Left), yielding the parity signal with 2ϕ periodicity. Its contrast, along with the populations without the analysis pulse (Right), are used to calculate the fidelity. A detailed scan, shown on the in-phase mode gate (Top), can be fitted with a sinusoid, but measuring only at points of maximum contrast, shown on the out-of-phase mode (Bottom), prevents over-estimation of the fidelity. Green and red points correspond respectively to $^{43}Ca^+$ or $^{88}Sr^+$ being bright, irrespective of the other.

Despite this we often measure fidelities which, while consistent with (in the sense that their lower error bar overlaps) 1, have upper error bars and mean values above 1. Figure 5.9 shows multiple fidelity measurements of gates on both modes taken over two days. The errorbars are calculated using standard error propagation of the confidence intervals of the binomially distributed bright and dark measurement probabilities. We find $F = 99.89(12)\%$ for the gate on the in-phase mode and $F = 99.80(10)\%$ on the out-of-phase mode. Most measurements were taken with $N = 10,000$ shots per point, though some had $N = 1,000$.

The unphysical fidelities occur for both systematic and statistical reasons. Firstly, as our SPAM errors are larger than our gate errors, small changes in beam pointing or ion temperature can cause the former to be measured slightly higher than their true values during the gate. Readout error normalisation then over-corrects the raw data, yielding artificially high fidelities. Additionally, perfectly stable and positive raw gate and SPAM errors when measured can subtract to give negative readout normalised gate errors, as each has independent statistical measurement fluctuations which can only be reduced by taking higher statistics. This problem is compounded by the erroneous assumption that both are described by normal, rather than binomial, statistics.

Teasing apart these two sources of error to reliably measure the gate fidelity unfortunately requires unreasonably high statistics which are difficult to achieve. The measurements in figure 5.9 were generally taken with $N = 10,000$ shots per point, though some had $N = 1,000$. The points with higher statistics took several minutes to run, so trying to significantly lower the uncertainties which scale as $\sqrt{N}$, would take so long that in practice systematic errors due drifts in the experiment would swamp any statistical gains. We must therefore use other methods to pin down the gate fidelity with confidence.

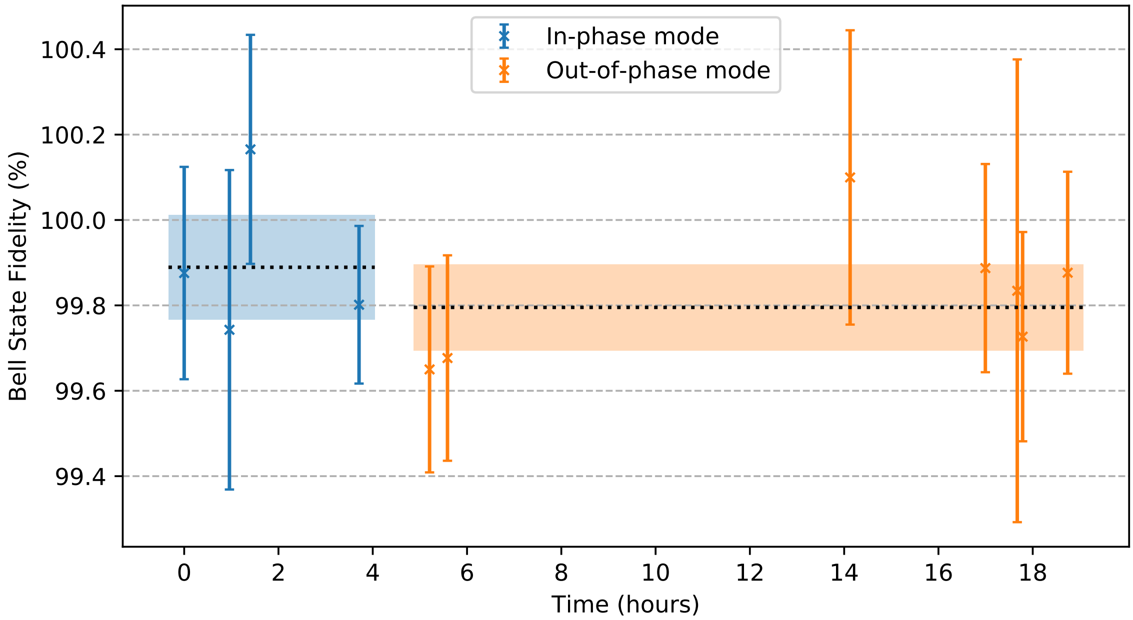

Figure 5.9: Bell state fidelity measurements taken on consecutive days. For each, the parity contrast was found using two or four point parity scans, and corrected with readout errors measured near contemporaneously. The dotted line and shading show the mean fidelity, weighted by the reciprocal variance of each point, and its uncertainty, for gates on both modes. These are $F = 99.89(12)\%$ for the in-phase mode and $F = 99.80(10)\%$ for the out-of-phase mode.

5.1.9 Full State Tomography

Whilst measuring the Bell state fidelity is a form of partial tomography, it doesn't measure the output state over a complete set of bases, but only the overlap with a particular state. While we can quantify the degree to which the state is, and is not, the Bell state, it doesn't tell us *which* state it is.

To fully characterise the two-qubit state after the gate we need to find its density matrix. Using a finite number of measurements we can reconstruct it by carrying out full state tomography [87, 88]. The density matrix ρ can be expanded as

$$\rho = \sum_{i}^{4^N} \lambda_i O_i \,, \tag{5.6}$$

into a set of mutually orthogonal Hermitian operators O_i, where N is the number of qubits. For two qubits, the natural choice for this set are the 16 tensor products of the Pauli matrices including the identity $O_i = \sigma_j \otimes \sigma_k$, with $\sigma_i = \{I, X, Y, Z\}$, known as the 'Pauli basis'. To find the coefficients λ_i and reconstruct the density matrix we need to measure the expectation values of these operators

$$\lambda_i = \mathrm{tr}[\rho O_i]/4\,, \tag{5.7}$$

with the divisor their Hilbert-Schmidt inner product $\mathrm{tr}[O_i O_j] = 4\delta_{ij}$. When we make a measurement in the experiment, we intrinsically measure in the $Z \otimes Z$-basis, and the probabilities $\{p_{\uparrow\uparrow}, p_{\uparrow\downarrow}, p_{\downarrow\uparrow}, p_{\downarrow\downarrow}\}$ we obtain are the diagonal elements of the density matrix. From this we can find the expectation value of the $Z \otimes Z$ operator, the $Z \otimes I$ and $I \otimes Z$ operators by taking the partial trace over one qubit, and $I \otimes I$. For the rest we need to apply single qubit rotations to the ions before measurement to transform the measurement basis from $Z \otimes Z$ to the operator in question. To measure the expectation value of $X \otimes Y$ for example, we need to rotate the ions with $U_{XY} = R_Y(\pi/2) \otimes R_X(-\pi/2)$ such that

$$\mathrm{tr}\Big[(U_{XY}\rho U_{XY}^\dagger)(Z\otimes Z)\Big] = \mathrm{tr}\Big[\rho(U_{XY}^\dagger(Z\otimes Z)U_{XY})\Big] = \mathrm{tr}[\rho(X\otimes Y)]\,, \tag{5.8}$$

where we can also find the expectation values of $X \otimes I$, $I \otimes Y$ and $I \otimes I$ for free [89]. In this way we need only measure at 9 different settings of the experiment to find all 16 expectation values.

Calculating the density matrix from these observations by inverting the linear equation 5.6, while straightforward, often fails to produce a physical density matrix, i.e. one that is Hermitian, with unit trace and whose eigenvalues are all real and non-negative [90]. The most common method around this is to do the opposite, and instead find from the space of physical density matrices the one most likely to have generated the observed data; Maximum Likelihood Estimation [91].

Following the procedure in [88, 92], the trial density matrix $\rho'(\mathbf{t})$ is parametrised with 16 real numbers $\mathbf{t} = \{t_1, t_2, \ldots, t_{16}\}$ in such a way as to ensure physicality. We then define the negative-log-likelihood function

$$\mathcal{L}(\mathbf{t}) = \sum_i \frac{(\lambda_i - \lambda_i')^2}{1 - \lambda'^2} \tag{5.9}$$

where λ_i are the observed, and $\lambda_i' = \mathrm{tr}[\rho'(\mathbf{t})O_i]/4$ are the trial expectation values, which quantifies how well $\rho'(\mathbf{t})$ fits the measured data [89]. If we numerically minimise $\mathcal{L}(\mathbf{t})$ we can find the optimal values for $\mathbf{t}$ and so the best estimate for the density matrix. As an initial guess to start the optimisation, we take the density matrix produced by linear inversion and use the 'Quick and Dirty' method described in [92] as a shortcut to make it physical, namely decomposing it into its eigenvalues and eigenvectors, setting any negative eigenvalues to zero and renormalising. The density matrix of the output state of our gate, as found by MLE is shown in figure 5.10.

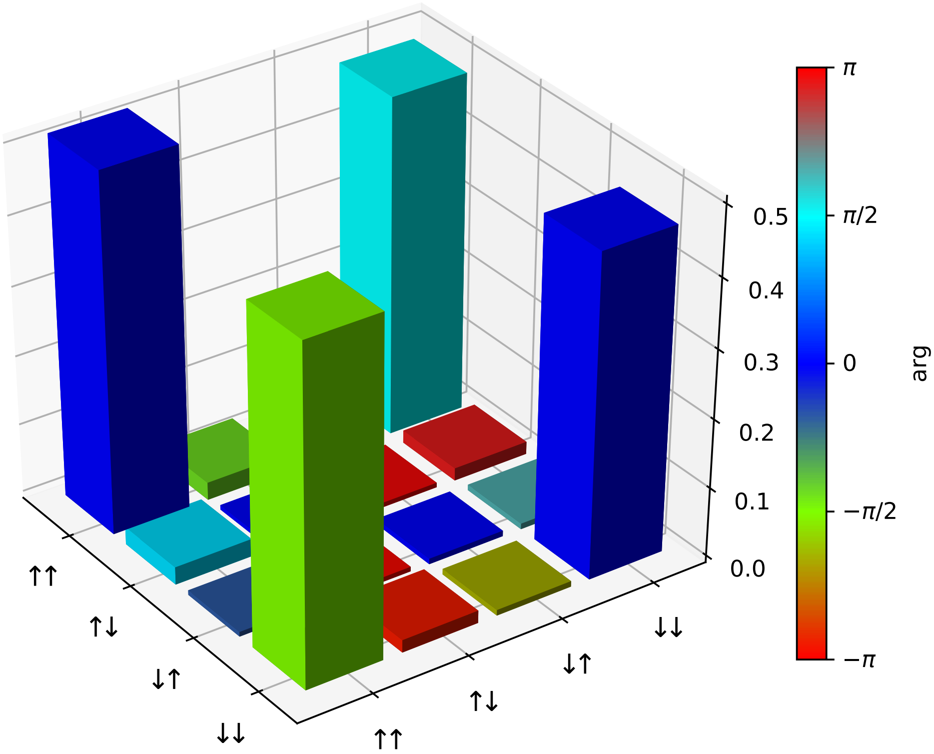

Figure 5.10: Reconstructed density matrix of the output of the gate on the out-of-phase mode, where the spin states are in the order $\{^{43}Ca^+, ^{88}Sr^+\}$. After readout error correction of the full tomography measurements, maximum likelihood estimation (MLE) was used to find the density matrix which best fit the data.

