

Quantum Computing Fundamentals with Mixed Qubit Types in $^{137}\text{Ba}^+$



Jamie D. Leppard
St Edmund Hall
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Michaelmas 2025

Acknowledgements

I would like to start by extending my heartfelt thanks to every member of the Oxford Ion Trapping group, past and present. You have all been a joy to work with, and I have learnt so much. If you are considering joining the Ion Trapping group in Oxford, then I'd like to say that you will not find a better group of physicists anywhere in the world. They are world-class in both ability and character. It has been an honour getting to know and work with you all. Thank you to David Lucas for putting together such an extraordinary group of people.

My supervisor, Chris Ballance, has been a continuous source of inspiration. Your passion for physics has kept me motivated throughout the PhD. Thank you for the opportunity to do this PhD, your support, and the freedom you have given me to explore such a large range of projects over the past four years.

The ABaQuSians, past and present, have been the most central part of my time in Oxford. Ana and Fabian, you have taught me everything I know. Without your guidance and kindness, this thesis would never have come to fruition – thank you. Andres, my brother in this experimental madness, thank you for the constant laughter, day in and day out. I've learnt so much from you, and there is nobody in the world I would rather flood an optics table with. As an aside, I must also acknowledge Mr Muscle Window & Glass – specifically bottle number 1021609. In the face of much advice not to employ your cleaning prowess, our mirrors are the shiniest they've ever been. Sophie and William, thank you for trusting me as I attempted to pass on my limited knowledge – and, more importantly, for being such a joy to spend time with, in and out of the lab. I am reassured in knowing the experiment will be in your capable hands. I consider you all lifelong friends.

Thank you, David Nadlinger, for your help in the lab and in maintaining our control system. Without it, I would still be collecting data to this day. Thank you, Kaitlin Gili, for your advice and support over the years. Our college lunches formed some pivotal moments early in my PhD, and while it has not surfaced in this thesis, I learnt so much from you when working on our QML project.

I must also thank those that I subjected to reading this thesis. Raghu, thank you for your support in the lab and rivalry outside of it. Without another football fan in Oxford, I would have truly lost my mind. I (almost) hope that Arsenal win a

trophy sometime in the (distant) future for you. Mario, you made our Texas trip one to remember. We made sure that Fort Worth knows that the Europeans are top dogs when it comes to eight-ball pool. Thank you for the endless humour – ‘all fart and no poo’ still tickles me to this day. Oana, you are one of the kindest people I have ever met. Thank you for your help with the two-qubit gate work and for always contributing new ideas when we were stuck.

Of course, a PhD is not possible without support outside of the lab. Freddie, thank you for keeping me above the breadline during my PhD. Working with you has given me invaluable experience that will play a huge role in my future career. Jonathan, thank you for the consistent interest in my work over years. It served as a reminder that in return for all the hard work, our research can make an impact in the world far beyond what we see up close.

Thank you to my incredible partner, Ahana. Thank you for the endless, unwavering love and support. Thank you for the unlimited patience as the lab dragged me away from our time together. I could not have done this without you. I am so lucky to have you, and cannot wait for our next adventure in life. Thank you to Seema, Pinaki, and Abhiraj. I have always felt such a welcome part of the family; you have been such a great source of advice and guidance over the years.

Mum and Dad, thank you for always believing in me and raising me to try my best in everything I do. Thank you for the constant love and the sacrifices you have made to give me such great opportunities in life; without them, I would not be here.

Abstract

Trapped-ion systems are leading candidates for quantum computing, offering long coherence times and high-fidelity gate operations. Scaling them, however, requires advanced capabilities such as mid-circuit measurement, qubit reset, and cooling, which together enable quantum error correction. One approach is the use of mixed qubit types, for example employing different ion species for cooling and computation. Recently, the *omg* blueprint [1] has gained attention as a method for coherently converting between qubit types within a single ion species, removing the overhead of mixed-species architectures while retaining their flexibility.

In this thesis, we realise mixed qubit types in $^{137}\text{Ba}^+$ using optical and metastable qubits. We demonstrate coherent control of both, as well as their coherent interconversion, achieving error rates at the 10^{-4} level suitable for NISQ applications. To enable selective control in multi-ion chains, we integrate a novel photonic chip for individual addressing of Raman beams [2]. We achieve cross-talk levels at or below 10^{-3} across all zones, and as low as 10^{-5} in some cases, establishing the scalability of our approach.

Building on this, we present a new protocol for heralded state preparation and measurement (SPAM) [3]. This approach combines high-fidelity measurement with coherent conversion between qubit types to mitigate errors associated with qubit loss, decay, and imperfect SPAM. We demonstrate record-low SPAM error rates across optical, metastable, and ground-state qubits, with a minimum of $5(4)\times 10^{-6}$. This protocol can be integrated with erasure conversion techniques to enable in-sequence qubit loss correction, providing a pathway to scalable error correction even with finite qubit lifetimes.

Finally, we introduce two-qubit operations through an implementation of the optical-transition dipole force (OTDF) gate [4]. This is the first implementation of the OTDF gate in $^{137}\text{Ba}^+$, or in any ion with hyperfine structure, and it is directly compatible with our existing 532 nm laser system. Crucially, this gate operates in the σ_Z basis, avoiding the need for phase coherence between the driving field and the qubit, making it an excellent candidate for use in mixed-qubit-type systems.

These results demonstrate the application of mixed-qubit-type architectures within a single ion species and establish a foundation for scaling trapped-ion quantum computers.

Contents

1	Introduction	1
1.1	Quantum computing	1
1.2	Protecting qubits	5
2	Ba⁺: A quantum playground	9
2.1	Ion species	10
2.2	Choice of isotope	12
2.3	Choice of qubit	13
2.3.1	Metastable qubits	13
2.3.2	Optical qubits	14
2.3.3	Converting between qubit types	15
2.3.4	Poly-qubits	17
2.4	Summary	18
3	Photoionisation and dissipative control	20
3.1	Photoionisation	20
3.2	Fluorescence and Doppler cooling	23
3.3	EIT Cooling	24
3.4	State preparation and measurement	27
4	Single qubit operations	30
4.1	Atom-light interactions	31
4.1.1	Dipole interactions	32
4.1.2	Coupling to the motion	36
4.1.3	Quadrupole interactions	38
4.1.4	Raman interactions	39
4.2	Implementing single qubit operations	41
4.2.1	Optical qubit	41
4.2.2	Metastable qubit	44
4.3	Calibrating single qubit operations	47
4.3.1	Ramsey spectroscopy	48

4.3.2	Robust phase estimation	49
4.4	Randomised benchmarking	56
4.4.1	Metastable qubit RBM	57
4.4.2	Optical qubit RBM	62
4.5	Converting between qubit types	69
4.5.1	Clock qubit	70
4.5.2	Hybrid qubit	71
4.5.3	Maintaining coherence	71
4.6	Summary	74
5	Individual addressing using a novel integrated photonic chip	76
5.1	Methods of individual addressing	77
5.2	Addressing apparatus	80
5.3	System characterisation	82
6	High-fidelity heralded state preparation and measurement	85
6.1	Introduction	86
6.2	Taxonomy of errors	87
6.3	The protocol	89
6.3.1	State preparation	89
6.3.2	Measurement	90
6.4	Scaling of data rates and errors	91
6.4.1	State preparation	91
6.4.2	Measurement	94
6.5	Observable Statistics	95
6.6	Implementation in $^{137}\text{Ba}^+$	98
6.6.1	Metastable level qubit	99
6.6.2	Optical qubit	100
6.6.3	Ground qubit	101
6.7	Results	103
6.7.1	Optical detection	103
6.7.2	Metastable level lifetime	104
6.7.3	Protocol results	106
6.7.4	Observable bias	107

7	OTDF Gates	112
7.1	Motivating an alternative two-qubit gate	113
7.2	Spin-dependent forces	115
7.2.1	Background	115
7.2.2	Generating a $\hat{\sigma}_z$ force	116
7.3	Entering the Lamb-Dicke regime	120
7.3.1	Sideband cooling	121
7.3.2	Sideband thermometry	122
7.4	Spin-dependent forces with a single ion	124
7.4.1	Phase tracking the motion	125
7.5	Choice of polarisation	128
7.6	Entanglement generation	129
7.7	Gate setup & implementation	131
7.7.1	Heating rate and motional coherence time	132
7.7.2	Mode calibration	134
7.7.3	The spectator mode	135
7.7.4	Balancing light shifts	137
7.8	Results and sources of error	137
7.8.1	Results	137
7.8.2	Sources of error	138
7.9	Summary	152
8	Outlook	154
Appendices		
A	Motional coherence times	158
B	Optical clock qubits	160
List of Figures		162
References		165

This is it, don't get scared now

— Kevin McCallister, Home Alone

1

Introduction

Contents

1.1 Quantum computing	1
1.2 Protecting qubits	5

1.1 Quantum computing

The classical computing revolution brought with it a series of paradigm shifts that transformed both modern science and daily life to a degree many would have considered unimaginable. Across scientific disciplines, a key shift has been in the ease with which one can construct a simulation of a system or idea. Previously limited to physical models or analogue electronics, simulations in the digital era allowed a machine of sufficient capabilities to simulate almost any classical system. Decades after its advent, advances in data processing have propelled the field of machine learning, again unlocking access to complex problems which can now be solved on a machine instead of physically in the laboratory. One of the most impactful of these machine learning models is AlphaFold [5], Google DeepMind's

model for finding the natural structure of proteins as well as predicting the structure of various protein-based interactions [6].

However, classical computing has a limit. It is only capable of efficiently simulating classical systems, but we know nature to be quantum. The number of parameters required to classically simulate an arbitrary quantum system scales exponentially, leading to a shortage in computational resources for a system of more than ~ 50 entangled spins [7]. This exponential scaling leaves us with both a problem and an opportunity - can we construct a system that can efficiently simulate *all* physical systems? Such a device could create another paradigm shift, at first for the scientific world, which may then echo to the wider global community. The issue was stated memorably by Richard Feynman in the 1980s in his keynote speech¹ in which he said “*And I’m not happy with all the analyses that go with just the classical theory, because nature isn’t classical, dammit, and if you want to make a simulation of nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem, because it doesn’t look so easy.*” [8].

We now commonly refer to the devices Feynman envisaged as “quantum computers” or “quantum simulators”. Though the first proposals for a quantum computer were made more than 40 years ago by Richard Feynman [8, 9], Paul Benioff [10], and then later by David Deutsch [11], we have not yet constructed one which is generally useful. The path to utility was succinctly described by DiVincenzo’s criteria [12]. These are summarised as follows

1. A scalable system of well characterised qubits.
2. The ability to prepare qubits in some fiducial state, e.g. $|\downarrow\rangle^{\otimes}$.
3. Long, relevant coherence times compared to the time required to perform operations.
4. The ability to perform a universal gate set.

¹Though he said himself “On the program it says this is a keynote speech – and I don’t know what a keynote speech is.”

5. The ability to perform qubit specific measurement.

Note that he does not quote a specific size of system for utility, but that if the system is scalable, then it should be possible to create a system of the relevant size for a given problem. There are two additional criteria regarding “flying” qubits for quantum communication, specifically to be able to interconvert stationary and flying qubits, as well as faithfully transmit flying qubits between locations. This idea of flying qubits points towards a potential quantum internet [13], in which many quantum computers are networked together allowing for advanced protocols such as distributed quantum computing [14, 15] and quantum cryptography [16] which have been demonstrated here at the University of Oxford.

An enormous amount of effort has been applied to meet each of these criteria across a variety of platforms, including superconducting qubits [17, 18], NV-centres [19], neutral atoms [20, 21], trapped ions [22] (the topic of this thesis), as well as proposals to combine multiple platforms [23]. Each platform has its own set of advantages and disadvantages, the extent and detail of which could fill an entire thesis.

In this thesis we focus on trapped atomic ions. Atomic platforms offer us qubits provided by nature. In this regard, all qubits are identical and therefore operations are reproducible as a system scales. This platform holds records in all categories regarding the fidelity of qubit operations [24, 25] and coherence times [26]. However, scaling trapped-ion systems to hundreds of qubits poses a challenge that has not yet been solved and is critical to achieving practical utility. This limitation is particularly significant because many quantum algorithms only provide an advantage when executed at scales far beyond the reach of classical systems, such as Shor’s algorithm [27] for prime factorisation.

This quest for larger scale devices has lead us to conceptualise a variety of “architectures” in trapped-ion quantum computing which utilise and interact with the quantum states within ions in different ways. The two most common architectures ideated in the literature are:

- The ‘QCCD’ architecture [28, 29] which aims to scale the complexity of the trap chip, increasing the number of trapping zones in which quantum logic operations can be performed. This enables the control of many small chains of ions in parallel, each of which has a simple and manageable structure.
- A networked architecture [30–32], enhanced by optical cavities [33]. Past a given limiting number of qubits, devices can be connected via flying qubits to network many small scale devices together. This then places a focus on the additional sixth and seventh of DiVincenzo’s criteria in order to achieve the first.