The value of reconstructing the density matrix arises when trying to understand the sources of error in the gate. For example a probabilistic process which produces a perfect Bell state most of the time, can have the same Bell state fidelity as one which produces an imperfect Bell state every time. To distinguish between these two extremes we can quantify from the density matrix measures of stochasticity and entanglement [93].

Quantifying the degree of order in the system is the linear entropy

$$S_L = \frac{4}{3}(1 - \mathrm{tr}\left[\rho^2\right]) , \tag{5.10}$$

which ranges from $S_L = 0$ for a pure state to $S_L = 1$ for a maximally mixed state (where

$\rho = I^{\otimes N}/2^N$). There are multiple different measures used to quantify the degree of entanglement [94], here we use tangle (concurrence squared) which for a general mixed state is

$$\tau = C^2 = [\max(0, \lambda_1 - \lambda_2 - \lambda_3 - \lambda_4)]^2 \,, \tag{5.11}$$

where λ_i are the eigenvalues in decreasing order of the matrix $R = \sqrt{\sqrt{\rho}\rho_{\mathrm{sf}}\sqrt{\rho}}$, and the spin-flipped state is

$$\rho_{\mathrm{sf}} = (Y \otimes Y)\rho^*(Y \otimes Y) \,, \tag{5.12}$$

where the complex conjugate is taken with ρ in the standard basis $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$ [95]. $\tau = 0$ for an unentangled separable state and ranges up to $\tau = 1$ for a maximally entangled state.

Whilst these measures describe different qualities of the system they are not unconnected, as an impure state must have a reduced amount of entanglement; indeed states which have the maximum amount of entanglement for a given purity, maximally entangled mixed states, form a boundary between physical and unphysical states on the entropy-tangle plane [96]. Our gate is plotted on this plane in figure 5.11, and its Bell state fidelity, linear entropy and tangle are $F = 98.1\%$, $S_L = 0.045$, $\tau = 0.932$. It lies on the line of Werner states, defined as a mixture between the maximally entangled and maximally mixed states

$$W(p) = p\,|\Phi\rangle\langle\Phi| + (1-p)(I \otimes I)/4 \,, \tag{5.13}$$

where p is the probability of producing $|\Phi\rangle$, a Bell state. It would thus be reasonable to identify that the infidelity in our gate is consistent with probabilistic processes which introduce decoherence (i.e. mixing), rather than stable miscalibration of the gate geometric phase or single qubit pulses.

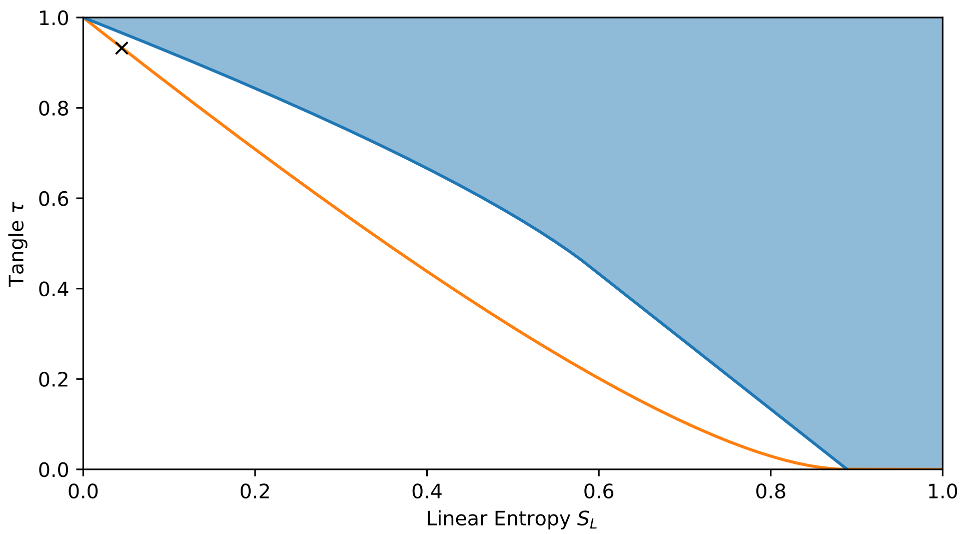

Figure 5.11: Properties of reconstructed density matrix (black x) on entropy-tangle plane. The blue line corresponds to maximally entangled mixed states [96], above which no physical state can exist (shaded region). The orange line corresponds to the Werner states, on which our state lies. A perfect Bell state would be in the top left corner.

5.1.10 Process Tomography

While state tomography assesses the output of the gate on a *particular* input state, in the context of a computational sequence we need to understand how it acts on an *arbitrary* input state. Thankfully we don't need to test the gate on all infinite input states, but instead on a linearly independent basis set from which we can work out where any arbitrary state will be mapped, in a procedure called process tomography [93, 97–99].

The gate, its errors, and any decoherence from the environment can be described by the completely positive trace preserving map $\mathcal{E}$, which transforms the input density matrix ρ to the output $\widetilde{\rho}$. For N qubits it can be decomposed into a set of operators E_m, which form a basis of the space of $2^N \times 2^N$ matrices, as

$$\mathcal{E}(\rho) = \sum_{m,n}^{4^N} \chi_{mn} E_m \rho E_n^\dagger = \widetilde{\rho}\,, \tag{5.14}$$

where χ is the process matrix. The basis of operators and the process matrix together completely describe the process.

To find χ, we start by preparing the system into each of the linearly independent input density matrices ρ_i, and measure after application of the gate the output $\mathcal{E}(\rho_i)$ using state tomography. For two qubits this requires 16 input states needing measurement in 9 bases each, resulting in 144 different settings of the experiment. Rather than the 'Pauli basis' as in the previous section, it is experimentally more convenient to choose for ρ_i the tensor products of the states $\{|z_+\rangle\langle z_+|, |z_-\rangle\langle z_-|, |x_+\rangle\langle x_+|, |y_+\rangle\langle y_+|\}$[3] (henceforth the 'input basis'), as they require for our two-qubit case at most a π or $\pi/2$ pulse to reach from our initially prepared $|z_+\rangle \otimes |z_+\rangle$ state. If we expand each output in this basis we get

$$\mathcal{E}(\rho_i) = \sum_j \lambda_{ij}\rho_j \,, \tag{5.15}$$

where λ is a matrix of observations. If we then write the action of the operators on each of the input states as $E_m \rho_i E_n^\dagger = \sum_j \beta_{ij}^{mn}\rho_j$, where β is fully specified by our choice of bases, we can express the process on each state as

$$\mathcal{E}(\rho_i) = \sum_{j,m,n} \chi_{mn}\beta_{ij}^{mn}\rho_j \,. \tag{5.16}$$

Comparing equations 5.15 and 5.16 we can form a matrix equation by combining the indices ij and mn, which flattens λ and χ into vectors and reshapes β in our case to a 256×256 matrix. We can solve this equation to find the process matrix as

$$\chi = \beta^{-1}\lambda \,. \tag{5.17}$$

Most often we wish to describe the process matrix in the more general 'Pauli Basis', especially as its expectation values are what we measure with state tomography. Choosing these operators for E_m we need to appropriately convert from the 'input

[3]Where $|z_+\rangle = |\uparrow\rangle = \left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$, $|z_-\rangle = \left(\begin{smallmatrix}0\\1\end{smallmatrix}\right)$, $|x_+\rangle = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix}1\\1\end{smallmatrix}\right)$ and $|y_+\rangle = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix}1\\i\end{smallmatrix}\right)$.

basis' during the above calculations.

The gate on the out-of-phase mode was measured during 9 consecutive process tomography experiments, each complete with 144 measurement points, of 500 shots each. Each experiment took approximately 21 minutes to run, with all except one directly preceded by a B field servo, for a total of around 3h 10 mins. As with state tomography in the previous section, linearly inverting these observations can often produce a result which is not physical. We therefore again used maximum likelihood estimation (MLE) to find the process most likely to have produced the combined process tomography data [100–102]. The result of this procedure, with MLE code written by David Nadlinger, is shown in figure 5.12.

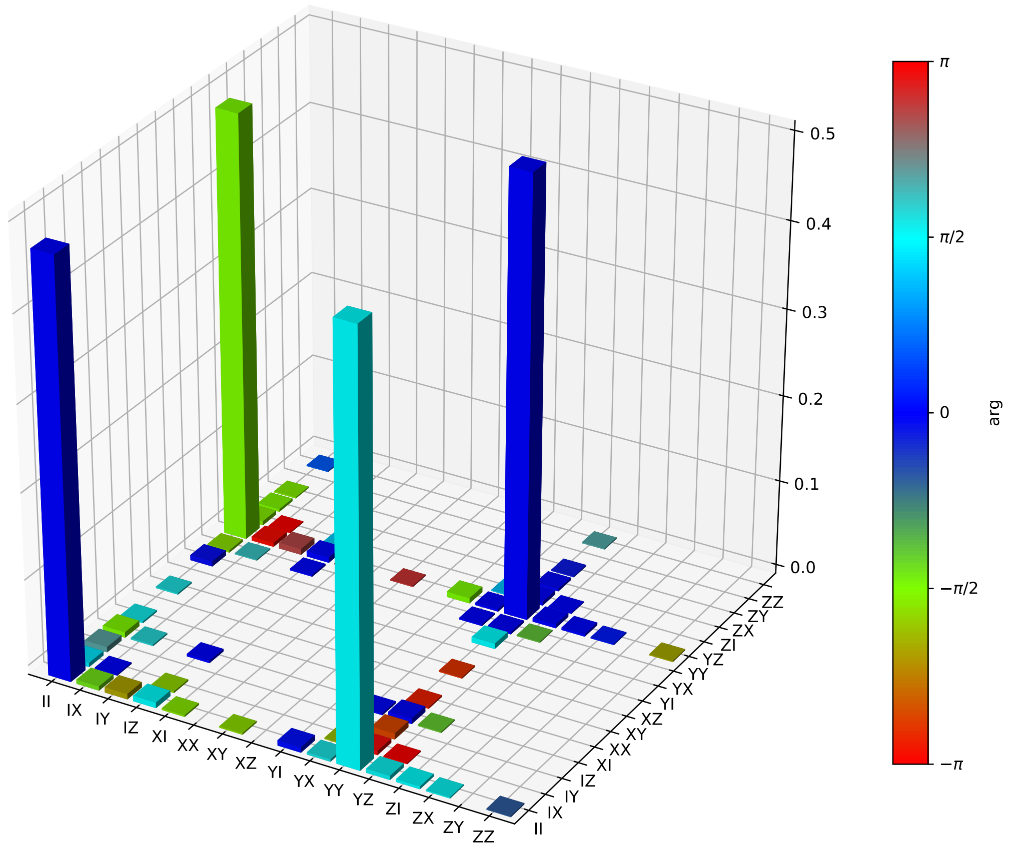

Figure 5.12: Reconstructed process matrix χ of the out-of-phase mode gate, from maximum likelihood estimation (MLE) of the data after readout error normalisation. It is expressed in the basis of tensor products of the identity and Pauli operators, with any entries of magnitude below 0.001 omitted for clarity.

Up to a global phase, the unitary describing the ideal gate is given by

$$
\begin{aligned}
G &= \begin{pmatrix} 1 & & & \\ & \pm i & & \\ & & \pm i & \\ & & & 1 \end{pmatrix} \\
U &= R_{XX}\left(\frac{\pi}{2}\right) \cdot \sqrt{G} \cdot R_{XX}(\pi) \cdot \sqrt{G} \cdot R_{XX}\left(\frac{\pi}{2}\right) \\
&= \frac{1}{\sqrt{2}} \left(I \otimes I \mp iY \otimes Y\right) ,
\end{aligned} \tag{5.18}
$$

where $R_{XX}(\theta) = R_X(\theta) \otimes R_X(\theta)$ and the sign of the imaginary elements of G match that of the gate detuning. We therefore expect the non-zero elements of the process matrix to be

$$\begin{aligned}\chi_{II,II} &= 1/2\\ \chi_{II,YY} &= \pm i/2\\ \chi_{YY,II} &= \mp i/2\\ \chi_{YY,YY} &= 1\end{aligned} \tag{5.19}$$

with the tensor products omitted for clarity, which is indeed what we see. Apart from verifying that we understand the gate process well, the process matrix can be used to characterise the quality of our gate. We first find the process fidelity $F_P = \mathrm{tr}[\chi_{\mathrm{ideal}}\chi_{\mathrm{meas}}]$, which is the overlap between the process measured from tomography and the ideal process. For our gate $F_P = 98.81(7)\%$. A more useful measure is the average gate fidelity $\overline{F}$, which is the state fidelity (as in equation 5.3) of the measured output states of the gate with respect to the ideal output states, averaged over all input states. This has a simple relation to the process fidelity as

$$\overline{F} = \frac{dF_P + 1}{d+1}\,, \tag{5.20}$$

where $d = 2^N$ is the dimension of the Hilbert space [86, 103]. For our gate $\overline{F} = 99.05(6)\%$.

The uncertainties on these numbers are obtained by bootstrapping, where we create an ensemble of simulated process tomography datasets, and find for each the maximum likelihood estimate for χ and then calculate the fidelities. The mean and standard deviation of these fidelities then give the values and uncertainties above. We use non-parametric bootstrapping, which generates the simulated data by resampling from the original observations with replacement, or equivalently sampling from a multinomial distribution with the probabilities of the original data. This method makes no assumptions about the quantum process, unlike parametric bootstrapping, which uses the original maximum likelihood estimated process to generate the simulated data.

The histogram of average gate fidelities generated by this method is shown in figure 5.13.

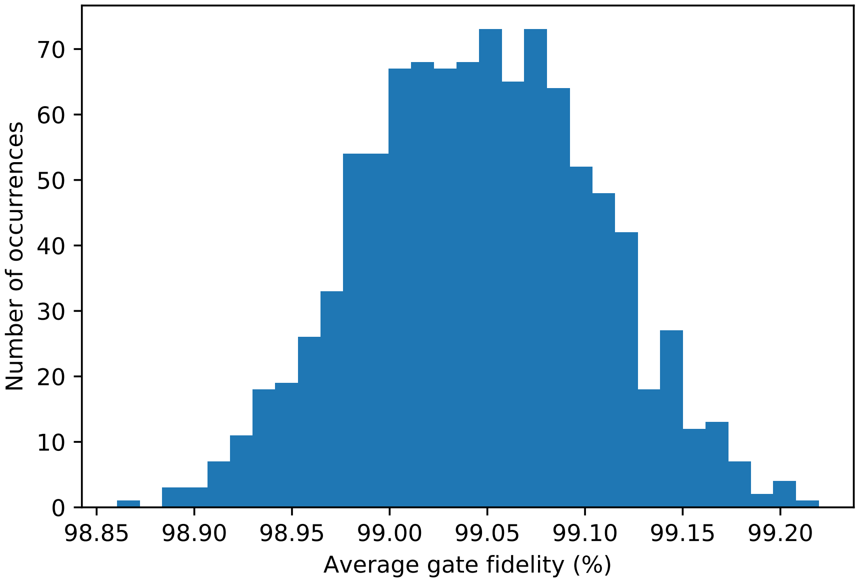

Figure 5.13: Non-parametric bootstrapped estimate of the average gate fidelity $\overline{F}$. The histogram shows the outcome of 1000 re-samples of the process tomography measurements, and has mean 99.05% and standard deviation 0.06%.

5.1.11 Randomised Benchmarking

In the preceding sections we have struggled to find a reliable number to quantify the quality of our gate. The methods used have required large amounts of data, and hence long collection times during which experimental drifts can skew the results. We also do not have confidence normalising out the SPAM errors, as they are larger than, and their drifts swamp our sensitivity to, the gate error.

The first method one might try to surmount these problems and amplify the measured gate error until it is the dominant source, is to do multiple gates in a row and divide the final measured error by the number of gates applied. While this does increase the gate error, the similar input states to each gate mean many sources of error will sum coherently (e.g. geometric phase error), which is highly unrepresentative of a real quantum computation. We instead want to randomise the state between applications of our gate with multiple single- and two-qubit gates, as then the coherent gate errors sum linearly. These randomising single and two-qubit gates however will have errors of their own, so to isolate the error due to our gate we first measure the error of a sequence of these randomising gates alone, and find the additional error of interleaving our gate under test within it. This method is known as randomised benchmarking [31,104–106].

Randomised benchmarking has a number of key advantages over other methods. While the number of measurements needed for process tomography scales exponentially with the number of qubits, it is only polynomial with randomised benchmarking, making it well suited for comparing gate quality between few-qubit and many-qubit systems. As well as providing the gate fidelity averaged over all input states (rather than just one as in Bell state fidelity), we can separate arbitrarily small gate errors from SPAM errors during the procedure. We can also ascertain whether the average gate error is stable over a long sequence or deteriorates with sequence length, an important characteristic for realistic length quantum algorithms. The downside is that it needs individual addressing of each qubit, which is often not available in trapped ion experiments, though subspace randomised benchmarking can be used instead [107]. Thankfully we get this ‘for free’, as using different species with very different qubit frequencies avoids the problem of cross-talk entirely.

We first choose for the set of randomising gates the two-qubit Clifford group, whose members are defined as the set of unitaries C which map every two-qubit Pauli gate (including identities) to another signed Pauli gate, e.g. $C_i(Z \otimes I)C_i^\dagger = -X \otimes Y$, for a particular Clifford C_i. This choice, while not completely general, is sufficient to completely depolarise the error between gates and crucially allows the outcome of gate sequences to be easily simulated, due to the Gottesman-Knill theorem [108]. The experimental procedure is then to cool and prepare the ions to $|{\downarrow\downarrow}\rangle$, apply a sequence of randomly chosen Clifford gates of length l, apply a random Pauli gate, apply a final

Clifford chosen to invert the action of all the preceding gates except the Pauli, and then measure whether the ions were in the state we expected. We repeat this with k different random sequences, each run N times, for each sequence length l. The random Pauli gate is inserted to randomise the measurement outcome so we are sensitive to a wider class of errors. Plotting how this survival probability P_s decays with l, we fit it to the exponential decay model

$$P_s(l) = Ap^l + B \tag{5.21}$$

with parameters

$$A = \frac{1}{\alpha} - \epsilon_m\,, \quad p = 1 - \alpha\epsilon_o\,, \quad B = 1 - \frac{1}{\alpha}\,. \tag{5.22}$$

The constant $\alpha = d/(d-1)$ is in our $d = 4$ dimensional case $\alpha = 4/3$ [106]. We can therefore extract the combined SPAM, random Pauli and final inverting gate error ϵ_m, and the average error per Clifford operation ϵ_o.

We repeat the above process a second time, but now interleave our test gate G' after each Clifford gate, and fit the results to the same model, with the final Clifford again chosen to invert all the preceding gates, including those interleaved, except the random Pauli. The sequence length l still counts the number of Clifford gates, but not the interleaved gates. Bearing in mind that ϵ_m can vary, we can then extract the error per gate as

$$\epsilon_{G'} = \frac{1}{\alpha}\left(1 - \frac{p'}{p}\right) \tag{5.23}$$

where p' is the parameter from the fit of the interleaved experiment. The average gate fidelity is then simply $\overline{F} = 1 - \epsilon_{G'}$.

Randomised benchmarking of our gate was carried out by Amy Hughes, with further details and discussion in her thesis. The interleaved gate G' was the gate as described in section 5.1.4, but with the initial and final spin-echo $\pi/2$ pulses removed, i.e. just the two gate pulses and the π-pulse in the middle; $G' = \sqrt{G} \cdot R_{XX}(\pi) \cdot \sqrt{G}$. The $\pi/2$ pulses are no longer necessary as the ions are no longer always in a z-basis eigenstate before each gate, and because in any optimised quantum algorithm they would be subsumed into surrounding single qubit operations.

The Clifford gates were physically implemented as a near optimal decomposition into X,Y and Z $\pm\pi/2$-rotations on each qubit, carried out simultaneously, along with G'. Z rotations are implemented by simply updating the phase of the local oscillator of the relevant qubit with no delay (see section 3.2), so we expect them to contribute no error. Identity operations I in a sequence of single qubit operations are skipped and replaced with the next operation on that qubit without a delay, reducing the length of the sequence[4]. Each Clifford contains between 0 and 3 two-qubit G' gates, with an average of 1.5 over the whole group. Pauli randomisations were implemented as simultaneous X or Y π-rotations on each qubit, which keeps the final state in the Z measurement basis.

[4]For example a sequence $(R_X(\frac{\pi}{2}) \otimes I) \cdot (I \otimes R_Y(-\frac{\pi}{2})) \cdot (R_Z(\frac{\pi}{2}) \otimes I)$ would be implemented as $(R_X(\frac{\pi}{2}) \otimes R_Y(-\frac{\pi}{2})) \cdot (R_Z(\frac{\pi}{2}) \otimes I)$.