In each of these cases, the choice of architecture can make significant headway on all of DiVincenzo’s criteria. However, there are still decisions to be made. In particular, the 5th criterion to perform qubit specific measurement has very far reaching implications. As a system scales, the need for efficient error correcting codes [34, 35] becomes inevitable [36], even with the excellent error rates demonstrated in state of the art systems [24, 25]. Such error correction requires that per ‘cycle’ of quantum logic (a set of single- and two-qubit operations), a set of error correcting operations must be run over the system to prevent errors from propagating uncontrollably. Each error correcting cycle will involve the *qubit specific measurement* of many auxiliary qubits. This brings us back to the choice remaining within the QCCD and networking architecture - how to implement qubit specific measurement. As with any complex system, complying with a new constraint results in a trade-off of different performance metrics. For example, an implementation that requires sequential measurement of each ion would incur a large time overhead in error-correction cycles. If coherence times are only just sufficient for a given use case, this could be the straw that breaks the camel’s back, degrading performance and, thus, utility.

1.2 Protecting qubits

In most ion systems measurement is performed by checking for fluorescence in an optical qubit [37] (encoded across the $S_{1/2}$ and $D_{5/2}$ manifolds). Bright implies projection into the $S_{1/2}$ level, and dark the $D_{5/2}$ level. This is a strong dissipative coupling which excites both the electronic and motional state of the ion. As such, there are two avenues for errors. Any coupling to the electronic state of an ion we do not wish to measure will destroy the quantum information within it. Some selective coupling must therefore be implemented. Secondly, if the measured ion is in a chain with others, the shared motional state means all ions will be excited in the motion. As a result, the damage to performance done through the measurement of our qubits extends beyond just the electronic state of the ions. If the measurement process were to heat the ions that were not subsequently cooled, the increased error rate in our two-qubit operations (which are mediated by the shared ion motion) would again break our quantum camel's back. To account for this, we add an additional criterion: *Logic gates with Markovian² error rates to enable deep quantum circuits.*

For state preparation and measurement (SPAM), a strong coupling to the environment is required to dissipatively prepare the ion into a known electronic configuration, and then to retrieve the configuration from the ion at the end of a quantum circuit. To achieve Markovian error rates, the strong coupling is required to prepare and maintain the motional state of the ions in a known Fock state, (typically $|n = 0\rangle$). Maintenance of the motional state is almost always required due to the inevitable heating that will occur during a quantum circuit if no further cooling is applied, aside from any excess dissipative coupling. This also means that resetting an ion's motional state is required irrespective of whether the ion is currently home to our delicate quantum information. So, we have an engineering challenge: *how do we protect quantum information from this strong coupling to the environment?*

²By Markovian we mean constant. That the error rate at the start of a circuit is the same as after the thousandth and millionth gates.

All approaches to this challenge are centred around having multiple modes of operation for the ions. By placing the ion in a state where it is insensitive to the light field driving operations on another, we can protect qubits from global operations which would destroy the quantum information they hold. We can do this with lasers used for readout and preparation of the qubit, allowing for qubit specific measurement and preparation. The most direct approach would be to target the ions we wish to measure with the measurement laser, and avoid ions that we do not wish to measure. The engineering effort involved in doing this is large, particularly in a QCCD approach which would then require a set of integrated optics components per qubit. A compromise would be to create measurement “zones” that ions are shuttled through, trading off engineering effort and parallelisability. However, such an approach could hinder system performance due to the large amount of time spent shuttling ions [38]. This allows for all forms of selectivity with one exception, preparing an active ion in a cold motional state.

To cool an ion which is presently storing quantum information, it must be sympathetically cooled by an auxiliary coolant ion [39]. Without individual addressing of the cooling lasers, which would compound the engineering effort, this makes the target ion susceptible to the light field applied to the coolant ion and, thus, the error rates non-Markovian. Such a restrictive requirement for implementing deep quantum circuits leads us to search for configurations of ions which can protect quantum information, even when laser light is applied directly to the ion storing this information.

One solution is to house two different atomic species in the same ion trap. The two species are sensitive to different wavelengths of laser light, thus, each is protected from operations on the other [40, 41]. With one species acting as coolant, and the other processing data, we can keep error rates Markovian. This two-species approach has further benefits: rather than moving ions in and out of light fields to protect their quantum state, we can map quantum information from one species to another,

protecting the information during state preparation and measurement operations. This reduces the engineering effort required to begin scaling our quantum system.

Unfortunately, this is not a free lunch. Adding a second atomic species to a system introduces a whole host of different fundamental and technical engineering challenges. Primarily, the two atomic species have different atomic masses. This mismatch leads to issues in sympathetic cooling due to unequal mode participation. It also increases the complexity of two-qubit gate operations, and ion transport [42–44]. Additionally, the extra atomic species requires an entirely different set of laser systems, adding significantly to the overhead of classical control hardware [1].

In this thesis we explore an alternative approach to the multi-species system by employing multiple coherently convertible qubit types within a single ion species [45]. This method is now commonly known as the *omg* blueprint [1], which utilises long-lived metastable states of trapped ions to gain the same utility of the mixed-species approach, with a much reduced overhead. The idea is to hide qubits in a metastable level, commonly the $D_{5/2}$ in Ca^+ , Sr^+ and Ba^+ ; or $F_{7/2}$ in Yb^+ [1], to exclude them from the $S_{1/2} \rightarrow P_{1/2}$ fluorescence cycle. Instead of selectivity being driven by ion species, we put our efforts into the selective shelving of quantum information to the metastable level, a trade-off that plays nicely with the pre-existing requirement for qubit selective quantum gates.

In this work, we focus on realising the *omg* blueprint in Ba^+ , demonstrating how multiple qubit types within a single ion species can be harnessed for scalable quantum information processing. The chapters that follow build progressively from the choice of qubits and the tools required to manipulate them, through to demonstrations of high-fidelity state preparation and measurement (SPAM), low cross-talk single-qubit addressing and single- and two-qubit operations. A brief outline of the thesis is given below.

Chapter 2: We motivate the use of barium, and in particular $^{137}\text{Ba}^+$, for use within the *omg* scheme. We then examine $^{137}\text{Ba}^+$'s hyperfine structure through simulation, identifying possible optical and metastable qubit choices. From these,

we select two optical–metastable combinations that are carried forward through the remainder of the thesis.

Chapter 3: We outline the fundamental dissipative operations used for trapping and cooling $^{137}\text{Ba}^+$. Variations of these operations form the backbone of quantum control, enabling state preparation and measurement.

Chapter 4: With the choice of qubits established, we investigate the implementation of coherent operations in each qubit type. This includes the calibration of parameters required for high-fidelity operation. We then benchmark their error rates and assess the accuracy with which one qubit type can be converted into another.

Chapter 5: In this chapter we demonstrate the integration of a novel photonic chip used for individual addressing of our Raman lasers. This enables selective qubit operations, a prerequisite for any useful quantum system. We demonstrate cross-talk levels of 10^{-3} or lower across the chip, which are sufficiently low to avoid being a limiting factor in many quantum systems.

Chapter 6: Next, we utilise our single qubit operations along high fidelity measurement to implement a novel protocol for heralded state preparation and measurement (SPAM). Using this protocol, we demonstrate the lowest SPAM error rates across all three qubit types - optical, metastable, and ground.

Chapter 7: We introduce two-qubit operations with an implementation of the optical-transition dipole force (OTDF) gate, a fundamental building block in scaling our trapped ion device. Possible sources of error are outlined along with next steps to improve the implementation.

Chapter 8: Conclusions are drawn and the next steps for the project are outlined.

Sometimes science is more art than science ...

— Rick Sanchez, Rick and Morty

2

Ba⁺: A quantum playground

Contents

2.1	Ion species	10
2.2	Choice of isotope	12
2.3	Choice of qubit	13
2.3.1	Metastable qubits	13
2.3.2	Optical qubits	14
2.3.3	Converting between qubit types	15
2.3.4	Poly-qubits	17
2.4	Summary	18

There are many choices we can make with respect to the levels we use within an atom to make up our qubit (or qudit). Research groups around the world are using different atomic species and levels within those species across ground, optical and metastable qubit. As such, the scientific community has not settled on a single best choice of qubit¹.

With that said, there are certainly choices we can make that will bring us closer to achieving the criteria set out by DiVincenzo, with very little engineering overhead arising from these choices. We investigate these low-cost choices from three

¹Ideally we never do as such a choice would indicate that there is nothing better to be found and progress in this domain would halt.

perspectives: an optical qubit-based setup, a metastable qubit-based setup, and, more briefly, a polyqubit/qudit setup. As touched upon in the previous chapter, both qubit types are required in an advanced quantum system for mid-circuit measurements [46]. The optical qubit is used for measurement, and the metastable qubit to protect active qubits from the measurement process. This is illustrated in Fig. 2.1, where the fluorescence cycle couples to the S_{1/2} level, which contains both ground-state qubit levels and corresponds to one of the optical qubit states. However, this cycle does not couple to the D_{5/2} metastable level at all.

The exact choice of how these qubits are utilised in coherent operations can have far reaching consequences to the fidelities easily achievable in a system (all systems can technically be improved with increased engineering effort) and therefore its capabilities.

2.1 Ion species

When selecting an ion species for quantum information processing, it is essential to balance long coherence times, convenient laser requirements, and compatibility with advanced qubit control schemes. Barium offers many advantages over other atomic species when implementing the *omg* scheme. In particular, the long-lived D_{5/2} manifold has a 30 s lifetime [1], more than sufficient for quantum operations, which are typically on the microsecond timescale. By comparison, calcium and strontium offer metastable lifetimes of only 1.2 s and 300 ms, respectively, thus limiting the qubit lifetime [47]. While ytterbium has exceptionally long metastable lifetimes of days to years, it requires an ultraviolet laser at 369 nm for fluorescence. Such laser sources are typically lower power, and the corresponding optical components are more expensive to manufacture at high quality compared to their visible counterparts. Additionally, transferring population to the F_{7/2} level requires either a two stage process via the D_{5/2} or driving the extremely narrow octupole transition with nHz linewidth [48]. Barium strikes a middle ground: it provides long-lived metastable states while keeping all but one laser in the visible spectrum, the exception being the

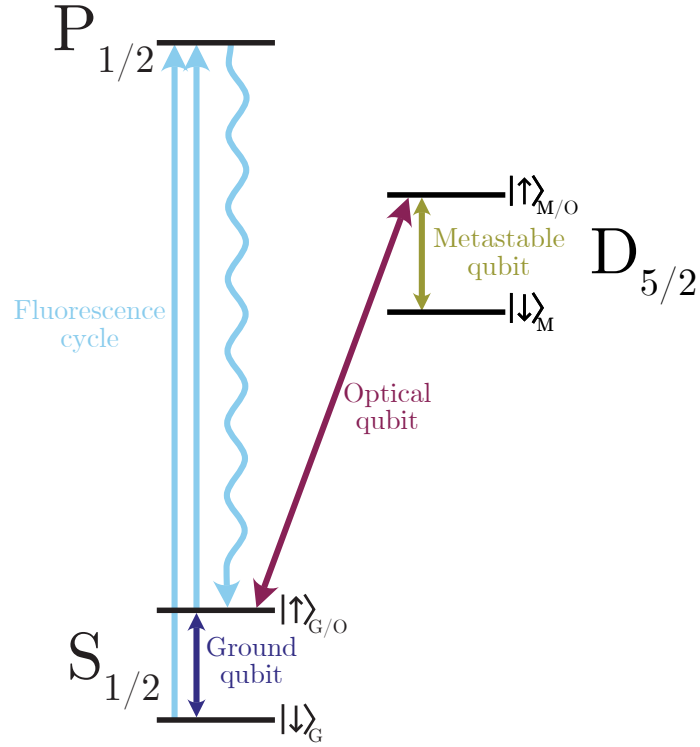


Figure 2.1: Level diagram of a typical trapped ion system implementing ground, optical, and metastable qubits. The ground qubit is defined entirely in the $S_{1/2}$ manifold and both states couple to the fluorescence cycle. The optical qubit is defined by one state in the ground $S_{1/2}$ level and one state in the metastable $D_{5/2}$ level. In this configuration a measurement can be performed by checking for fluorescence. Alternatively, an ion can be protected from measurement by storing information in a metastable qubit, stored entirely in $D_{5/2}$, which does not couple to this fluorescence cycle.