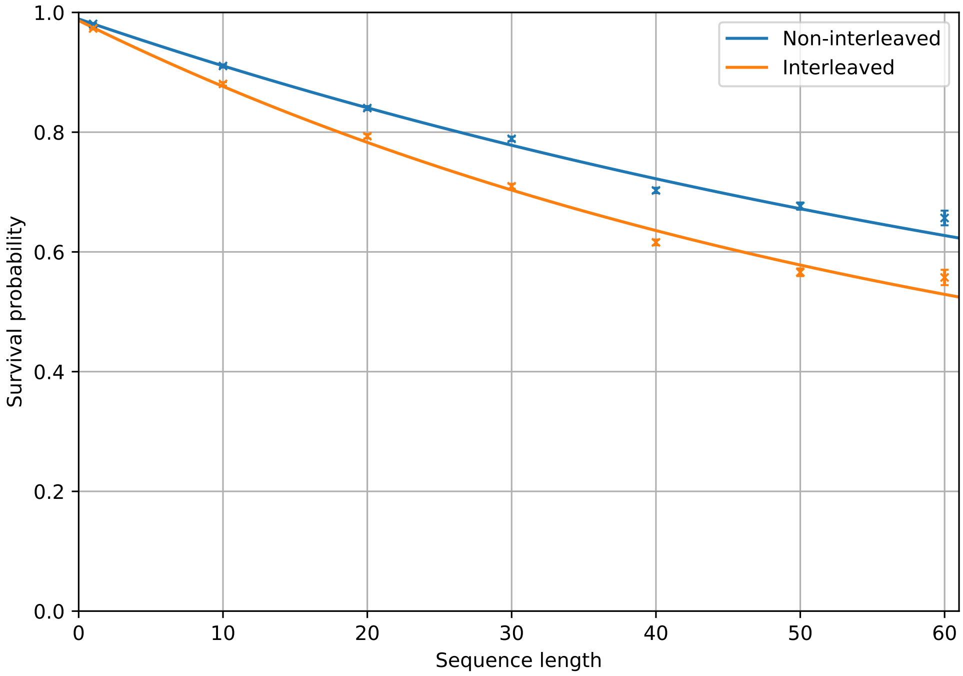

Figure 5.14: Randomised benchmarking of the out-of-phase mode gate, showing the decay in probability of producing the correct outcome for increasingly long sequences of Clifford gates. The extra error of interleaving our gate (orange) fits to an average gate fidelity of $\overline{F} = 99.62(3)\%$.

The results are shown in figure 5.14. They are a combination of multiple randomised benchmarking experiments taken over two days, each measuring different, and not necessarily overlapping, sets of sequence lengths. Each of these experiments used $k = 15$ different random sequences per sequence length l, each run $N = 100$ times. Combining this data, sequences of length $l = \{1, 10, 20, 30, 40, 50, 60\}$ were run with in total $k = \{120, 120, 120, 120, 105, 60, 15\}$ randomisations, which is reflected in the size of the errorbars.

From the least squares fit we find an error per Clifford operation in the non-interleaved sequences of $\epsilon_o = 0.84(3)\%$, the SPAM/Pauli/inversion error averaged over both interleaved and non-interleaved sequences of $\epsilon_m = 1.2(3)\%$ and the error per interleaved gate $\epsilon_{G'} = 0.38(3)\%$. The uncertainties on these figures are derived from the variances of the estimated fit parameters, except for $\epsilon_{G'}$, where a non-parametric bootstrapping technique was used, analogous to that in the previous section. We therefore find an average gate fidelity $\overline{F} = 99.62(3)\%$.

It must be noted that this figure is the average over 60 gates. While we hope the error per gate is constant with sequence length, in general heating and other effects will increase this error over time, especially in our somewhat high heating rate trap. Repeating the fit but only up to sequence lengths of 30 gates yields a slightly higher fidelity of $\overline{F} = 99.66(5)\%$. Further discussion of how the error per gate changes with sequence length in these experiments will be presented in Amy Hughes' Thesis.

5.1.12 Discussion

In principle we expect these varying methods of assessing the gate to give fidelities which agree to within their statistical uncertainties. In practice however systematic effects cause them to differ significantly. Table 5.3 summarises the Bell state fidelities and average gate fidelities obtained from the four methods used. In general we expect the average gate fidelities to be higher than the Bell state fidelities, as when averaging over all input states some have one or both qubits orthogonal to the gate basis[5].

	Bell State Fidelity (F)	Average Gate Fidelity ($\overline{F}$)
Bell State Tomography	99.80(10)%	-
Full State Tomography	98.1%	-
Process Tomography	-	99.05(6)%
Randomised Benchmarking	-	99.62(3)%

Table 5.3: Summary of the fidelities measured on the out-of-phase mode gate using a variety of methods.

We can see that both state tomography and process tomography give fidelities below what we expect from the other measurements, which can be attributed to several fac-

[5] i.e. the gate in principle does nothing to $|y_+\rangle \otimes |y_+\rangle$.

tors. Firstly, the state tomography experiments were carried out before the gate, single qubit pulses and SPAM were fully optimised, with the errors of the latter measuring 60% higher than usual, averaging $\bar{\epsilon} = 0.48(9)\%$. Though the data was readout error normalised before further processing, poorly controlled SPAM errors are often highly variable and can therefore lead to unreliable fidelity estimates.

We also suspect that, even after optimisation, inconsistent single qubit pulse areas during the pre- and post-gate rotations are a source of the discrepancy in fidelity. This is because over the 500 shots per point in each experiment the RF amplifiers warm up and change their amplification. We can only account for a few tenths of a percent infidelity from this, so is too small to be the dominant source.

The rest is likely due to experimental drifts during the state and process tomography experiments, which lasted up to $\sim$3 hours. Whilst between Bell state tomography experiments, each lasting $\sim$10 minutes, many experimental parameters were tuned up, this was not possible during long periods of data collection. It is likely drifts in laser frequencies (e.g. in the dark resonance cooling) and trap frequencies or laser instabilities (e.g. in the Raman lasers) are the leading sources of error.

Bell State Tomography is also not an ideal method to pin down the magnitude of the gate infidelity. As explained in section 5.1.8, readout normalisation can be fraught with difficulty and can produce fidelities consistent with but greater than 1. Taking enough statistics to overcome this leads to the same difficulties described above.

Randomised benchmarking therefore provides our best estimate of the gate fidelity. We can both amplify the gate errors with comparatively low statistics, and separate the SPAM and single qubit rotation errors. The downside however of applying repeated gates and averaging their error is that if the error per gate is not constant, it is difficult to assign a clear value to the error for one gate. If the increase in error per gate is largely due to heating, this could be mitigated in future experiments by adding a third ion to sympathetically cool the crystal between two-qubit gates.

With these caveats in mind we believe the remaining error on the in-phase mode gate is primarily due to mode heating errors from our large anomalous heating rate, of $\dot{\bar{n}} \approx 100$ quanta/s for a single $^{43}Ca^{+}$. The remaining error on the out-of-phase mode gate is likely due to both heating errors and Kerr cross coupling errors from hot radial modes (see section 3.6.1). In future we could reduce the latter by sideband cooling the radial modes with the 674 nm laser.

Comparison to previous work

This gate mechanism, with the beams tuned between the resonances of two species but insensitive to their qubit frequencies, was first carried out between $^{40}Ca^{+}$-$^{43}Ca^{+}$ and published in [109], in a previous iteration of the apparatus. Notably, the trap had a significantly lower heating rate of $\dot{\bar{n}} \approx 1$ quanta/s for a single $^{43}Ca^{+}$, was operated at a low magnetic field of ∼2 G, and used only a single 850 nm beam for readout (using an EIT method for $^{40}Ca^{+}$) which gave average SPAM errors of $\bar{\epsilon} = 5.5\%$.

A proof of principle experiment between $^{43}Ca^{+}$-$^{88}Sr^{+}$ with this mechanism was presented in [39] using the apparatus described in this work, with the exception of a preliminary readout scheme for $^{88}Sr^{+}$ instead of shelving with the 674 nm laser. This involved a 408 nm laser resonant with the $S_{1/2}$-$P_{3/2}$ transition, which gave large SPAM errors similar to the $^{40}Ca^{+}$-$^{43}Ca^{+}$ experiment.

Compared to these two previous experiments which achieved fidelities of $F = 99.8(6)\%$ and $F = 83(3)\%$ respectively, the mixed species gate described here has lower errors, with tighter bounds. This is mainly due to better readout on both species, especially the addition of the 674 nm shelving laser, better cooling of the $^{43}Ca^{+}$-$^{88}Sr^{+}$ crystal, and many small reliability and stability improvements. These experiments were also carried out at 146 G, in comparison to the low field used in the $^{40}Ca^{+}$-$^{43}Ca^{+}$ experiments, which allows access to the clock qubit of $^{43}Ca^{+}$. Though our gate mechanism cannot be used on it directly, microwave pulses can be used to transfer from the stretch to the clock qubit, as described in section 2.7.

5.2 $^{88}\mathrm{Sr}^{+}$– $^{88}\mathrm{Sr}^{+}$ Mølmer-Sørensen Gate

Before conducting the mixed species gate described above, we conducted an entangling gate between two $^{88}\mathrm{Sr}^{+}$ ions, as a test of our new 674 nm narrow-linewidth quadrupole laser system and $^{88}\mathrm{Sr}^{+}$ operations in general. A further motivation was a planned $^{43}\mathrm{Ca}^{+}$-$^{88}\mathrm{Sr}^{+}$ gate using both the Raman and 674 nm beams for a mixed species Mølmer-Sørensen gate in future experiments. The gate described here was carried out on the optical qubit with a 674 nm gate beam made up of two tones symmetric about the qubit carrier frequency. Each sideband is detuned from the axial stretch motional mode on which we conduct the gate by the gate detuning δ, in the direction away from the carrier, creating an off-resonant Mølmer-Sørensen force. As in the light-shift gate, this drives loops in phase space which acquire a geometric phase, here in the $\sigma_\phi \otimes \sigma_\phi$ basis[6]. Choosing a geometric phase of $\pi/2$ then implements an entangling operation between the two ions. This mechanism is independent of the ion spacing, but does depend on the qubit frequency. Though the sidebands create light-shifts on the qubit carrier, they are in opposite directions and cancel if they have equal Rabi frequency.