1762 nm laser driving the $S_{1/2}$ to $D_{5/2}$ transition, which is also direct. The absence of deep-blue or ultraviolet lasers greatly increases the efficiency and ease of use of the laser systems. Barium also offers a plethora of low-field optical and metastable clock qubits in isotopes with hyperfine structure. At particular magnetic fields, such transitions are first-order insensitive to magnetic-field fluctuations, almost completely removing a major source of decoherence [47]. Taken together, these features make barium an especially attractive platform for realising the *omg* blueprint in a scalable quantum computer.

2.2 Choice of isotope

With the choice of barium as our ion species established, we must now determine which isotope to use. The options available are summarised in Ref. [47], but in practice we find that there are only three sensible choices: ¹³⁸Ba⁺, ¹³⁷Ba⁺ or ¹³³Ba⁺. ¹³⁸Ba⁺ has no nuclear spin, giving it a simple level structure, but no access to magnetically insensitive clock qubits. Additionally, the Zeeman splitting in the D_{5/2} manifold is uniform across states, making an implementation of qudits or poly-qubits in just the D_{5/2} manifold technically challenging due to the lack of frequency selectivity.

All naturally occurring isotopes of barium with non-zero nuclear spin have nuclear spin $I=3/2$. Consequently, they have the same hyperfine and Zeeman structure, with slight variations in the size of the hyperfine splitting. This means there is no real choice to be made between them and so we choose the most abundant naturally occurring isotope ¹³⁷Ba⁺.

This leaves the radio-isotope ¹³³Ba⁺ which has a half-life of ~ 10.5 years, but nuclear spin $I=1/2$. The lower nuclear spin provides a much simpler level structure, which is advantageous for Doppler cooling and state preparation [49, 50]. However, when implementing a qudit in the metastable level of ¹³³Ba⁺, we only have access to two hyperfine levels, meaning some states must be separated by just Zeeman splitting. Therefore the transitions have a relatively high magnetic field sensitivity compared to options in ¹³⁷Ba⁺. There are also more metastable clocks to choose from in ¹³⁷Ba⁺ vs ¹³³Ba⁺ with the only choices we have found in ¹³³Ba⁺ being $|D_{5/2}, F=2, m_F=+1\rangle \leftrightarrow |D_{5/2}, F=3, m_F=+1\rangle$ at 16.473 G and $|D_{5/2}, F=2, m_F=+2\rangle \leftrightarrow |D_{5/2}, F=3, m_F=+2\rangle$ at 32.951 G. Additionally, as mentioned in Ref. [47], working with small $\sim 100 \mu\text{g}$ atomic sources to meet radiation safety requirements posed the challenge of using ablation loading instead of a conventional atomic oven. Unfortunately, we were not able to reproduce an ablation target reliably enough to use ¹³³Ba⁺. Other research groups have since put effort into

creating ablation targets [51, 52], though it remains a challenge for individual labs to recreate.

2.3 Choice of qubit

Having chosen the ion species and isotope, we now turn to the selection of atomic states that form our qubits. In this section, we provide an overview of the available metastable and optical qubit options, as well as possible pairings within a single system. We then pursue two of these choices in Chapter 4, focusing on single-qubit operations in conjunction with the optical transition dipole force (OTDF) gate, as detailed in Chapter 7. The simulations in this section, carried out using the Python package of Ref. [53], are used to calculate the properties of the different transitions.

2.3.1 Metastable qubits

Much of the background regarding metastable qubits, in particular clock qubits, has been covered in my predecessor's thesis (Ref. [47]). Therefore, this section will be a brief overview. Metastable qubits in ¹³⁷Ba⁺ offer a range of low-field clock transitions which can exhibit coherence times of 10+ s, and are fundamentally limited by the 30 s lifetime of the D_{5/2} manifold. Such long coherence times are possible as these qubits are driven with either microwave radiation or a two-photon Raman process, neither of which is sensitive to common mode laser phase noise. In this work and previous work we choose to utilise the 3.333 G clock between $|D_{5/2}, F = 3, m_F = +2\rangle$ and $|D_{5/2}, F = 2, m_F = +2\rangle$. There are many other choices available as outlined in Appendix A of Ref. [47], however the 3.333 G qubit is an excellent choice for the following reasons:

- The low 3.333 G field is very easy to generate experimentally, and keeps Zeeman levels in the ground close together, making Doppler cooling easier.
- The nearest spectator is ~ 2.4 MHz away, meaning off-resonant excitation will be minimal for gates on the timescale of 10 μ s, fast enough for our application.

S _{1/2} state	D _{5/2} state	Field insensitive point (G)	$d^2 f/dB^2$ (kHz/G ²)	Matrix element Cm/h
1, 0⟩	2, 0⟩	1.3471	0.2	0.139
2, 0⟩	3, -1⟩	1.6353	26.4	0.794
1, 0⟩	3, -1⟩	1.6971	26.8	1.31
2, 0⟩	2, 0⟩	4.1796	2.1	1.09
2, -2⟩	4, 0⟩	4.6159	35.7	0.269
2, -2⟩	3, -2⟩	8.5406	35.2	2.52
1, -1⟩	1, 0⟩	8.8244	72.6	2.43
2, 1⟩	1, 0⟩	8.9014	71.7	0.454

Table 2.1: Transition properties of optical clock qubits between 1 G and 10 G in ¹³⁷Ba⁺. Atomic states are written at $|F = i, m_F = j\rangle = |i, j\rangle$.

All metastable qubits in ¹³⁷Ba⁺ can be driven using 532 nm Raman transitions, making them easily addressable which allows us to scale our device to many qubits. This is discussed further in Chapter 5. One benefit of using purely metastable qubits for quantum information processing is that they are natively insensitive to our measurement process. As a result, the metastable qubit only requires mapping to an optical qubit when a measurement is to be performed. Thus, there is no requirement to maintain qubit coherence in the mapping process - we only measure qubit populations. This is not true in an optical-centric approach as coherence must be maintained when moving to a metastable configuration to protect a qubit from measurement. We will see in the next section how this adds some extra considerations to our approach.

2.3.2 Optical qubits

Another option is to perform all of our quantum information processing in the optical qubit and use a metastable qubit simply for storage, as illustrated previously in Fig. 2.1. If we were to take such an approach, which optical qubit would be best? Firstly, we can look for available optical clock transitions. As with the metastable clocks, there are a plethora of options to choose from at low field (we choose the cutoff to be <10 G) which are summarised in Table 2.1. Among the choices in Table 2.1, a few choices can be eliminated, in particular those with small matrix elements (<1 Cm/h) as they will exhibit particularly low Rabi frequencies. There is then a convenient

choice of optical qubit: $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$. Although the matrix element is lower than many other options, it is suitable for our state preparation scheme using π -polarised light outlined in Sec. 3.4. Additionally, the second order field sensitivity of this qubit is an order of magnitude lower than all other choices (except for the very slow $|S_{1/2}, F = 1, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ clock at 1.3471 G).

However, unlike the metastable transitions, our optical qubit cannot be driven with either a two-photon Raman process, or microwave radiation. Thus, the qubit is sensitive to common mode noise from the driving laser. This places a technical limitation on the coherence time of the qubit, irrespective of how insensitive it is to magnetic field noise. Even in state of the art atomic clock systems, Ramsey durations are limited to a few seconds [54, 55].

Given this limitation, and the requirement to have a metastable qubit with reasonable coherence time for storage, we instead turn our attention to the metastable qubit. More specifically, how we can convert between the metastable qubit and optical qubit.

2.3.3 Converting between qubit types

When working with multiple qubit types, we must be able to coherently move information between different configurations, in this instance between optical and metastable configurations². The challenge here is choosing choosing compromises between qubit encodings and pairings that both individually are good, while also can be efficiently and coherently mapped. One option would be to take the greedy approach, using both the metastable clock at 3.333 G and an optical qubit of low field sensitivity. In Table 2.2 we show the lowest sensitivity quadrupole transitions at 3.333 G. From this there appears to be a clear winner which is once again $|S_{1/2}, F = 2, m_F = 0\rangle \rightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ which has a first order sensitivity of

²Note that when working a metastable-centric system, coherent conversion isn't required since qubits need only be mapped to optical configurations for measurement. However, in the hybrid optical/metastable system we are working with, this coherent conversion is a must for protecting qubits from the destructive measurement process.

S _{1/2} state	D _{5/2} state	dF/dB (kHz/G)	Matrix element Cm/h
$ S_{1/2}, F = 2, m_F = 0\rangle$	$ D_{5/2}, F = 2, m_F = 0\rangle$	1.31	1.10
$ S_{1/2}, F = 1, m_F = 0\rangle$	$ D_{5/2}, F = 2, m_F = 0\rangle$	1.95	0.342
$ S_{1/2}, F = 2, m_F = -2\rangle$	$ D_{5/2}, F = 4, m_F = 0\rangle$	47.2	0.274
$ S_{1/2}, F = 1, m_F = 0\rangle$	$ D_{5/2}, F = 3, m_F = -1\rangle$	56.9	1.29
$ S_{1/2}, F = 2, m_F = 0\rangle$	$ D_{5/2}, F = 3, m_F = -1\rangle$	60.1	0.939
$ S_{1/2}, F = 2, m_F = -2\rangle$	$ D_{5/2}, F = 2, m_F = -1\rangle$	117	1.46

Table 2.2: Magnetic field sensitivity dF/dB and matrix elements for transitions between hyperfine sublevels of the S_{1/2} and D_{5/2} states at 3.333 G.

just ~ 1.31 kHz/G. As a point of comparison for how poor a choice we could make, the largest magnetic field sensitivity is ~ 4.9 MHz/G in the $|S_{1/2}, F = 1, m_F = 1\rangle \rightarrow |D_{5/2}, F = 3, m_F = 3\rangle$ transition, a factor of $\sim 3700\times$ more sensitive.

If we were to choose this metastable and optical qubit pairing, we must then consider how we map between the optical and metastable qubits. Unfortunately there is a large downside to this. The qubits do not share any states, and we cannot map directly from either of the metastable qubit states to the $|D_{5/2}, F = 2, m_F = 0\rangle$ state without resorting to a magnetically sensitive quadrupole transition. The details of this mapping are provided in Sec. 4.5 where we characterise the error from such mapping pulses. In summary, the requirement to transition via magnetically sensitive qubits can harm qubit coherence and fidelity greatly, causing us to falter on one the most important of DiVincenzo's criteria.

An alternative, and less greedy, approach would be to look for a metastable and optical qubit which can be easily and directly interconverted and that both exhibit relatively low magnetic-field sensitivity. In this instance, we must consider the first-order sensitivity of both the metastable and the optical qubit. We do so in Table 2.3 for a magnetic field of 3.333 G, we also include the nearest neighbour to the metastable qubit. While the metastable $|D_{5/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 3, m_F = -1\rangle$ qubit and optical $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ would be very stable, the close by $|D_{5/2}, F = 4, m_F = +1\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ qubit limits the speed with which we can drive it before off-resonant excitation becomes an issue. The only

M Qubit	O Qubit	M dF/dB (kHz/G)	O dF/dB (kHz/G)	Nearest M qubit Distance (MHz)
$ 2, 0\rangle \leftrightarrow 3, -1\rangle$	$ 2, 0\rangle \leftrightarrow 2, 0\rangle$	59	1.31	$ 4, +1\rangle \leftrightarrow 2, 0\rangle$ (0.46)
$ 2, 0\rangle \leftrightarrow 4, +1\rangle$	$ 2, 0\rangle \leftrightarrow 2, 0\rangle$	173	1.31	$ 3, -1\rangle \leftrightarrow 2, 0\rangle$ (0.46)
$ 2, -1\rangle \leftrightarrow 4, 0\rangle$	$ 2, -2\rangle \leftrightarrow 4, 0\rangle$	164	47.2	$ 3, -2\rangle \leftrightarrow 2, -1\rangle$ (0.28)
$ 1, 0\rangle \leftrightarrow 2, 0\rangle$	$ 2, 0\rangle \leftrightarrow 2, 0\rangle$	274	1.31	$ 2, -1\rangle \leftrightarrow 1, 0\rangle$ (5.1)
$ 2, -1\rangle \leftrightarrow 3, -2\rangle$	$ 2, -2\rangle \leftrightarrow 2, -1\rangle$	272	116.7	$ 4, 0\rangle \leftrightarrow 2, -1\rangle$ (0.28)

Table 2.3: Metastable and optical qubit transitions with their magnetic field sensitivities and nearest-neighbour metastable qubits.

qubit combination which meets our criteria at 3.333 [G] is $|D_{5/2}, F = 1, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ and $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$. This holds for all fields up to 10 G with one exception: the $|S_{1/2}, F = 1, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ optical clock at 1.3471 G, which suffers from particularly slow transitions due to the low matrix element as discussed before.