5.2.1 Bichromatic Beam

To create the two tones necessary for the gate we applied two RF frequencies to the last acousto-optic modulator (AOM) in the 674 nm optical chain. The 200 MHz AOM is in a single pass configuration with lenses either side set a focal length away from the AOM centre. Along with the fibre coupling lens, this improves coupling into the trap going fibre placed afterwards, as the two tones have only an angular rather than positional displacement going into the fibre tip. The two frequencies of light

[6]Where ϕ is a direction on the equator of the Bloch sphere.

superimpose and interfere, creating a beam at the mean optical frequency amplitude modulated at half the difference frequency. Due to the non-linear response of the AOMs diffracted optical power with input RF power, we get harmonics of this beat frequency that are undesirable for our gate mechanism. Though far off-resonant, these harmonics are suppressed by driving the AOM with 12% of the normal RF power, so the optical power in the first harmonic is ~10% that of the beat note itself, as measured by a photodiode[7] before the trap. With the same photodiode we also measure the depth of the amplitude modulation of the bichromatic tone, where anything less than unity modulation depth implies a power imbalance between the two tones. As we want their Rabi frequencies equal, we adjust the amplitude of the RF sources to counter the frequency response of the AOM, which in general is not flat. The bichromatic beam enters the vacuum system at close to 45° to the trap axis, so has significant projection onto both the axial and radial motional modes. The geometry with respect to the other beams and the magnetic field is shown in figure 4.7.

5.2.2 Experimental Sequence

Akin to section 5.1.3, each experiment consists of Doppler cooling, state-preparation, sideband cooling, the gate pulse, optional analysis pulse and readout, repeated N times per setting of the analysis pulse phase. We prepare into the $|\downarrow\rangle$ state of the optical qubit for both ions, and do the single qubit analysis pulse with the (motionally sensitive) global 674 nm laser, rather than with RF as in the previous section. Here the $|\uparrow\rangle$ qubit state is measured dark rather than the usual bright, as it is out of the Doppler cooling cycle, and as the two ions are indistinguishable we can only measure the probabilities $\{p_{\uparrow\uparrow}, p_{\uparrow\downarrow}+p_{\downarrow\uparrow}, p_{\downarrow\downarrow}\}$ (see section 2.7.1). Unlike in the mixed species experiment, the gate pulse drives a single loop in phase space, and is not placed inside a spin-echo sequence. Since the $|\downarrow\rangle$ state we start in is already in an equal superposition of all four ϕ-basis

[7] Thorlabs DET025A/M, 2GHz bandwidth

states[8] on which the gate acts this is no longer necessary, but it does increase our sensitivity to several sources of error, notably spin dephasing error. We do still shape the gate pulse, again with $t_{\text{shape}} = 2\,\mu\text{s}$, primarily to reduce off-resonant excitation of the qubit carrier. Since we could not use the B servo to trim the magnetic field, as it relies on measuring the $^{43}\text{Ca}^{+}$ stretch qubit frequency, and we were worried about slow drifts of the 674 nm laser locking cavity, we servoed the 674 nm laser frequency directly on the $^{88}\text{Sr}^{+}$ qubit by measuring its frequency in a Ramsey experiment every 10 minutes. To reduce the effects of magnetic field noise from the lab electrical mains on the gate, each shot of the experiment was triggered to run at the same phase in the 50 Hz cycle, with the downside that this significantly slows down data collection.

Cooling

We Doppler cool the $^{88}\text{Sr}^{+}$-$^{88}\text{Sr}^{+}$ crystal for 0.5 ms before state-preparation and pulsed sideband cooling using the 674 nm laser. As we have projection onto all motional modes we sequentially cool them in the order {radial COM upper, radial COM lower, radial rocking upper, radial rocking lower, axial COM, axial stretch}; in increasing importance of their temperature to the gate. As modes earlier in the sequence are heated by photon recoil during the cooling of later modes, we repeat this sequence again with half the number of sideband pulses, for a total of 50 pulses per mode. We linearly increase the pulse duration each pulse in an attempt to match the increasing sideband π-time[9] as the crystal cools. The number of pulses and the initial and final sideband pulse durations were set empirically on the measured ion temperature. We achieve temperatures of roughly $\bar{n} \approx 0.07$ for the stretch mode, and $\bar{n} \approx 0.7$ for the other modes.

[8] i.e. $|{\downarrow\downarrow}\rangle = \left(|{\rightarrow\rightarrow}\rangle + |{\rightarrow\leftarrow}\rangle + |{\leftarrow\rightarrow}\rangle + |{\leftarrow\leftarrow}\rangle\right)/2$.

[9] More properly called the maximum inversion time.

5.2.3 Implementation

Mode frequency	2.324 MHz
Gate detuning	13.5 kHz
Gate duration	74 μs
Gate beam power	10 mW
Lamb-Dicke parameter	0.023

Table 5.4: Parameters for the single loop $^{88}Sr^{+}$-$^{88}Sr^{+}$ gate on the stretch mode. The beam power is the time-average for the 674 nm bichromatic beam, focussed to a 60 μm diameter spot on the ion. The Lamb-Dicke parameter is calculated from experimental parameters, and includes the mode eigenvector magnitude.

We performed gates on the stretch mode of the $^{88}Sr^{+}$-$^{88}Sr^{+}$ crystal, with the parameters shown in table 5.4. As only one of the trap endcaps was connected to the 'tickle' drive, the purely symmetric stretch mode frequency could not directly be found. We instead found the axial COM mode frequency and used the relation $\omega_{\text{str}} = \sqrt{3}\omega_z$. The rest of the parameters were set by looking at the state populations while scanning the gate duration, as in section 5.1.7, and also while scanning the gate detuning. The behaviour was not symmetric over positive and negative detunings, which suggested an undiagnosed light-shift offsetting the qubit carrier and motional sideband frequencies. Despite using identical Rabi frequencies, this light-shift was only seen on the gate dynamics, i.e. when using the bichromatic beam, and not during carrier operations, which pointed to its source being either a slight imbalance in the Rabi frequencies of the two tones, or coupling to other levels. This was corrected by moving the centre of the bichromatic tone by 3 kHz.

5.2.4 Results

The results of a Bell state tomography fidelity measurement are shown in figure 5.15. We fit the parity fringe using a maximum likelihood method based on the binomial probability of each point, and find a fidelity $F = 95.9(4)\%$. SPAM errors were not measured concurrently with this experiment, so we do not normalise them out here. This was a preliminary result and was not optimised further in the interest of time. We can account for a few tenths of a percent infidelity from the imperfect analysis pulse, as the residual temperature on all modes causes carrier depletion on the motionally sensitive 674 nm pulse, but the source of the rest was not properly investigated.

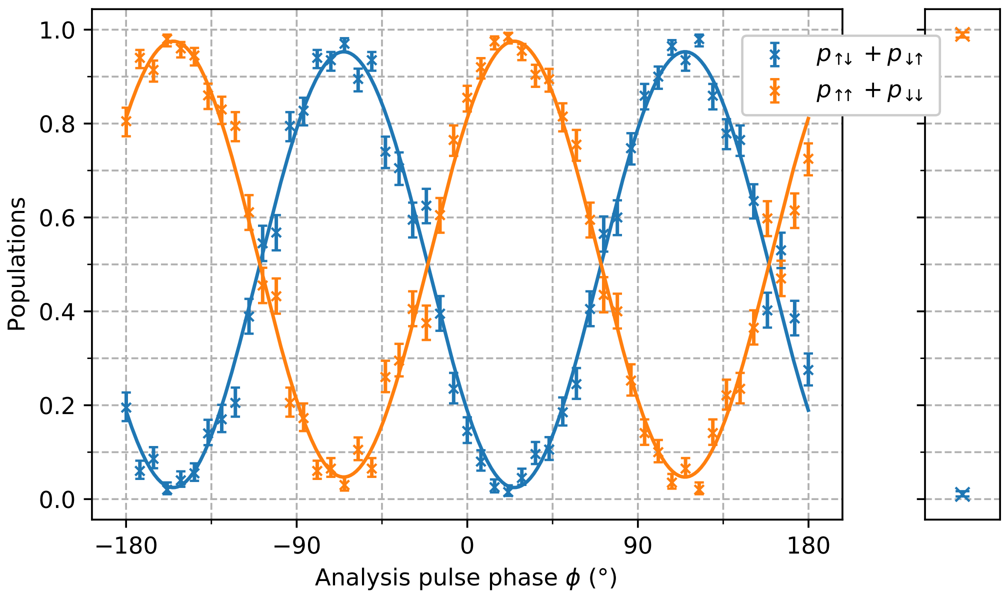

Figure 5.15: Bell state fidelity measurement of the $^{88}Sr^{+}$-$^{88}Sr^{+}$ gate, without readout error correction. The even and odd spin state populations after the analysis pulse (Left) are fitted using a maximum likelihood method to find the parity contrast. Together with the populations directly after the gate (Right), the fidelity is $F = 95.9(4)\%$.

During the period of this experiment we had large undiagnosed problems with the magnetic field stability, causing unpredictable and sometimes catastrophic noise bursts. For this reason when measuring the Bell state fidelity we used a full scan of the parity fringe rather than a four or two point scan as described in section 5.1.8, as any noise would be more apparent. A bad temperature controller on the 422 nm laser locking cavity also caused substantial drifts on the Doppler beams, giving unstable ion fluorescence and cooling. These drifts were corrected every 10 minutes by scanning the 422 nm resonance and setting the frequency back to the usual half fluorescence point, but faster drifts were not corrected.

These problems were discovered and solved before taking the final data for the mixed species gate experiment in the previous section.

6

Conclusion

In this thesis I have described the construction of a trapped ion apparatus for experiments on $^{43}Ca^+$ and $^{88}Sr^+$ ions. We have demonstrated an entangling Mølmer-Sørensen (MS) operation between the optical qubits of two $^{88}Sr^+$ ions, as a test of the system for future work. We have also demonstrated an entangling operation between two different ion species, $^{43}Ca^+$ and $^{88}Sr^+$, using a pair of beams independent of either qubit frequency and derived from a single laser. This light-shift gate scheme relies on the proximity of the S-P transition frequencies in the two species, separated by only 20 THz, with the gate beams tuned between them having off-resonant Raman couplings of similar magnitude to both. Multiple tomography methods along with randomised benchmarking were used to asses the quality of the gate, with its errors found to be almost an order of magnitude below the threshold of one of the best error correcting protocols, the surface code [110,111]. It therefore demonstrates one of the fundamental operations of a mixed species trapped ion networked quantum computer.

6.1 Comparison to Other Work

The first mixed species quantum logic experiments in trapped ions were demonstrations of quantum logic spectroscopy, where the transitions of ions without good laser cooling, state preparation or state detection are probed by transferring their state to an auxiliary 'logic' ion [112–114].

Subsequently, the first mixed species entangling gates were carried out between $^{40}Ca^{+}$-$^{43}Ca^{+}$, using the same scheme used in this work, and $^{9}Be^{+}$-$^{25}Mg^{+}$, using a Raman beam pair for each species in a MS interaction, published concurrently [109,115]. They achieved fidelities of $F = 99.8(6)\%$, limited by uncertainties in the SPAM error, and $F = 97.9(1)\%$, limited by SPAM, photon scattering and heating errors, respectively. A similar mixed species Raman MS gate between $^{171}Yb^{+}$-$^{138}Ba^{+}$ has also been reported, with approximate fidelity $F = 60\%$ limited by the extremely high heating rate of the trap [116].