2.3.4 Poly-qubits

Another active area of research in our lab is utilising higher dimensional quantum systems to increase the size of the Hilbert space we have available for a given number of ions. Systems are often restricted to the typical two-level qubits for the sake of simplicity, as well as access to clock transitions. Relaxing this constraint opens up the possibility of using many more states within each ion. This is something that our team has worked on theoretically [56] and is now breaking into experimentally. Much of the literature on this focusses on using qudits [57–60], but we choose to work with virtual qubits, or poly-qubits, where a higher dimensional system is mapped to many logical qubits [56, 61]. The difference between the two is not in hardware, but in how the Hilbert space is utilised. Poly-qubits encode multiple qubits in a single ion, changing how quantum algorithms are implemented, as poly-qubits thus work with existing

qubit based algorithms. This re-framing of the Hilbert space unlocks a different set of quantum operations. For example, an operation between two qubits and an auxiliary ion used for parity measurement can be done with just one entangling gate [56].

The details of implementing poly-qubits in ¹³⁷Ba⁺ will be left to the thesis of Andres Vazquez-Brennan. As a brief overview/teaser, choosing a set of states with which to form a qudit or a poly-qubit can be seen as an extension of the compromises we have made above. In our case we have chosen three levels for convenient mapping between different qubit types, this is similar to optimising for a qutrit. Extending to four levels for a two-poly-qubit ion requires further compromise and consideration of how connecting transitions will be utilised. As before there is no best qudit, but I hope Andres can convince you of our choices.

2.4 Summary

I must place one final emphasis that *there is no best qubit*. However, through careful consideration we have uncovered a few options with which to explore. The choices within ¹³⁷Ba⁺ that we will carry forward with us into the next few chapters are as follows³.

- A metastable clock at 3.333 G between $|D_{5/2}, F = 3, m_F = +2\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = +2\rangle$. This can be used stand-alone with Mølmer-Sørensen gates for two-qubit interactions with only a clock qubit for measurement. Alternatively it connected to an optical qubit using complex mapping sequences.
- An optical clock at 4.1796 G between $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$. This is compatible with OTDF gates which we detail in Chapter 7. The optical qubit is susceptible to phase noise from its driving laser, and so below a given sensitivity there is no added benefit. It can be

³In some chapters data was taken before such choices and optimisations were made, in particular chapters 5&6. Since we do not aim for high-fidelity qubit operations in these chapters the choice makes very little difference.

connected to the metastable $|D_{5/2}, F = 1, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$ for storage while protecting from readout lasers in the *omg* scheme.

Are you trapping ions, or are they trapping you?

— Every ion trapper, at some point in their career

3

Photoionisation and dissipative control

Contents

3.1 Photoionisation	20
3.2 Fluorescence and Doppler cooling	23
3.3 EIT Cooling	24
3.4 State preparation and measurement	27

In this chapter, we discuss the use of each laser involved in trapping and cooling $^{137}\text{Ba}^+$. The lasers used in our ion trap apparatus are illustrated in Fig. 3.1(a) along with the levels that they couple to in Fig. 3.1(b/c). This includes the 1762 nm quadrupole laser and the 532 nm Raman laser, which will be discussed in Chapters 4&7. In particular, we describe the tones applied to match resonances within the ion, as well as the relative geometry to the magnetic field (also shown in Fig. 3.1), which is necessary to achieve a particular polarisation at the ion.

3.1 Photoionisation

We load $^{137}\text{Ba}^+$ from a neutral atomic barium source. The neutral barium is ionised via a two-stage photoionisation process shown in Fig. 3.1(b): 553 nm light drives the $6s^2 \rightarrow 6s6p$ transition, and a 405 nm laser ejects the electron from the atom




Figure 3.1: (a) Diagram of beam geometry. The 493 nm (Doppler), 650 nm, and 614 nm beams are all 90° to the B-field. The 493 nm (Doppler) beam has linear polarisation such that interactions with the ion are π -polarised. The $|S_{1/2}, F = 1\rangle \rightarrow |P_{1/2}, F = 1\rangle$ sideband is turned off when performing state-preparation or EIT cooling, allowing the Doppler beam to optically pump population into $|S_{1/2}, F = 2, m_F = 0\rangle$. The 493 nm (EIT) beam enters from the other side of the trap creating a larger Δk with the pump beam along the radial modes. There is a slight skew to the nominally 45° angle providing weak projection along the axial mode to cool these as well. The photoionisation lasers (405 nm and 553 nm) are counter-propagating to the Doppler lasers for ease of setup and initial trapping. (c) Neutral barium level diagram for two stage photoionisation. (c) $^{137}\text{Ba}^+$ level diagram and corresponding lasers. Details of the 532 nm and 1762 nm are provided in Sec. 7.5 and Sec. 4.2.1 respectively.




Figure 3.2: Neutral fluorescence resonance. The 553 nm is detuned from the ^{138}Ba peak. We adjust the laser frequency by tuning the 1106 nm laser frequency in software. The main peak is ^{138}Ba , the inset shows the ^{135}Ba and ^{137}Ba resonance at 540 MHz predicted in Ref. [65].

as described in Ref. [62]. Alternatively, it is possible to photoionise Barium with a combination of 791 nm and 337 nm light [63, 64], however, we prefer the easily accessible 405 nm laser for our set up. To trap exclusively $^{137}\text{Ba}^+$, we must be selective in our photoionisation (PI) process. Without this selectivity, the natural abundance of ^{138}Ba in the atomic source would flood the trap. We achieve this selectivity in two ways. First, we detune the 553 nm first-stage PI beam by 540 MHz from the ^{138}Ba resonance, targeting the small peak shown in Fig. 3.2, which corresponds to the ^{137}Ba resonance frequency. The second step is to verify the ion species via a fluorescence measurement before merging it into the main ion chain. We accomplish this by using the trap’s segmented electrodes to shuttle each ion to the trap centre for inspection (and moving the existing chain to a separate zone) as it enters the trap. Only if the ion reaches a fluorescence threshold for the $S_{1/2} \rightarrow P_{1/2}$ transition is it accepted and merged with the current chain. Details of this sequence can be found in Ref. [47].

3.2 Fluorescence and Doppler cooling

Once trapped, we utilise the $S_{1/2} \leftrightarrow P_{1/2}$ cycling transition to cool, prepare and measure the state of the ion. Due to the hyperfine structure in $^{137}\text{Ba}^+$, we must apply multiple tones of both 493 nm and 650 nm light to the ion. In each case this is achieved by applying the relevant tones to an EOM. In the case of the 493 nm laser, an 8.037 GHz and a 6.56 GHz tone is applied, where the sum terms then construct all four desired sidebands. These four tones are depicted in Fig. 3.1(c). Some of the power is also dissipated into off-resonant tones, such as second-order sidebands, which we can ignore. For the 650 nm laser, we apply just one tone at 430 MHz which allows the carrier, red, and blue sidebands to target all four $|D_{3/2}, F = 0, 1, 2, 3\rangle \rightarrow |P_{1/2}, F = 2\rangle$ transitions given sufficient power broadening from ~ 1000 saturation intensities of power.

We use two modes of operation when Doppler cooling: “idle” and “cold”. When “idle” we apply $\sim 100I_{\text{sat}}$ of 493 nm light far detuned to -100 MHz from resonance. This allows hot ions to be cooled either when first trapped, or after some event which can cause heating such as a collision or post-measurement.

When beginning an experimental sequence, for example a randomized benchmarking sequence, we cool as close to the Doppler limit as possible before applying any sub-Doppler cooling. This uses the “cold” configuration which uses $\sim 0.5I_{\text{sat}}$ of 493 nm light detuned by $\Gamma/2$ where Γ is the FWHM of the Doppler profile shown in Fig. 3.3. We predict a Doppler temperature $\bar{n} \approx 15$ in the lower frequency radial mode distribution using $\frac{\hbar\Gamma}{2} = \hbar\omega(\bar{n} + \frac{1}{2}) \cos(\theta) \cos(\phi)$, where θ, ϕ are the polar and azimuthal angles of projection onto the mode (both are 45°). This prediction is consistent with the initial temperature measured when characterising our EIT cooling later in this chapter (Sec. 3.3), where we measure a Doppler temperature of 12(5) quanta.




Figure 3.3: Resonances with the 493 nm laser for (a) “idle” Doppler cooling where the laser is placed far detuned and significantly power broadened to capture and cool hot ions. (b) “cold” Doppler cooling where the laser uses minimal intensity $\sim 0.5I_{\text{sat}}$ to minimise the linewidth of the resonance for optimal cooling. The detuning of the laser is set to match the $\Gamma/2$ from the peak centre where Γ is the FWHM. This is measured to be 2.3 MHz and predicts temperatures of $\bar{n}_{\text{axial}} = 34$, $\bar{n}_{\text{radial lower}} = 15$, $\bar{n}_{\text{radial higher}} = 13$ for mode frequencies of 0.5 MHz, 1.6 MHz and 1.8 MHz respectively.

3.3 EIT Cooling¹

After cooling to the Doppler limit, we further cool the ions using the electromagnetically-induced-transparency (EIT) scheme presented in Ref. [66]. This procedure uses four lambda-schemes setup between two 493 nm beams detuned from resonance as illustrated in Fig. 3.4. In this configuration, the $|S_{1/2}, F = 1\rangle \rightarrow |P_{1/2}, F = 1\rangle$, 6.56 GHz, sideband is switched off, causing the main Doppler beam shown in Fig. 3.1 to act as a pump into the $|S_{1/2}, F = 2, m_F = 0\rangle$ state (this is also how state prep is

¹Note this was implemented by William Cutler for his first year project.




Figure 3.4: EIT cooling level diagram. The π -polarised 493 nm light acts as the pump beam with detuning Δ_{Pump} . This forms four three-level systems with the σ^{\pm} -polarised probe beam with detuning Δ_{Probe} . As the detuning of the probe beam is adjusted, the resonances in Fig. 3.5 are observed. When the detuning is placed red detuned of a resonance by 1.5 MHz (the radial mode frequency), then rapid cooling occurs as seen in Fig. 3.6.

performed as explained in the next section). A much less powerful beam labelled “EIT” in Fig. 3.1 acts as the probe beam which creates the cooling cycle. The two beams enter the trap from opposite sides in the radial direction to give a strong coupling to the modes we use for gates via the difference in \vec{k}_{Pump} and \vec{k}_{Probe} . We also add a slight skew to the EIT beam to add axial projection, cooling these modes as well². We set the pump detuning to 80 MHz above resonance with $S_{1/2} \rightarrow P_{1/2}$ transition to reduce the background scattering which the pump would otherwise drive. In Fig. 3.5 we scan the detuning of the probe beam and observe a set of

²Ideally the EIT beam would enter from the exact opposite direction to the Doppler beam to maximise $\Delta\vec{k}$. This is where the photoionisation beams enter from and so would alter our well established procedure to reset the system in the event of Doppler beams misaligning. Some historical choices must unfortunately remain.




Figure 3.5: Measurement of the ions’ fluorescence as the probe detuning is scanned. Each broad peak contains a pair of resonances to $|S_{1/2}, F = 2, m_F = 0\rangle$ as illustrated in Fig. 3.4. At 78 MHz, the resonances are associated with $|S_{1/2}, F = 2, m_F = -1\rangle$ and $|S_{1/2}, F = 1, m_F = 1\rangle$. At 82 MHz, the other two resonances occur with $|S_{1/2}, F = 2, m_F = +1\rangle$ and $|S_{1/2}, F = 1, m_F = -1\rangle$. Smaller peaks on the slope are due to other resonances within the hyperfine structure which we ignore for the purposes of cooling.

resonances associated with the four transitions in Fig. 3.4.

We choose to drive the EIT with the probe beam red detuned by 1.5 MHz from the resonance at 77 MHz to prioritise the cooling of the radial modes with frequencies between 1–2 MHz. We then measure the cooling rate from EIT on the lower frequency radial mode of a single ion using the sideband ratio method outlined in Sec. 7.3.2. The result is shown in Fig. 3.6. We measure a cooling rate of $4.0(6) \times 10^3$ q/s fitting an exponential model and final temperature $\bar{n} = 0.66(6)$ quanta. To perform final stages of cooling on a specific mode (e.g. for two-qubit gates) we then apply resolved-sideband cooling [67]. This brings the lower radial mode of a single ion to 0.40(5) quanta, and the rocking mode of two ions down to 0.004(4) quanta, both measured with sideband thermometry [68] (See Sec. 7.3.2 for details.).

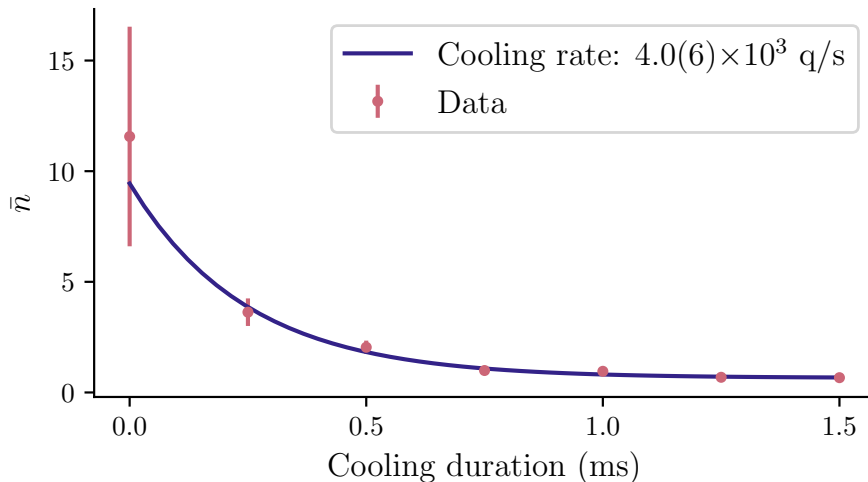


Figure 3.6: Temperature measurements vs EIT cooling time. Starting from an initial temperature of 12(5) quanta and cooling to 0.66(6) quanta in just over 1 ms. Fitting an exponential model to this we measure a cooling rate of $4.0(6) \times 10^3$ q/s.

3.4 State preparation and measurement

Before any coherent operations can take place, we must initialise our qubit into a fiducial state. In trapped-ion systems, this is typically achieved through optical pumping. In our case, we pump into the $|S_{1/2}, F = 2, m_F = 0\rangle$ state using π -polarised light, with only the sidebands required to drive the $|S_{1/2}, F = 2\rangle \rightarrow |P_{1/2}, F = 2\rangle$ and $|S_{1/2}, F = 1\rangle \rightarrow |P_{1/2}, F = 2\rangle$ transitions. This configuration takes advantage of the forbidden $|S_{1/2}, F = 2, m_F = 0\rangle \rightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ transition. The transitions utilised for state-preparation are depicted in Fig. 3.7.

We characterise the optical pumping error by performing four 1762 nm π -pulses shelving population to the $D_{5/2}$ level and checking for fluorescence. The use of several 1762 nm π -pulses exponentially reduces the error from imperfect mapping leaving only the optical pumping error. If sufficient fluorescence is detected, this indicates that the qubit was not properly prepared. The result of this is shown in Fig. 3.8 where we measure a steady state error of 2.0(3)%. We perform optical pumping for 25 μ s just beyond the time required to reach steady state. This can be seen clearly in Fig. 3.8.

After a desired quantum circuit has been run, we measure the state of the qubit by




Figure 3.7: Diagram of the lasers used for State preparation. The 493 nm beam uses only the $|S_{1/2}, F = 2\rangle \rightarrow |P_{1/2}, F = 2\rangle$ and $|S_{1/2}, F = 1\rangle \rightarrow |P_{1/2}, F = 2\rangle$ transitions with pure π -polarised light. This leaves a single forbidden transition from $|S_{1/2}, F = 2, m_F = 0\rangle \rightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ causing population to be pumped into $|S_{1/2}, F = 2, m_F = 0\rangle$.




Figure 3.8: Measured bright fraction vs optical pumping duration. After optical pumping several quadrupole pulses drive population from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $D_{5/2}$. If the state preparation is perfect then no bright population is measured. An exponential decay is fit to the data showing a time constant of $3.77 \mu\text{s}$ and a steady state error (bright population) of $2.0(3)\%$. In experiments throughout this thesis we perform optical pumping for well beyond the $3.77 \mu\text{s}$ time constant for $25 \mu\text{s}$ where steady state is approximately reached.

converting the qubit to an optical configuration and checking for fluorescence. Bright indicates projection into the $S_{1/2}$ and dark the $D_{5/2}$. When working with a metastable qubit we convert between qubit types using a 1762 nm pulse. To reduce the error from imperfect mapping, we use a discretised variant of the procedure used in Ref. [69] where alternating cycles of 1762 nm and 493 nm light pump population into the $D_{3/2}$ level, ready for fluorescence detection using the 493 nm and 650 nm beams.

In Chapter 6 we push the fidelity of our state preparation by several orders of magnitude by utilising a heralding scheme to remove unsuccessful preparation and measurement attempts.

All of his co-workers were gone, what could it mean?

— The Narrator, The Stanley Parable

4

Single qubit operations

Contents

4.1	Atom-light interactions	31
4.1.1	Dipole interactions	32
4.1.2	Coupling to the motion	36
4.1.3	Quadrupole interactions	38
4.1.4	Raman interactions	39
4.2	Implementing single qubit operations	41
4.2.1	Optical qubit	41
4.2.2	Metastable qubit	44
4.3	Calibrating single qubit operations	47
4.3.1	Ramsey spectroscopy	48
4.3.2	Robust phase estimation	49
4.4	Randomised benchmarking	56
4.4.1	Metastable qubit RBM	57
4.4.2	Optical qubit RBM	62
4.5	Converting between qubit types	69
4.5.1	Clock qubit	70
4.5.2	Hybrid qubit	71
4.5.3	Maintaining coherence	71
4.6	Summary	74

In Chapter 2, we performed some theoretical analysis of the $^{137}\text{Ba}^+$ ion to inform our choice of metastable and optical qubits. In this chapter, we implement single-qubit operations in the qubits we have chosen and compare their performance. Such

single-qubit operations lay the foundation for our quantum computer, which, when combined with a two-qubit gate, provide access to arbitrary unitaries. In this thesis, we place focus on laser-based gates in optical and metastable qubits. Implementing our gates with lasers allows for fast, albeit limited, scaling, which is the topic of Chapter 5. We begin by outlining the theory of single-qubit operations and follow this with the details of how we choose our laser geometry for both quadrupole and two-photon Raman transitions. With this theory in place, the single-qubit operations are calibrated using robust phase estimation (RPE) [70] and characterised with randomized benchmarking (RBM) [71]. Finally, the chapter closes by considering how we can map between qubit types to integrate the OTDF gate in Chapter 7 into our *omg* scheme. This mapping is critical to using the $D_{5/2}$ level to protect information from mid-circuit measurements.

Unless otherwise stated, the metastable qubit is $|D_{5/2}, F = 3, m_F = +2\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = +2\rangle$ and the optical qubit is $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$, both implemented at 3.33 G.

4.1 Atom-light interactions

The Hamiltonian of a free two-level atom can be expressed as

$$\tilde{H}_A = \frac{\hbar}{2}\omega_0\sigma_z, \quad (4.1)$$

where the energies are expressed as an offset from the mean. The interaction between the ion and a light field can be expressed through the Hamiltonian

$$H_{AE} = - \int d^3\vec{r} \vec{P}^\perp(\vec{r}) \cdot \vec{E}^\perp(\vec{r}), \quad (4.2)$$

where \vec{P}^\perp is the transverse component of the polarisation field of the atom, and \vec{E}^\perp is an external electric field¹ [72].

¹We have dropped the self-interacting term in Eq. 4.2 for simplicity since it will be ignored. For a full treatment see Ref. [72], chapter 9.5.

For a hydrogen like atom, with a central singly charged nucleus located at the origin, and an electron positioned at \vec{r}_e the polarisation field is given by,

$$P(\vec{r}) = -e\vec{r}_e \int_0^1 ds \delta^3(\vec{r} - s\vec{r}_e). \quad (4.3)$$

Taking the expansion of the effect of the delta function on $\vec{E}^\perp(\vec{r})$,

$$H_{AE} = \int d^3\vec{r} \int_0^1 ds e\vec{r}_e (\delta^3(\vec{r}) - s[\vec{r}_e \cdot \delta(\vec{r} - s\vec{r}_e)\nabla] + \dots)\vec{E}^\perp(\vec{r}), \quad (4.4)$$

and evaluating the integrals we find

$$H_{AE} = er_{e,\alpha} E_\alpha^\perp(0) - \frac{1}{2}e(r_{e,\alpha}r_{e,\beta})\partial_\alpha E_\beta^\perp(0) + \dots \quad (4.5)$$

Finally, we write $-er_{e,\alpha} = d_\alpha$, $-\frac{1}{2}e(r_{e,\alpha}r_{e,\beta}) = Q_{\alpha,\beta}$ and transform to the centre of mass of the atom \vec{r}_n ² leaving,

$$H_{AE} = -d_\alpha E_\alpha^\perp(\vec{r}_n) + Q_{\alpha,\beta}\partial_\alpha E_\beta^\perp(\vec{r}_n) + \dots \quad (4.6)$$

The first term is the electric dipole interaction (ED), and the second is the electric quadrupole interaction (EQ). A similar expansion can be applied to the moments of a magnetic field, but in this thesis, we are only concerned with transitions driven by electric fields.

4.1.1 Dipole interactions

The dipole interaction is utilised for our dissipative operations, such as cooling and state preparation. In this regime, the lifetime of the excited state is small relative to a typical Rabi frequency, and we can formulate the interaction through the optical Bloch equations [72, 73]. For the purposes of this chapter, we are interested in the dipole interaction as a simple basis from which to learn about coherent quantum interactions. The physics of dipole interactions can then be mapped onto the following sections on quadrupole interactions and two-photon Raman transitions, and the light shift is used in Chapter 7 to construct two-qubit gates.

²This transformation is rather involved but can be found in Ref. [72] chapter 9.6.

The electric dipole interaction is

$$H_{ED} = -\vec{d} \cdot \vec{E}(\vec{r}, t). \quad (4.7)$$

For an oscillating electric field with polarisation components $\vec{\epsilon}$ this can be written as

$$H_{ED} = -\vec{d} \cdot \vec{\epsilon} E_0 \cos(\vec{k} \cdot \vec{r} - \omega t + \phi). \quad (4.8)$$

Following the treatment in Ref. [74], we then take the second quantisation of the Hamiltonian to obtain³,

$$H_{ED} = \sum_{i,j} \hbar \Omega_{i,j} \cos(\vec{k} \cdot \vec{r} - \omega t + \phi) |i\rangle \langle j|, \quad (4.9)$$

where the Rabi frequency $\Omega_{i,j}$ is defined as,

$$\Omega_{i,j} = \frac{-E_0 \langle i | \vec{d} \cdot \vec{\epsilon} | j \rangle}{\hbar}. \quad (4.10)$$

The dipole operator \vec{d} has odd parity, and thus the diagonal elements of $\Omega_{i,j}$ are zero. As such, this Hamiltonian describes the interaction between different levels in the presence of the light field. Depending on the frequency of the light, ω , relative to the characteristic energy between a pair of states, $\omega_{i,j}$, this may describe driving a transition (Rabi flopping), or a shift in energies i.e. entering a “dressed” picture.

Later, we will discuss the coupling between the position operator \vec{r} and the atoms energy, but for now, we can consider the atom stationary with $\vec{r} = \vec{0}$. To express H_{ED} in a form which is useful for calculating light shifts and Rabi oscillations, we start by expanding H_{ED} in terms of two counter-rotating terms,

$$H_{ED} = \sum_{i,j} \frac{\hbar}{2} \Omega_{i,j} (e^{-i(\omega t + \phi)} + e^{i(\omega t + \phi)}) |i\rangle \langle j|. \quad (4.11)$$

We then take the full two level Hamiltonian, $\tilde{H} = \tilde{H}_A + \tilde{H}_{ED}$, and transform into the rotating frame with angular frequency ω .

$$\tilde{H} = \frac{\hbar}{2} \delta \sigma_z + \frac{\hbar}{2} \Omega (e^{-i(\omega t + \phi)} + e^{i(\omega t + \phi)}) \cdot (e^{-i\omega t} |\uparrow\rangle \langle \downarrow| + e^{i\omega t} |\downarrow\rangle \langle \uparrow|) \quad (4.12)$$

³To work the expression into this form, left and right multiply by $\mathbf{1} = \sum_i |i\rangle \langle i|$

Finally, we apply the rotating-wave approximation (RWA) by neglecting fast-rotating terms ($\omega \gg \Omega, \delta$), we obtain the Hamiltonian:

$$\tilde{H} \approx \frac{\hbar}{2} \delta \sigma_z + \frac{\hbar}{2} \Omega \left(e^{-i\phi} |\downarrow\rangle \langle \uparrow| + e^{i\phi} |\uparrow\rangle \langle \downarrow| \right) \quad (4.13)$$

where $\delta = \omega - \omega_0$. This Hamiltonian is time-independent, and the result propagator is thus,

$$\begin{aligned} U(t) &= e^{-\frac{i}{\hbar} \tilde{H} t} \\ &= \mathbb{1} \cos\left(\frac{\Omega_{\text{eff}} t}{2}\right) + \frac{i}{\Omega_{\text{eff}}} \left(\delta \sigma_z - \Omega e^{i\phi} \sigma_+ - \Omega e^{-i\phi} \sigma_- \right) \sin\left(\frac{\Omega_{\text{eff}} t}{2}\right), \end{aligned} \quad (4.14)$$

where $\Omega_{\text{eff}} = \sqrt{\Omega^2 + \delta^2}$. This Hamiltonian has the effect of driving coherent transitions between the states $|\downarrow\rangle$ and $|\uparrow\rangle$. To gain an intuition of how a qubit behaves under this unitary, and the utility it can provide, we consider the case when $\delta=0$ and $\delta \gg \Omega$.