A three ion mixed species gate has been demonstrated in a $^{9}Be^{+}$-$^{40}Ca^{+}$-$^{9}Be^{+}$ crystal, using a Raman beam pair for $^{9}Be^{+}$ and a narrow linewidth quadrupole laser for $^{40}Ca^{+}$, in a MS scheme producing a maximally entangled three-qubit state with $F = 93.8(5)\%$ [117]. In the same report, a Bell state between the two $^{9}Be^{+}$ ions with the $^{40}Ca^{+}$ as a spectator was also created, with fidelity $F = 97.8(4)\%$. Both fidelities were mainly limited by decoherence of the ions' motion.

Two $^{9}Be^{+}$-$^{25}Mg^{+}$ crystals in separate trapping zones, have been used to demonstrate deterministic quantum gate teleportation [118]. Mixed element gates, using the same mechanism as [115], transfer the initial entanglement between the two $^{25}Mg^{+}$ ions onto the two $^{9}Be^{+}$ ions, using only local operations and classical communication. The teleported controlled-NOT operation was achieved with a fidelity in the 95% confidence interval $\{84.5\%, 87.2\%\}$, and the $^{9}Be^{+}$-$^{25}Mg^{+}$ gates contained within it with $F = 97(1)\%$.

Finally, mixed species MS gates between $^{40}Ca^{+}$ and $^{88}Sr^{+}$ have recently been been reported using a narrow-linewidth quadrupole laser for each species, achieving a fidelity of $F = 94.3(3)\%$ [119]. Applying the gate to only the $^{88}Sr^{+}$ ions in a $^{88}Sr^{+}$-$^{40}Ca^{+}$-$^{88}Sr^{+}$ crystal yielded a fidelity of $F = 95.7(3)\%$, with both gates limited by cryocooler vibrations and laser phase instabilities.

6.1.1 Discussion

The results presented in this work are an extension of the mechanism demonstrated in the $^{40}Ca^{+}$-$^{43}Ca^{+}$ experiments, from mixed isotope to mixed element operations. This yields significant advantages in reduced cross-talk, independent operation and species specialisation in a mixed species ion trap node. A comparison with this experiment can be found in section 5.1.12.

This work achieves the highest yet reported fidelity for a mixed species gate, and reduces the error for mixed element gates by an order of magnitude. This is in large part due to the simplicity of the scheme, needing only two frequencies separated by AOMs from one laser to operate simultaneously on both species, independent of either qubit frequency. As with other schemes using Raman beam pairs we only care about their frequency difference, which, along with their $\sim$10 THz detuning from either species, avoids the need for frequency locking of the gate laser to an external cavity. We also avoid complications managing the relative phase of beams from multiple lasers, or their phase coherence to each qubit.

A significant downside to this, and all, light-shift gates is the inability to apply it directly to first order magnetic field insensitive 'clock' qubits, on which the long coherence times in trapped ions have been reported. We can get around this by mapping back and forth from the 'clock' qubit with microwave pulses; indeed we choose our isotope of $^{43}Ca^{+}$ and operate at its 'clock' qubit's field insensitive point for this reason. This process however still suffers from errors, from off-resonant transitions during the mapping and magnetic field fluctuations during the time spent in the magnetically sensitive qubit for the gate.

6.2 Outlook

Many of the physics barriers involved in building a practical quantum computer have been surpassed in recent years, in small systems of a few ions. The remaining challenge of growing these systems while keeping their performance, is largely one of engineering and design. As more ions need to be managed, preventing the size and complexity of their control systems from ballooning is of paramount importance. The relative simplicity of our mixed species entangling operation could therefore prove to be a useful building block.

6.2.1 Improvement

The technical requirements for this scheme can be reduced further. Currently we use two phase-locked frequency-doubled Ti:Sapphire lasers for the gate and motional sideband operations, as they produce sufficient power with low noise at 402 nm and bridge the $\sim$3 GHz $^{43}Ca^{+}$ qubit frequency. We could instead use direct diodes in the ultraviolet, and achieve the same power and noise characteristics through injection locking and spectral filtration with a diffraction grating [120]. We can further drop the requirement of needing $\sim$3 GHz separated beams, by sideband cooling on the Zeeman qubit of $^{88}Sr^{+}$, using an AOM to bridge the $\sim$400 MHz gap, or use the 674 nm quadrupole laser to sideband cool on the optical qubit.

To lower the currently crippling overhead in ion number imposed by quantum error correction, the errors of qubit operations in general need to be reduced substantially. While the gate described here approaches the state-of-the-art fidelity of two-qubit gates across all physical platforms, we can take steps to improve it further [15, 16, 121, 122]. We currently do not cool the radial modes below the Doppler limit, and indeed expect them to be heated by the sideband cooling on the other modes. This causes Kerr cross coupling errors for gates on both modes, especially the out-of-phase mode. We could reduce this error in our existing setup by sideband cooling these modes with the 674 nm quadrupole laser.

Another source of error is the anomalous heating of the ion during the gate. While we already use a macroscopic trap with large ion-electrode separation, we could reduce this further by operating under cryogenic conditions, though this would require a substantial redesign of the apparatus and introduce further complexity. Though in principle different materials or surface treatments could be used for the trap electrodes, the effect on the anomalous heating rate is poorly understood.

6.2.2 Application to Other Ion Pairs

This scheme is not limited to the particular choice of Ca^+ and Sr^+, though they make a complementary pair for a mixed species system [123]. While using pairs of isotopes of the same element would readily work, pairs of different elements need to be chosen carefully.

In particular a mass ratio close to 1 is desirable, to keep the magnitude of the motional mode eigenvectors similar for both species, for efficient sympathetic cooling and gate operations. For our gate mechanism it is also necessary to choose a pair whose S-P transitions are not too far separated, as the gate Rabi frequency scales with detuning to these transitions Δ, and Raman laser intensity I, as $\Omega \propto I/\Delta$.

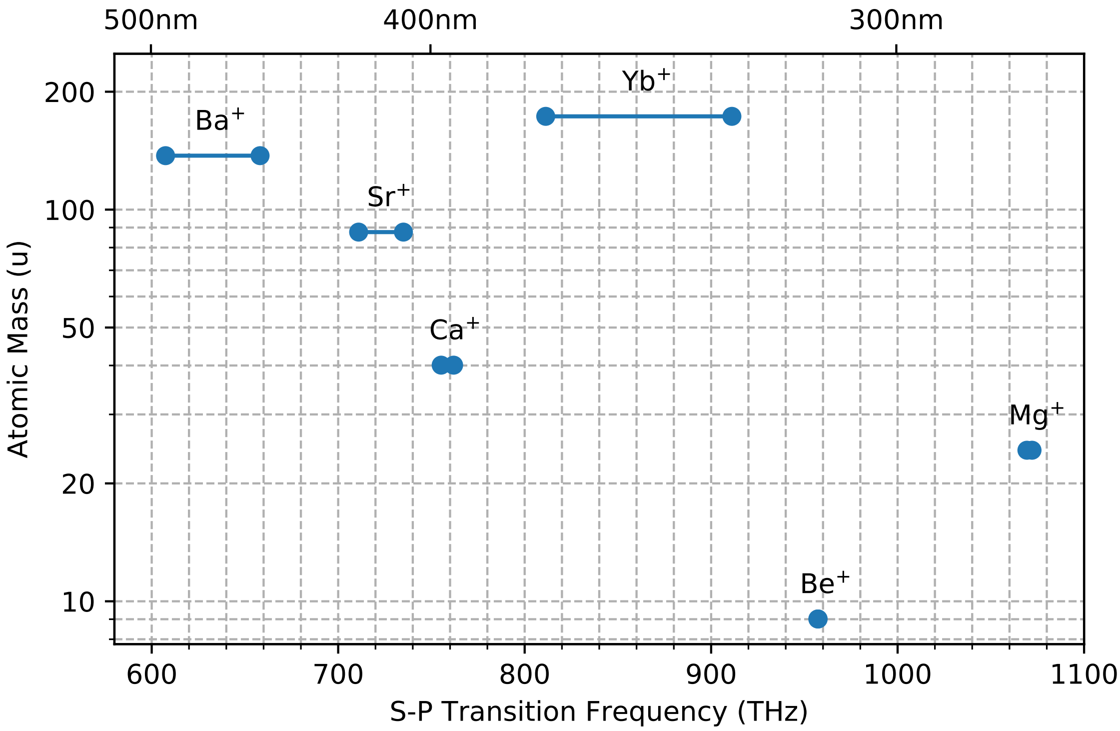

Figure 6.1: Masses and transition frequencies of ions commonly selected for trapped ion quantum information experiments. Each species has two points, representing its $S_{1/2}$ - $P_{1/2}$ and $S_{1/2}$ - $P_{3/2}$ transitions. Similar atomic masses are important for mixed species ion pairs to be efficiently cooled, while close S-P transition frequencies are favourable for entangling operations with the mechanism described in this work. The masses are shown on a logarithmic scale, so equal mass ratios are separated by equal vertical distances.

Figure 6.1 shows the S-P transition frequencies and masses of common ion species used in trapped ion quantum information experiments. Ca^+ and Sr^+ have a mass ratio of $\sim$2.2 and their transitions have a minimum separation of 20 THz. For reference, for our 50 μs gate we currently use a total of 120 mW of Raman laser power focussed to 50 μm diameter spots.

Ba^+-Sr^+ would also make an appropriate pair, with a slightly more favourable mass ratio of $\sim$1.6 but an increased transition separation of 53 THz. Gates on this pair of ions, ignoring differences in the magnitude of coupling matrix elements, would only require a $\sim$2.5$\times$ increase in laser intensity, easily achievable with smaller beam diameters or a modest increase in power.

Gates on Ca^{+}-Yb^{+} or Be^{+}-Mg^{+} are also possible with this mechanism, though their large mass imbalance and frequency gap respectively could prove problematic.

6.2.3 Future Work

An initial application of this work would be to conduct an experiment extending the memory time of the magnetically sensitive Zeeman qubit in $^{88}Sr^{+}$. Using the gate to conduct a swap operation between the two species, and microwave pulses to transfer to the $^{43}Ca^{+}$ 'clock' qubit and back again, the effective coherence time of the information stored on the Zeeman qubit could be increased by 'hiding' it in the 'clock' qubit. This would only be a useful technique for long delay times, when the coherence of the Zeeman qubit would otherwise fall below the combined errors from the swap and microwave transfer operations.

A further extension would be to combine this work with the recent fast gate experiments conducted in the same apparatus [19]. Using the same gate mechanism but with amplitude shaped pulses, carefully chosen to remove dependence on the initial optical phase, gate durations over an order of magnitude faster than in this work were reported. Though applying this scheme to a mixed species crystal may increase the complexity of finding appropriate pulse shapes, in principle similar speeds should be achievable.