On resonance

When $\delta=0$, the propagator simplifies to

$$U(t) = \begin{pmatrix} \cos\left(\frac{\Omega t}{2}\right) & i e^{i\phi} \sin\left(\frac{\Omega t}{2}\right) \\ i e^{-i\phi} \sin\left(\frac{\Omega t}{2}\right) & \cos\left(\frac{\Omega t}{2}\right) \end{pmatrix}. \quad (4.15)$$

The unitary evolution, $U(t)$, describes a rotation by an angle $\theta = \Omega t$ about the axis in the equatorial plane defined by the phase ϕ of the driving field. This axis can be written as

$$\hat{n} = \cos(\phi) \vec{x} + \sin(\phi) \vec{y},$$

so that the propagator corresponds to the rotation operator

$$U(t) = e^{-i \frac{\Omega}{2} \vec{n} \cdot \vec{\sigma}},$$

where $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$ is the vector of Pauli matrices. Several familiar operations arise from particular choices of the driving time t and phase ϕ :

- For $\phi = 0$ and $\Omega t = \pi$ (which we refer to as a π -pulse), the unitary becomes an X gate (bit-flip):

$$U = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}.$$

- For $\phi = \pi/2$ and $\Omega t = \pi$ (π_Y -pulse), the unitary becomes a Y gate (bit- and phase-flip):

$$U = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

- For $\phi = 0$ and $\Omega t = \pi/2$ ($\pi/2$ -pulse), the gate corresponds to a $\pi/2$ rotation around the x -axis:

$$U = \begin{pmatrix} \cos\left(\frac{\pi}{4}\right) & i \sin\left(\frac{\pi}{4}\right) \\ i \sin\left(\frac{\pi}{4}\right) & \cos\left(\frac{\pi}{4}\right) \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}.$$

By combining these operations, with arbitrary phase control, we obtain arbitrary single-qubit operations. With the addition of a two-qubit gate, we have a universal gate set.

In the case of dipole transitions within Ba^+ , the lifetimes of the excited $P_{1/2}$, and $P_{3/2}$, states are too short relative to Ω and the typical pulse durations accessible to maintain any qubit coherence. However, the Hamiltonian derived here can be mapped to quadrupole transitions between $S_{1/2}$, and $D_{5/2}$, which do have a sufficiently long lifetime to maintain coherence.

Far off resonance

In the case $\delta \gg \Omega$, then Eq. 4.14 approximates to

$$\begin{aligned} U(t) &\approx \mathbb{1} \cos\left(\frac{\Omega_{\text{eff}} t}{2}\right) + i \sin\left(\frac{\Omega_{\text{eff}} t}{2}\right) \sigma_z \\ &= e^{\frac{i\Omega_{\text{eff}} t}{2} \sigma_z}. \end{aligned} \quad (4.16)$$

which drives a Z rotation with angular frequency $\Omega_{\text{eff}} = \sqrt{\Omega^2 + \delta^2} \approx \delta + \Omega^2/2\delta$. In this regime, the terms driving Rabi oscillations have been dropped, leaving only a shift in phase. The δ term from the Taylor expansion appears due to the choice of the laser's frame. Transforming into the atom's rotating frame leaves only

$$\tilde{U}(t) = e^{\frac{i\Omega^2}{4\delta} \sigma_z}, \quad (4.17)$$

which corresponds to a shift in the energy of the Z eigenstates by $\hbar\Omega^2/4\delta$. This is the light shift, or A.C. Stark shift.

This operation is sometimes undesirable in trapped-ion quantum processors [69]. After a change of basis, an uncontrolled σ_z rotation can propagate through a system in the same manner as bit-flip errors, which must then be corrected [34]. Additionally, due to the complex hyperfine structure of $^{137}\text{Ba}^+$, off-resonant coupling to other dipole transitions can produce uncontrolled energy shifts when driving a qubit, leading to the qubit having a different resonant frequency when idle and active – an effect that must be accounted for.

However, the light shift also has utility, as it forms the basis of the two-qubit gate in our system (Section 7.2.2) and was used to characterise our addressing system in Chapter 5. It can also be used to enable ion selectivity on transitions driven by global fields, by shifting specific ions in or out of resonance with the driving field [75].

4.1.2 Coupling to the motion

In the previous section, we have neglected the coupling to the position operator \vec{r} . Reintroducing this, Eq. 4.13 is modified to

$$\tilde{H} \approx \frac{\hbar}{2} \delta\sigma_z + \frac{\hbar}{2} \Omega \left(e^{-i(\vec{k}\cdot\vec{r}+\phi)} |\downarrow\rangle \langle\uparrow| + e^{i(\vec{k}\cdot\vec{r}+\phi)} |\uparrow\rangle \langle\downarrow| \right). \quad (4.18)$$

Our ions are in a quantised harmonic well (the trapping potential), with $3N$ motional modes, where N is the number of ions. For a single mode, we can express $\vec{k}\cdot\vec{r}$ as

$$\begin{aligned} \vec{k}\cdot\vec{r} &= z_0 \vec{k}\cdot\vec{n}(a^\dagger + a) \\ &= \eta(a^\dagger + a), \end{aligned} \quad (4.19)$$

where $z_0 = \sqrt{\langle 0|z^2|0\rangle}$ describes the spatial extent of the ground state of motion [76]. The vector \vec{n} describes the direction of the motional mode of the ion, and these parameters are combined to give the Lamb-Dicke parameter η . This can be easily extended to multiple motional modes and multiple ions by using $\vec{k}\cdot\vec{r}_n$ where n is the ion number, and summing over the various modes under the definition $\vec{r}_n = z_0 \sum_i \vec{n}_{i,n}(a_i^\dagger + a_i)$.

We also need to modify the two-level atom Hamiltonian to include the harmonic potential. Equation 4.1 becomes

$$\tilde{H}_A = \frac{\hbar}{2}\omega_0\sigma_z + \hbar\omega_m\left(a^\dagger a + \frac{1}{2}\right), \quad (4.20)$$

where ω_m is the frequency of the motional mode. This addition is relevant as we then substitute Eq. 4.19 into Eq. 4.18 in the interaction picture of the motion to find

$$\tilde{H} \approx \frac{\hbar}{2}\delta\sigma_z + \frac{\hbar}{2}\Omega \sum_{n,n'} \left(e^{-i(\omega_m(n'-n)t+\phi)} |n'\rangle \langle n'| e^{-i\eta(a^\dagger+a)} |n\rangle \langle n| \otimes |\downarrow\rangle \langle \uparrow| + \text{H.c.} \right). \quad (4.21)$$

This then provides another resonance condition: if $\delta = \omega_m(n' - n)$, the sideband between motional states $|n\rangle$ and $|n'\rangle$ is driven. Coupling to motional sidebands allows us to perform resolved sideband cooling and thermometry [68]. Additionally, the Rabi frequency Ω is modified according to

$$\begin{aligned} \Omega_{n,n'} &= \Omega \langle n'| e^{-i\eta(a^\dagger+a)} |n\rangle \\ &= \Omega e^{-\eta^2/2} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta^{|n'-n|} \mathcal{L}_{n_{<}}^{|n'-n|}(\eta^2), \end{aligned} \quad (4.22)$$

where $\mathcal{L}_i^j(\eta^2)$ is the generalised Laguerre polynomial, and $n_{<}(n_{>})$ is the smaller (larger) of n and n' [76, 77].

The modification of the Rabi frequency due to ion temperature can have a significant impact on the operations we perform, particularly for the optical qubit, which has an unavoidable sensitivity to the motion. In Fig. 4.1, we show the adjusted Rabi frequency for both the optical qubit driven by 1762 nm light and a counter-propagating Raman configuration with 532 nm light for our trap geometry (Fig. 3.1), where there is a 45° angle to the radial modes. The simulation is with a single radial mode (the lower-frequency mode) which has frequency $\omega_m = 2\pi \times 1.6$ MHz. For all single qubit operations in this section with metastable qubits, we use a co-propagating geometry which has $\eta=0$, removing the motional coupling completely.

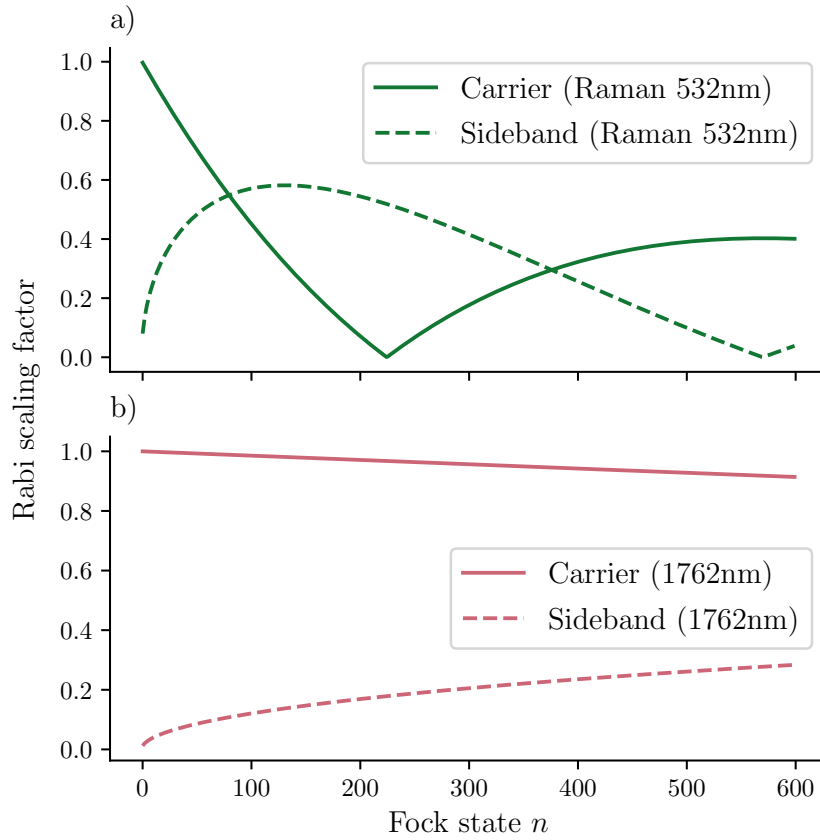


Figure 4.1: Calculation of the scaling factor $\langle n' | e^{-i\eta(a^\dagger+a)} | n \rangle$ for the carrier and first sideband of: (a) a two photon Raman process with counter-propagating 532 nm beams (see Sec. 4.1.4 and Sec. 4.2.2). (b) a quadrupole transition using 1762 nm light (see Sec. 4.1.3 and Sec. 4.2.1). In both cases there is a 45° projection onto a radial mode with frequency $\omega = 2\pi \times 1.6$ MHz.

4.1.3 Quadrupole interactions

Returning to the multipole expansion of the atom-light interaction Hamiltonian (Eq. 4.6), for some pairs of states, the electric dipole matrix element $\langle i | \vec{d} \cdot \vec{\epsilon} | j \rangle / \hbar = 0$. However, there is still the quadrupole coupling to the gradient of the electric field. In our system, the quadrupole transitions of interest are between $S_{1/2}$ and $D_{5/2}$.

For an oscillating electric field with polarisation components $\vec{\epsilon}$, the electric

quadrupole (EQ) Hamiltonian can be written as,

$$\begin{aligned}
H_{EQ} &= -\frac{1}{2} e\vec{r}_e \cdot \vec{\partial}_{\vec{r}_n} (\vec{r}_e \cdot \vec{E}(\vec{r}_n)) \\
&= -\frac{eE_0}{2} (\vec{r}_e \cdot \vec{\epsilon}) (\vec{r}_e \cdot \partial_{\vec{r}_n} [\cos(\vec{k} \cdot \vec{r}_n - \omega t + \phi)]) \\
&= \frac{eE_0}{2} (\vec{r}_e \cdot \vec{\epsilon}) (\vec{r}_e \cdot \vec{k}) \sin(\vec{k} \cdot \vec{r}_n - \omega t + \phi).
\end{aligned} \tag{4.23}$$

Again left- and right-multiplying by $\mathbb{1} = \sum_i |i\rangle\langle i|$ we take the second quantisation of the Hamiltonian to find,

$$\begin{aligned}
H_{EQ} &= \sum_{i,j} \hbar \Omega_{i,j} \sin(\vec{k} \cdot \vec{r}_n - \omega t + \phi) |i\rangle\langle j| \\
\Omega_{i,j} &= \frac{eE_0}{2\hbar} \langle i | (\vec{r}_e \cdot \vec{\epsilon}) (\vec{r}_e \cdot \vec{k}) | j \rangle,
\end{aligned} \tag{4.24}$$

which has the same form as the dipole Hamiltonian H_{ED} (Eq. 4.9) up to a phase shift, switching $\cos(\cdot)$ to $\sin(\cdot)$ and a modified matrix element/Rabi frequency. As a result, all the same physics regarding Rabi flopping and the light shift applies to the quadrupole coupling.

4.1.4 Raman interactions

In many quantum computing systems, it is desirable to drive electric-dipole forbidden transitions – for example, our metastable clock qubit. We choose to do this using a Raman interaction, which is treated as follows. Consider a three-level atom as shown in Fig. 4.2. There are two levels associated with the qubit, $|\downarrow\rangle$ and $|\uparrow\rangle$, which both couple to $|e\rangle$ via an electric-dipole transition, but the direct transition is electric-dipole forbidden. We define the resonance frequency of the transitions to $|e\rangle$ as $\omega_{0,e}$ and $\omega_{1,e}$, and apply two lasers with frequency ω_0 and ω_1 . The Hamiltonian describing this system is

$$\begin{aligned}
H &= H_0 \\
&+ \hbar \left(\Omega_{\downarrow,e;0} \cos(\vec{k}_0 \cdot \vec{r} - \omega_0 t + \phi_0) + \Omega_{\downarrow,e;1} \cos(\vec{k}_1 \cdot \vec{r} - \omega_1 t + \phi_1) \right) (|\downarrow\rangle\langle e| + |e\rangle\langle\downarrow|) \\
&+ \hbar \left(\Omega_{\uparrow,e;0} \cos(\vec{k}_0 \cdot \vec{r} - \omega_0 t + \phi_0) + \Omega_{\uparrow,e;1} \cos(\vec{k}_1 \cdot \vec{r} - \omega_1 t + \phi_1) \right) (|\uparrow\rangle\langle e| + |e\rangle\langle\uparrow|),
\end{aligned} \tag{4.25}$$




Figure 4.2: An example three level system driven by a Raman transition. The lasers connect the states $|\downarrow\rangle$ and $|\uparrow\rangle$ via a dipole allowed transition to the excited auxiliary level $|e\rangle$. The dipole-forbidden transition $|\downarrow\rangle \leftrightarrow |\uparrow\rangle$ has a resonance frequency of ω . The two lasers are detuned such that $\omega_1 - \omega_0 = \omega + \delta$. This detuning is several orders of magnitude smaller than Δ .

where $\Omega_{i,e;j}$ represents the Rabi frequency for the transition $|i\rangle \leftrightarrow |e\rangle$ due to laser j . In this instance, it is most convenient to transform into the frame of the atom, using the unitary $U(t) = \sum_{j=\{\downarrow,\uparrow,e\}} e^{i\hbar\omega_j t} |j\rangle \langle j|$. Additionally exponentiating the $\cos(\cdot)$ terms and making the rotating wave approximation we find,

$$\begin{aligned}
 H \approx & \frac{\hbar}{2} \left(\Omega_{\downarrow,e;0} e^{i(\vec{k}_0 \cdot \vec{r} + \Delta t + \phi_0)} + \Omega_{\downarrow,e;1} e^{i(\vec{k}_1 \cdot \vec{r} + (\Delta + \omega + \delta)t + \phi_1)} \right) |\downarrow\rangle \langle e| \\
 & + \frac{\hbar}{2} \left(\Omega_{\uparrow,e;0} e^{i(\vec{k}_0 \cdot \vec{r} + (\Delta - \omega)t + \phi_0)} + \Omega_{\uparrow,e;1} e^{i(\vec{k}_1 \cdot \vec{r} + (\Delta + \delta)t + \phi_1)} \right) |\uparrow\rangle \langle e| + \text{H.c.}
 \end{aligned} \tag{4.26}$$

We then make the James-Jerk approximation [78] to find the effective time averaged Hamiltonian:

$$\begin{aligned}
H_{\text{eff}} = & \frac{\hbar}{4\Delta} \left(\Omega_{\downarrow,e;0}\Omega_{\uparrow,e;1}e^{i(\Delta\vec{k}\cdot\vec{r}+\delta\phi)}e^{-i\delta t} \right. \\
& + \Omega_{\downarrow,e;1}\Omega_{\uparrow,e;0}e^{-i(\Delta\vec{k}\cdot\vec{r}+\delta\phi)}e^{i(2\omega+\delta)t} \\
& + \Omega_{\downarrow,e;0}\Omega_{\downarrow,e;0}e^{-i\omega t} \\
& \left. + \Omega_{\uparrow,e;1}\Omega_{\uparrow,e;1}e^{i\omega t} \right) |\downarrow\rangle\langle\uparrow| + \text{H.c} + \text{L.S},
\end{aligned} \tag{4.27}$$

where we have made the approximation that $\Delta \gg \omega, \delta$, neglected light shift (L.S) terms for brevity, and $\Delta\vec{k} = \vec{k}_0 - \vec{k}_1$, $\delta\phi = \phi_0 - \phi_1$ ⁴. When close to resonance, only the first term is of interest. It has the form of Rabi flopping between $|\uparrow\rangle$ and $|\downarrow\rangle$, as in Eq. 4.13, but with a modified Rabi frequency

$$\Omega = \frac{\Omega_{\downarrow,e;0}\Omega_{\uparrow,e;1}}{2\Delta}. \tag{4.28}$$

The remaining three represent off-resonant cross-terms from the two beams, as well as each individual beam with itself. This unintuitive effect of driving Raman transitions with a single beam has been observed in our system, see Ref. [47], Fig. 8.3.

4.2 Implementing single qubit operations

The system we present in this thesis is largely an explorative one. We are implementing many different qubits, qubit types and operations, often choosing which is most appropriate for a given demonstration. In this light, we present the current state of the system as one which aims to maximise flexibility to explore a variety of different research avenues, often simultaneously.

4.2.1 Optical qubit

The optical qubit is driven directly with 1762 nm light, utilising the quadrupole coupling described in Sec. 4.1.3.

⁴The “ Δ ” in $\Delta\vec{k}$ and “ δ ” in $\delta\phi$ have no relation to the detunings Δ and δ .

We wish to access states in all hyperfine manifolds of the $S_{1/2}$ and $D_{5/2}$ of $^{137}\text{Ba}^+$. In the $S_{1/2}$ at low field this requires a frequency range of over 8.037 GHz and ~ 150 MHz in the $D_{5/2}$. The primary reason for covering the range over the ground hyperfine splitting is to implement virtual qubits and qudits [56], which will be the topic of my colleague Andres Vazquez-Brennan's thesis.

To cover such a large frequency range, we use an electro-optic modulator⁵ (EOM) for phase and frequency control, combined with an acousto-optic modulator⁶ (AOM) in a single-pass configuration for optical switching. The 1762 nm laser is locked to a high-finesse cavity such that the carrier lies between the $|S_{1/2}, F = 1\rangle$ and $|S_{1/2}, F = 2\rangle$ hyperfine manifolds, which are 4.5 GHz and -3.5 GHz detuned, respectively. We then use the red sideband of the EOM to drive transitions from the $|S_{1/2}, F = 2\rangle$ manifold and the blue sideband to drive from $|S_{1/2}, F = 1\rangle$. This results in negligible cross talk or off-resonant excitation from the other sideband, due to the asymmetry in detuning to each state.

To generate tones between 3.5 GHz and 4.5 GHz with coherent phase control, we take a very direct approach: multiplying, filtering, and amplifying the signal from an Urukul in the sinara-hardware ecosystem by $12\times$. This can have the downside of amplifying errors in phase due. The error could originate from floating point precision in code, though these are negligible, or more likely due to timing errors from the clock network driving the control system. However, with the system clocked by a Rb source⁷ with a one-second Allan variance $< 2\times 10^{-11}$ this is not a concern, even with the multiplication.

Continuing on the quest for flexibility, we also wish to drive $\Delta m = 0, \pm 1, \pm 2$ transitions. The coupling to such transitions is determined by the quadrupole matrix element (Eq. 4.24), which is dependent upon both the polarisation of the driving field and the geometry of this field with the atom's quantisation axis, defined

⁵iXblue MPZ-LN-10

⁶ISOMET M1205-P80L-0.6

⁷Stanford Research Systems, Model FS725 Rubidium Frequency Standard

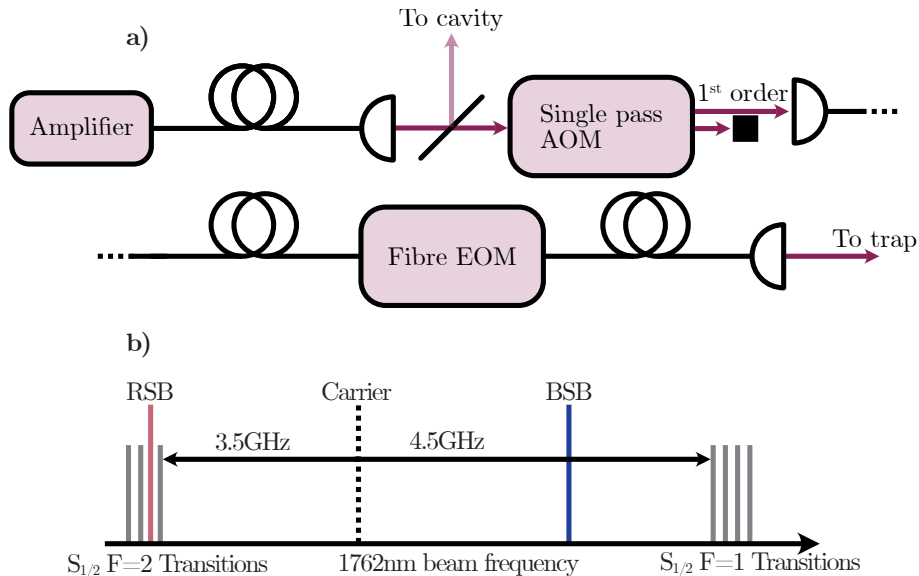


Figure 4.3: a) Diagram of the 1762 nm optical setup between laser amplifier and trap used for switching and frequency control. A 90:10 pickoff is taken with 10% being used for locking to the high-finesse cavity. The remaining light is passed through an AOM for optical switching (ISOMET M1205-P80L-0.5,-1,-2) and the first order coupled to a fibre. Frequency control is performed using a fibre EOM (iXBlue MPZ-LN-10). b) How frequency control is achieved using the EOM. When driven at ~ 3.5 GHz the red sideband (RSB) can be used to drive transitions from $|S_{1/2}, F = 2\rangle$ to $D_{5/2}$. At ~ 4.5 GHz the blue sideband (BSB) drives transitions from $|S_{1/2}, F = 1\rangle$ to $D_{5/2}$. The remaining light in the unused sideband and carrier are sufficiently far off resonance to not coherently move population, though some small light shift will be induced.

by the magnetic field vector \vec{B} . Following the treatment in Ref. [79], the matrix element in Eq. 4.24 can be expressed as

$$\langle i | (\vec{r}_e \cdot \vec{\epsilon})(\vec{r}_e \cdot \vec{k}) | j \rangle = \sum_{\Delta m} \epsilon_m k_n c_{l,m}^{(\Delta m)} \begin{pmatrix} j_i & 2 & j_j \\ -m_i & \Delta m & m_j \end{pmatrix} \langle i || r^2 C_{\Delta m}^{(2)} || j \rangle, \quad (4.29)$$

where the bracket represents the Wigner-3j symbol and the right hand term is the reduced matrix element. It is the first term which represents the geometry and polarisation factor that we are interested in. Defining this as $g^{(\Delta m)}(\phi, \gamma, \theta) = \epsilon_m k_n c_{l,m}^{(\Delta m)}$ ⁸, with the wave vector \vec{k} at an angle ϕ to \vec{B} , polarisation angle γ , and

⁸For the definitions of $c_{l,m}^{(\Delta m)}$ see the appendix of Ref. [79].