The natural continuation of this work and our current focus in the lab, is to leverage the recent and substantial progress entangling $^{88}Sr^{+}$ ions in neighbouring ion trap systems [27]. These results demonstrated remote entanglement generation in multi-zone surface traps with sufficient fidelity and rate, that we can attempt to build an elementary quantum network. This would consisting of two mixed species trapped ion nodes with a photonic link, and could demonstrate distillation and transfer of remote entanglement from the ancilla/interconnect $^{88}Sr^{+}$ ions, to the $^{43}Ca^{+}$ logic ions [29].

Achieving high fidelity entanglement between two logic ions in separate traps, through no direct entanglement operation between them, would mark an important milestone on the way towards scalable trapped ion quantum computing.

Bibliography

[1] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 1982. [Cited on page 1]

[2] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London, Series A*, 400(1818):97–117, 1985. [Cited on page 1]

[3] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London, Series A*, 439(1907):553–558, 1992. [Cited on page 1]

[4] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219, Philadelphia, PA, USA, 1996. [Cited on page 1]

[5] P. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, Santa Fe, NM, USA, 2002. [Cited on page 1]

[6] P. W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):R2493–R2496, 1995. [Cited on page 2]

[7] A. M. Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77(5):793–797, 1996. [Cited on page 2]

[8] Y. Cao, J. Romero, and A. Aspuru-Guzik. Potential of quantum computing for drug discovery. *IBM Journal of Research and Development*, 62(6):6:1–6:20, 2018. [Cited on page 2]

[9] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017. [Cited on page 2]

[10] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000. [Cited on page 2]

[11] C. R. Monroe and D. J. Wineland. Quantum computing with ions. *Scientific American*, 299(2):64–71, 2008. [Cited on page 3]

[12] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20):4091–4094, 1995. [Cited on pages 3 and 41]

[13] 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. [Cited on page 3]

[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. [Cited on pages 3, 12, and 70]

[15] 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. [Cited on pages 3 and 136]

[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 $^{9}Be^{+}$ ion qubits. *Physical Review Letters*, 117(6):060505, 2016. [Cited on pages 3 and 136]

[17] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Experimental comparison of two quantum computing architectures. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3305–3310, 2017. [Cited on page 3]

[18] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos. Universal digital quantum simulation with trapped ions. *Science*, 334(6052):57–61, 2011. [Cited on page 3]

[19] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas. Fast quantum logic gates with trapped-ion qubits. *Nature*, 555(7694):75–78, 2018. [Cited on pages 3 and 139]

[20] N. Friis, O. Marty, C. Maier, C. Hempel, M. Holzäpfel, P. Jurcevic, M. B. Plenio, M. Huber, C. Roos, R. Blatt, and B. Lanyon. Observation of entangled states of a fully controlled 20-qubit system. *Physical Review X*, 8(2):021012, 2018. [Cited on page 3]

[21] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. 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):259–328, 1998. [Cited on pages 3, 40, 52, 62, and 104]

[22] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, 2002. [Cited on page 3]

[23] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Physical Review Letters*, 78(16):3221–3224, 1997. [Cited on page 4]

[24] L. M. Duan, B. B. Blinov, D. L. Moehring, and C. Monroe. Scalable trapped ion quantum computation with a probabilistic ion-photon mapping. *Quantum Information and Computation*, 4(3):165–173, 2004. [Cited on page 4]

[25] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013. [Cited on page 4]

[26] 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. [Cited on page 4]

[27] 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. *Physical Review Letters*, 124(11):110501, 2020. [Cited on pages 4 and 139]

[28] D. J. Larson, J. C. Bergquist, J. J. Bollinger, W. M. Itano, and D. J. Wineland. Sympathetic cooling of trapped ions: A laser-cooled two-species nonneutral ion plasma. *Physical Review Letters*, 57(1):70–73, 1986. [Cited on page 4]

[29] 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. [Cited on pages 5 and 139]

[30] Q. A. Turchette, D. Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A*, 61(6):063418, 2000. [Cited on pages 11, 52, and 62]

[31] M. A. Sepiol, A. C. Hughes, J. E. Tarlton, D. P. Nadlinger, T. G. Ballance, C. J. Ballance, T. P. Harty, A. M. Steane, J. F. Goodwin, and D. M. Lucas. Probing qubit memory errors at the part-per-million level. *Physical Review Letters*, 123(11):110503, 2019. [Cited on pages 12 and 119]

[32] C. J. Ballance. *High-Fidelity Quantum Logic in* Ca^+. PhD thesis, University of Oxford, 2014.
[Cited on pages 14, 35, 43, 47, 52, 53, 55, 56, 60, 62, 63, 73, 80, 89, 96, and 105]

[33] C. J. S. Donald. *Development of an Ion Trap Quantum Information Processor*. PhD thesis, University of Oxford, 2000. [Cited on page 15]

[34] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83(10):5025–5033, 1998. [Cited on pages 15 and 18]

[35] D. F. V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B*, 66(2):181–190, 1998.
[Cited on pages 16 and 17]

[36] A. Steane. The ion trap quantum information processor. *Applied Physics B*, 64(6):623–642, 1997. [Cited on pages 16 and 17]

[37] D. Kielpinski, B. E. King, C. J. Myatt, C. A. Sackett, Q. A. Turchette, W. M. Itano, C. Monroe, D. J. Wineland, and W. H. Zurek. Sympathetic cooling of trapped ions for quantum logic. *Physical Review A*, 61(3):032310, 2000.
[Cited on page 16]

[38] S. Kotler, N. Akerman, N. Navon, Y. Glickman, and R. Ozeri. Measurement of the magnetic interaction between two bound electrons of two separate ions. *Nature*, 510(7505):376–380, 2014. [Cited on pages 16 and 41]

[39] V. M. Schäfer. *Fast Gates and Mixed-Species Entanglement with Trapped Ions.* PhD thesis, University of Oxford, 2018. [Cited on pages 17, 35, 47, 52, 59, 63, 73, 74, 79, 81, 90, and 125]

[40] J. B. Wübbena, S. Amairi, O. Mandel, and P. O. Schmidt. Sympathetic cooling of mixed-species two-ion crystals for precision spectroscopy. *Physical Review A*, 85(4):043412, 2012. [Cited on pages 17 and 22]

[41] J. P. Home. Quantum science and metrology with mixed-species ion chains. In E. Arimondo, P. R. Berman, and C. C. Lin, editors, *Advances in Atomic, Molecular and Optical Physics*, volume 62 of *Advances In Atomic, Molecular, and Optical Physics*, pages 231–277. Academic Press, 2013. [Cited on pages 20, 36, and 99]

[42] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75(1):281–324, 2003. [Cited on page 22]

[43] D. T. C. Allcock, T. P. Harty, M. A. Sepiol, H. A. Janacek, C. J. Ballance, A. M. Steane, D. M. Lucas, and D. N. Stacey. Dark-resonance Doppler cooling and high fluorescence in trapped $^{43}\mathrm{Ca}^{+}$ ions at intermediate magnetic field. *New Journal of Physics*, 18(2):023043, 2016. [Cited on page 22]

[44] T. P. Harty. *High-Fidelity Microwave-Driven Quantum Logic in Intermediate-Field $^{43}Ca^{+}$.* PhD thesis, University of Oxford, 2013. [Cited on pages 24, 25, and 70]

[45] B. H. Bransden and C. J. Joachain. *Physics of Atoms and Molecules; 2nd ed.*, pages 638–639. Prentice-Hall, Harlow, 2003. [Cited on page 27]

[46] R. E. Raab and O. L. De Lange. *Multipole Theory in Electromagnetism: Classical, Quantum, and Symmetry Aspects, with Applications.* International series of monographs on physics. Clarendon Press, Oxford, 2005. [Cited on page 27]

[47] R. Loudon. *The Quantum Theory of Light; 2nd ed.* Oxford Science Publications. Clarendon Press, Oxford, 1997. [Cited on page 27]

[48] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics, Volume 2.* Wiley, New York, NY, 1977. [Cited on page 27]

[49] L. D. Barron and C. G. Gray. The multipole interaction Hamiltonian for time dependent fields. *Journal of Physics A*, 6(1):59–61, 1973. [Cited on page 28]

[50] D. F. V. James and J. Jerke. Effective Hamiltonian theory and its applications in quantum information. *Canadian Journal of Physics*, 85(6):625–632, 2007. [Cited on page 33]

[51] R. Ozeri, C. Langer, J. Jost, B. DeMarco, A. Ben-Kish, B. Blakestad, J. Britton, J. Chiaverini, W. Itano, D. Hume, D. Leibfried, T. Rosenband, P. Schmidt, and D. Wineland. Hyperfine coherence in the presence of spontaneous photon scattering. *Physical Review Letters*, 95(3):030403, 2005. [Cited on page 35]

[52] D. J. Wineland, M. Barrett, J. Britton, J. Chiaverini, B. DeMarco, W. M. Itano, B. Jelenković, C. Langer, D. Leibfried, V. Meyer, T. Rosenband, and T. Schätz. Quantum information processing with trapped ions. *Philosophical Transactions of the Royal Society A*, 361(1808):1349–1361, 2003. [Cited on page 35]

[53] E. T. Jaynes and F. W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, 1963. [Cited on page 39]

[54] F. W. Cummings. Stimulated emission of radiation in a single mode. *Physical Review*, 140(4A):A1051–A1056, 1965. [Cited on page 39]

[55] A. Sørensen and K. Mølmer. Quantum computation with ions in thermal motion. *Physical Review Letters*, 82(9):1971–1974, 1999. [Cited on pages 42 and 46]

[56] D. Leibfried, B. DeMarco, V. Meyer, D. M. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422(6930):412–415, 2003. [Cited on pages 42, 47, and 104]

[57] A. Bermudez, P. O. Schmidt, M. B. Plenio, and A. Retzker. Robust trapped-ion quantum logic gates by continuous dynamical decoupling. *Physical Review A*, 85(4):040302, 2012. [Cited on page 42]

[58] R. Ozeri. The trapped-ion qubit tool box. *Contemporary Physics*, 52(6):531–550, 2011. [Cited on page 42]

[59] M. V. Berry. Quantal phase factors accompanying adiabatic changes. *Proceedings of the Royal Society of London, Series A*, 392(1802):45–57, 1984. [Cited on page 42]

[60] P. J. Lee, K.-A. Brickman, L. Deslauriers, P. C. Haljan, L.-M. Duan, and C. Monroe. Phase control of trapped ion quantum gates. *Journal of Optics B*, 7(10):S371–S383, 2005. [Cited on pages 43, 48, and 89]

[61] P. Carruthers and M. M. Nieto. Coherent states and the forced quantum oscillator. *American Journal of Physics*, 33(7):537–544, 1965. [Cited on page 43]

[62] J. Home. *Entanglement of Two Trapped-Ion Spin Qubits*. PhD thesis, University of Oxford, 2006. [Cited on pages 43 and 104]

[63] W. Magnus. On the exponential solution of differential equations for a linear operator. *Communications on Pure and Applied Mathematics*, 7(4):649–673, 1954. [Cited on page 44]

[64] S. Blanes, F. Casas, J. A. Oteo, and J. Ros. The Magnus expansion and some of its applications. *Physics Reports*, 470(5-6):151–238, 2009. [Cited on page 44]

[65] A. Sørensen and K. Mølmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A*, 62(2):022311, 2000. [Cited on page 46]

[66] I. A. Boldin, A. Kraft, and C. Wunderlich. Measuring anomalous heating in a planar ion trap with variable ion-surface separation. *Physical Review Letters*, 120(2):023201, 2018. [Cited on pages 52 and 62]

[67] Q. A. Turchette, C. J. Myatt, B. E. King, C. A. Sackett, D. Kielpinski, W. M. Itano, C. Monroe, and D. J. Wineland. Decoherence and decay of motional quantum states of a trapped atom coupled to engineered reservoirs. *Physical Review A*, 62(5):053807, 2000. [Cited on page 53]

[68] X. R. Nie, C. F. Roos, and D. F. V. James. Theory of cross phase modulation for the vibrational modes of trapped ions. *Physics Letters A*, 373(4):422–425, 2009. [Cited on page 54]

[69] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Physical Review A*, 75(4):042329, 2007. [Cited on page 56]

[70] D. Nadlinger. Entanglement between trapped strontium ions and photons. Master's thesis, ETH Zürich and University of Oxford, 2016. [Cited on page 59]

[71] S. R. Woodrow. Linear Paul trap design for high-fidelity, scalable quantum information processing. Master's thesis, University of Oxford, 2015. [Cited on page 60]

[72] S. T. Gulde. *Experimental Realization of Quantum Gates and the Deutsch-Jozsa Algorithm with Trapped $^{40}Ca^{+}$ ions*. PhD thesis, University of Innsbruck, 2003. [Cited on pages 60 and 61]

[73] D. Leibfried. Individual addressing and state readout of trapped ions utilizing rf micromotion. *Physical Review A*, 60(5):R3335–R3338, 1999. [Cited on page 61]

[74] Q. A. Turchette, C. S. Wood, B. E. King, C. J. Myatt, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland. Deterministic entanglement of two trapped ions. *Physical Review Letters*, 81(17):3631–3634, 1998. [Cited on page 61]

[75] Magnetic response of stainless steels. Technical Report 11, Atlas Speciality Metals, 2008. [Cited on page 61]

[76] T. O'Meara, A. Zverev, and H. Blinchikoff. Realization of a filter with helical components. *IRE Transactions on Component Parts*, 9(4):213–214, 1962. [Cited on page 63]

[77] J. D. Siverns, L. R. Simkins, S. Weidt, and W. K. Hensinger. On the application of radio frequency voltages to ion traps via helical resonators. *Applied Physics B*, 107(4):921–934, 2012. [Cited on page 63]

[78] D. T. C. Allcock. *Surface-Electrode Ion Traps for Scalable Quantum Computing*. PhD thesis, University of Oxford, 2011. [Cited on page 67]

[79] 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. [Cited on page 71]

[80] R. W. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31(2):97–105, 1983. [Cited on page 74]

[81] R. V. Pound. Electronic frequency stabilization of microwave oscillators. *Review of Scientific Instruments*, 17(11):490–505, 1946. [Cited on page 74]

[82] R. J. Nourshargh. Miniaturised laser systems for ion trap quantum computing. Master's thesis, University of Oxford, 2016. [Cited on page 76]

[83] D. M. 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. [Cited on page 77]

[84] D. R. Leibrandt, M. J. Thorpe, M. Notcutt, R. E. Drullinger, T. Rosenband, and J. C. Bergquist. Spherical reference cavities for frequency stabilization of lasers in non-laboratory environments. *Optics Express*, 19(4):3471, 2011. [Cited on page 84]

[85] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, I. V. Inlek, K. W. Lee, Q. Quraishi, and C. Monroe. Coherent error suppression in multiqubit entangling gates. *Physical Review Letters*, 109(2):020503, 2012. [Cited on page 95]

[86] A. Gilchrist, N. K. Langford, and M. A. Nielsen. Distance measures to compare real and ideal quantum processes. *Physical Review A*, 71(6):062310, 2005. [Cited on pages 104 and 117]

[87] C. F. Roos, G. P. T. Lancaster, M. Riebe, H. Häffner, W. Hänsel, S. Gulde, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Bell states of atoms with ultralong lifetimes and their tomographic state analysis. *Physical Review Letters*, 92(22):220402, 2004. [Cited on page 108]

[88] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White. Measurement of qubits. *Physical Review A*, 64(5):052312, 2001. [Cited on pages 108 and 110]

[89] M. Riebe. *Preparation of Entangled States and Quantum Teleportation with Atomic Qubits*. PhD thesis, University of Innsbruck, 2005. [Cited on pages 109 and 110]

[90] Z. Hradil. Quantum-state estimation. *Physical Review A*, 55(3):R1561–R1564, 1997. [Cited on page 109]

[91] K. Banaszek, G. M. D'Ariano, M. G. A. Paris, and M. F. Sacchi. Maximum-likelihood estimation of the density matrix. *Physical Review A*, 61(1):010304, 1999. [Cited on page 110]

[92] M. S. Kaznady and D. F. V. James. Numerical strategies for quantum tomography: Alternatives to full optimization. *Physical Review A*, 79(2):022109, 2009. [Cited on page 110]

[93] A. G. White, A. Gilchrist, G. J. Pryde, J. L. O'Brien, M. J. Bremner, and N. K. Langford. Measuring two-qubit gates. *Journal of the Optical Society of America B*, 24(2):172, 2007. [Cited on pages 111 and 113]

[94] M. B. Plenio and S. Virmani. An introduction to entanglement measures. *Quantum Information and Computation*, 7(1):1–51, 2007. [Cited on page 112]

[95] W. K. Wootters. Entanglement of formation of an arbitrary state of two qubits. *Physical Review Letters*, 80(10):2245–2248, 1998. [Cited on page 112]

[96] W. J. Munro, D. F. V. James, A. G. White, and P. G. Kwiat. Maximizing the entanglement of two mixed qubits. *Physical Review A*, 64(3):030302, 2001. [Cited on pages 112 and 113]

[97] I. L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *Journal of Modern Optics*, 44-11(12):2455–2467, 1997. [Cited on page 113]

[98] J. L. O'Brien, G. J. Pryde, A. Gilchrist, D. F. V. James, N. K. Langford, T. C. Ralph, and A. G. White. Quantum process tomography of a controlled-not gate. *Physical Review Letters*, 93(8):080502, 2004. [Cited on page 113]

[99] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt. Process tomography of ion trap quantum gates. *Physical Review Letters*, 97(22):220407, 2006. [Cited on page 113]

[100] J. Fiurášek and Z. Hradil. Maximum-likelihood estimation of quantum processes. *Physical Review A*, 63(2):020101, 2001. [Cited on page 115]

[101] J. Řeháček, Z. Hradil, E. Knill, and A. I. Lvovsky. Diluted maximum-likelihood algorithm for quantum tomography. *Physical Review A*, 75(4):042108, 2007. [Cited on page 115]

[102] G. C. Knee, E. Bolduc, J. Leach, and E. M. Gauger. Quantum process tomography via completely positive and trace-preserving projection. *Physical Review A*, 98(6):062336, 2018. [Cited on page 115]

[103] M. A. Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Physics Letters A*, 303(4):249–252, 2002. [Cited on page 117]

[104] J. Emerson, R. Alicki, and K. Życzkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B*, 7(10):S347–S352, 2005. [Cited on page 119]

[105] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, 2008. [Cited on page 119]

[106] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, D. Leibfried, and D. J. Wineland. Randomized benchmarking of multiqubit gates. *Physical Review Letters*, 108(26):260503, 2012. [Cited on pages 119 and 120]

[107] C. H. Baldwin, B. J. Bjork, J. P. Gaebler, D. Hayes, and D. Stack. Subspace benchmarking high-fidelity entangling operations with trapped ions. *Physical Review Research*, 2(1):013317, 2020. [Cited on page 119]

[108] D. Gottesman. The Heisenberg representation of quantum computers. In *Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, pages 32–43, Hobart, Australia, 1998. [Cited on page 119]

[109] C. J. Ballance, V. M. Schäfer, J. P. Home, D. J. Szwer, S. C. Webster, D. T. C. Allcock, N. M. Linke, T. P. Harty, D. P. Aude Craik, D. N. Stacey, A. M. Steane, and D. M. Lucas. Hybrid quantum logic and a test of Bell's inequality using two different atomic isotopes. *Nature*, 528(7582):384–386, 2015. [Cited on pages 125 and 134]

[110] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012. [Cited on page 133]

[111] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg. Surface code quantum computing with error rates over 1%. *Physical Review A*, 83(2):020302, 2011. [Cited on page 133]

[112] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland. Spectroscopy using quantum logic. *Science*, 309(5735):749–752, 2005. [Cited on page 134]

[113] T. Rosenband, P. Schmidt, D. Hume, W. Itano, T. Fortier, J. Stalnaker, K. Kim, S. Diddams, J. Koelemeij, J. Bergquist, and D. Wineland. Observation of the $^1S_0 \rightarrow {}^3P_0$ clock transition in $^{27}Al^+$. *Physical Review Letters*, 98(22):220801, 2007. [Cited on page 134]

[114] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Frequency ratio of Al^+ and Hg^+ single-ion optical clocks; metrology at the 17th decimal place. *Science*, 319(5871):1808–1812, 2008. [Cited on page 134]

[115] 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. [Cited on page 134]

[116] I. V. Inlek, C. Crocker, M. Lichtman, K. Sosnova, and C. Monroe. Multispecies trapped-ion node for quantum networking. *Physical Review Letters*, 118(25):250502, 2017. [Cited on page 134]

[117] V. Negnevitsky, M. Marinelli, K. K. Mehta, H. Y. Lo, C. Flühmann, and J. P. Home. Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register. *Nature*, 563(7732):527–531, 2018. [Cited on page 134]

[118] 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. [Cited on page 134]

[119] C. D. Bruzewicz, R. McConnell, J. Stuart, J. M. Sage, and J. Chiaverini. Dual-species, multi-qubit logic primitives for Ca^+/Sr^+ trapped-ion crystals. *npj Quantum Information*, 5(1):102, 2019. [Cited on page 134]

[120] V. M. Schäfer, C. J. Ballance, C. J. Tock, and D. M. Lucas. Optical injection and spectral filtering of high-power ultraviolet laser diodes. *Optics Letters*, 40(18):4265, 2015. [Cited on page 136]

[121] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014. [Cited on page 136]

[122] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11(1):369–395, 2020. [Cited on page 136]

[123] 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. [Cited on page 137]