Figure 4.4: Geometry factor from Eq. 4.30. In experiment, we use an angle $\phi = \pi/4$ and $\gamma = \pi/4$ to enable all of $\Delta m = 0, \pm 1, \pm 2$ transitions.

phase θ , evaluating we find

$$\begin{aligned}
 g^{(0)}(\phi, \gamma, \theta) &= -\frac{1}{2} \cos(\gamma) \sin(2\phi), \\
 g^{(\pm 1)}(\phi, \gamma, \theta) &= \frac{1}{\sqrt{6}} \left[i \sin(\gamma) \sin(\phi) e^{i\theta} \mp \cos(2\phi) \cos(\gamma) \right], \\
 g^{(\pm 2)}(\phi, \gamma, \theta) &= \frac{1}{\sqrt{6}} \left[\frac{1}{2} \cos(\gamma) \sin(2\phi) \mp i \sin(\gamma) \sin(\phi) e^{i\theta} \right].
 \end{aligned} \tag{4.30}$$

With linear polarisation, the magnitudes of the geometry factors are shown in Fig. 4.4. We choose to use an angle of $\pi/4$ to \vec{B} , enabling strong $\Delta m = 0$ transitions. By adjusting the polarisation, we can also enable $\Delta m = \pm 1, 2$ transitions allowing for full flexibility to work with multiple qubit types.

4.2.2 Metastable qubit

The metastable qubit is driven using pairs of 532 nm lasers in a two-photon Raman process, as previously described. A diagram illustrating this configuration, as well as how it can be applied to a ground-state qubit, is shown in Fig. 4.5. When driving motion, either for sideband cooling or two-qubit gates, there must be some difference in k -vector between the two beams. However, for single-qubit operations, it is desirable that this $\Delta \vec{k}$ is zero, such that errors due to thermal occupation and heating have no impact. This leads us to use three beams, which we label “co-propagating”, “counter-propagating”, and “metastable”. The terms co-propagating




Figure 4.5: Level diagram with the 532 nm Raman lasers. These beams can be used to drive electric-dipole forbidden transitions in both the $S_{1/2}$ by using a frequency difference between the two beams of ~ 8 GHz, or in the $D_{5/2}$ with a frequency difference of ~ 60 MHz. In this thesis we focus on the metastable qubits.

and counter-propagating are with respect to the metastable beam, which gets its name as it has a frequency offset to the other two, allowing for metastable transitions to be driven. The orientation of the beams is restrained by considerations for our two-qubit gate; details are provided in Sec. 7.5. This essentially constrains the two beam paths to be counter-propagating and orthogonal to the axial modes of the ion chain, in order to reduce excitation of these modes. Additionally, the co-propagating and counter-propagating beams, which drive the two-qubit gate, have linear, vertical polarisation.

This leaves two degrees of freedom in the polarisation angle and phase of the metastable beam. However, we must also overlap the metastable and co-propagating beams. This is best done in one of two ways: either coupling to the same fibre before delivery to the ions, or combining on a polarising beam splitter (PBS) with




Figure 4.6: Simulation of the matrix element for three different Raman transitions ($\Delta m = 0, \pm 1$) as the polarisation of the “metastable” beam is adjusted with linear vertical polarisation in the “co-propagating” beam. In implementation we prioritise the $|F = 2, m_F = +2\rangle \rightarrow |F = 3, m_F = +2\rangle$ clock transition setting the angle to $\gamma = \pi/2$. This also allows for the “metastable” and “co-propagating” beams to be combined without loss in power.

orthogonal polarisations. Coupling to the same fibre has the downside of discarding 50% of the optical power in the process. Additionally, we have the restriction of using PCF fibres in our experiment, which are polarisation-maintaining for vertical polarisation aligned to the fibre. This forces the metastable beam to have either the same, or orthogonal polarisation to the co-/counter-propagating beams. We choose to have orthogonal polarisation between the beams to maximise the Rabi frequency of the metastable clock transition, while maintaining the ability to map to other locations in the metastable level.

The matrix element of the clock and two other example transitions are plotted in Fig. 4.6 for the counter-propagating geometry described, with an angle of 45° to the B-field. On the x-axis, we adjust the polarisation angle while retaining linear polarisation, and see that the clock transition is fastest when the two beams are orthogonal in polarisation.

To generate the desired frequency offset between the Raman beam pairs, we pass the 532 nm light through a series of AOMs, as illustrated in Fig. 4.7. This generates a ~ 50 MHz frequency difference between the metastable and each of the

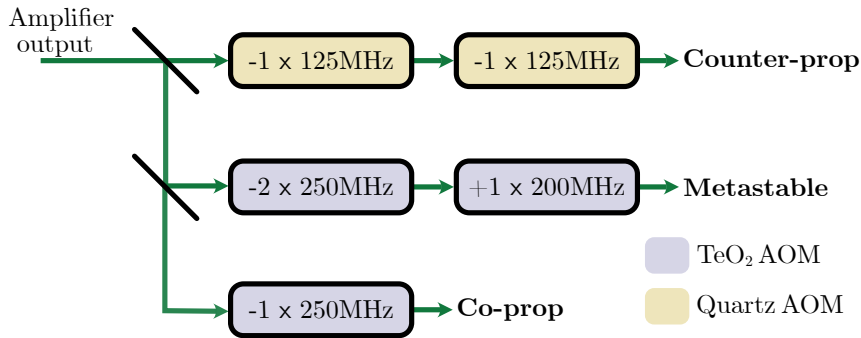


Figure 4.7: Diagram of the AOMs used to switch and frequency control the counter-propagating, co-propagating and metastable beams. The multiplier of the AOM frequency indicates the order of the beam and if the AOM is a single- or double-pass. The metastable beam is centred 50 MHz away from the other two. This frequency difference allows us to drive the metastable transition shown in Fig. 4.5. When targetting a specific transition, we adjust the frequency of the double-pass AOM in the metastable beam path.

co- and counter-propagating beams, allowing us to drive the metastable transitions shown in Fig. 4.5. The majority of the light is sent to the Quartz AOMs for the counter-propagating beam, which is the most powerful. We use Quartz AOMs in this case as they allow us to work with very powerful beams (>1 W of continuous power with ~ 1 mm radius). The high power allows us to produce a high intensity, despite a large beam radius causing the counter-propagating beam to act globally on all ions. The metastable and co-propagating beams are intended to be addressed beams in the long term (see Sec. 5.1). When addressed they have a considerably smaller spot-size, and so a lower power can be used to achieve the same intensity. As a result, we do not need to use Quartz AOMs and can utilise the higher bandwidth and less polarisation-sensitive TeO_2 AOMs.

4.3 Calibrating single qubit operations

When calibrating the single-qubit operations, we must know both the resonant frequency of the atomic transition, ω_0 , and the Rabi frequency, Ω , with which it is driven for a given laser intensity. Inaccuracies in each of these parameters contribute significantly to coherent or systematic errors. In our experiment, we use robust phase estimation (RPE) [70, 80] to calibrate both the qubit resonance and the Rabi

frequency. Before the details and advantages of RPE are given, it is important to understand its simpler form: Ramsey spectroscopy.

4.3.1 Ramsey spectroscopy

Consider a qubit with resonance ω_0 , and a driving laser slightly off resonance with angular frequency ω . After transforming into the rotating frame of the drive, the free atom Hamiltonian is given by

$$\tilde{H}_A = \frac{\hbar}{2}\delta\sigma_z, \quad (4.31)$$

where $\delta = \omega_0 - \omega$. We can see that even when idle, a Z rotation is being applied, with amplitude proportional to the detuning error δ . We can utilise this idle rotation to measure δ and correct $\omega \rightarrow \omega_0$ (or as close as possible to ω_0 within experimental constraints).

With the qubit in the $|+\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$ state, waiting some time τ will result in a phase shift $|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + e^{i\phi}|\downarrow\rangle)$, where $\phi = \delta\tau$. Subsequently performing a rotation $R(\pi/2, \phi_M)$ into a given measurement basis then gives a probability of measuring $|\downarrow\rangle$ of

$$P(|\downarrow\rangle) = \frac{1}{2}(1 + \cos(\phi - \phi_M)). \quad (4.32)$$

Utilising several different measurement angles ϕ_M to measure in the $X, Y, -X$, and $-Y$ bases, provides an estimate of ϕ . For a given measurement time τ , we can be certain in ϕ up to aliasing, meaning we know $\phi = \tilde{\phi} + 2n\pi$. The integer n represents the number of full rotations we have aliased through. This method of estimating ϕ raises two important and competing questions:

1. How do we choose τ to ensure $n = 0$?
2. How do we choose τ to give maximum sensitivity to δ ?

The questions are competing: increasing τ allows the phase ϕ to accumulate more significantly, which improves sensitivity to δ and enables full use of the measurement range $\phi \in (0, 2\pi]$. However, if τ becomes too large, the accumulated phase will

exceed 2π , leading to aliasing. In this case, we may infer the wrong value of δ , potentially underestimating it by up to 100%!

If we settle on a single value of τ , which does not cause aliasing, the uncertainty of our measurement will scale with the *standard quantum limit* (SQL) [81, 82]. This is the statement that binomial projection statistics from N samples of identical qubits in the same measurement basis poses a limit on the accuracy of the measurement as

$$\sigma^2(\downarrow) = \frac{P(\downarrow)(1 - P(\downarrow))}{N}, \quad (4.33)$$

meaning that the standard deviation of δ will scale as $\sigma(\delta) \propto 1/\sqrt{N}$. Such scaling means that attempts to calibrate our experimental parameters with increasing precision result in diminishing returns, and thus a trade-off between lengthy calibration routines, or imprecise parameters. The solution to this, RPE, offers a set of choices of τ which allow for the much more favourable Heisenberg scaling; $\sigma(\delta) \propto 1/N$. While we have focussed on measuring detuning errors, the same logic applies to errors in pulse area, or in Rabi frequency.

4.3.2 Robust phase estimation

RPE utilises an exponential scaling in the time τ to place a bound on the value of δ . Such scaling allows us to increase ϕ beyond 2π while retaining information about its true value. Eventually, a limit is reached in how much τ can be scaled before incoherent errors have detrimental effects on the measurement. At this limit, the final measurements provide an uncertainty $\sigma(\delta) \propto 1/N$, with Heisenberg scaling [80].

This scheme has been extended in the literature to include adaptive techniques [83, 84], which can give further speed-ups in estimation time. However, our experiment utilises an sCMOS camera for readout (see Ref. [3] for details), with images arriving asynchronously over USB⁹. As such, the static RPE protocol is perfect for our use case.

⁹It is possible to process these images in real-time using an FPGA and the CoaXPress data standard implemented by the ORCA-Fusion BT C15440-20UP. An implementation of this is currently in development in Oxford.

Returning to Eq. 4.32, we can again insert four different values of $\phi_M = 0, \pi/2, \pi, 3\pi/2$, which correspond to measurements in the $X, Y, -X$, and $-Y$ bases, respectively. For a small time $\tau = \tau_0$, measuring in each basis allows us to estimate ϕ . Taking more measurements for a given value of τ will increase our certainty of ϕ according to the SQL. Instead, we retain the estimate of ϕ , but allow for a large uncertainty of $\pm\pi/2$, well within our measurement certainty. Once this bound is established, it is best to increase τ exponentially: $\tau = 2^{i-1}\tau_0$ for iterations $\{1, 2, \dots, K\}$. We go into the next iteration with a maximum permissible region within which ϕ could lie. Each iteration then places a further bound on ϕ , each time to within $\pi/2^{i-1}$. At each iteration, a phase Φ_i is calculated from the measurements in each basis as,

$$\Phi_i = \arctan\left(\frac{P_{\pi/2} - P_{3\pi/2}}{P_0 - P_\pi}\right). \quad (4.34)$$

This is then normalised against the length of this iteration as,

$$\phi_i = \frac{1}{2^{i-1}}\Phi_i. \quad (4.35)$$

The normalised phases ϕ_i are then used to calculate the final value of ϕ with Algorithm 1. A visualisation of Algorithm 1 is provided in Fig. 4.8. The true phase

Algorithm 1 Estimate Phase from ϕ_i

```

function ESTIMATEPHASE( $\{(\phi_i)\}_{i=1}^K$ )
     $\phi = 0$  ▷ Initial estimate
    for  $i = 1$  to  $K$  do
         $\Delta = \frac{\pi}{2^{i-1}}$  ▷ Maximum allowed shift
        while  $\phi_i < \phi - \Delta$  do
             $\phi_i = \phi_i + 2\Delta$ 
        end while
        while  $\phi_i > \phi_{i-1} + \Delta$  do
             $\phi_i = \phi_i - 2\Delta$ 
        end while
         $\phi = \phi_i$  ▷ Update current estimate
    end for
    return  $\phi$ 
end function

```

we aim to measure with the initial Ramsey duration is 1 rad, and the permissible

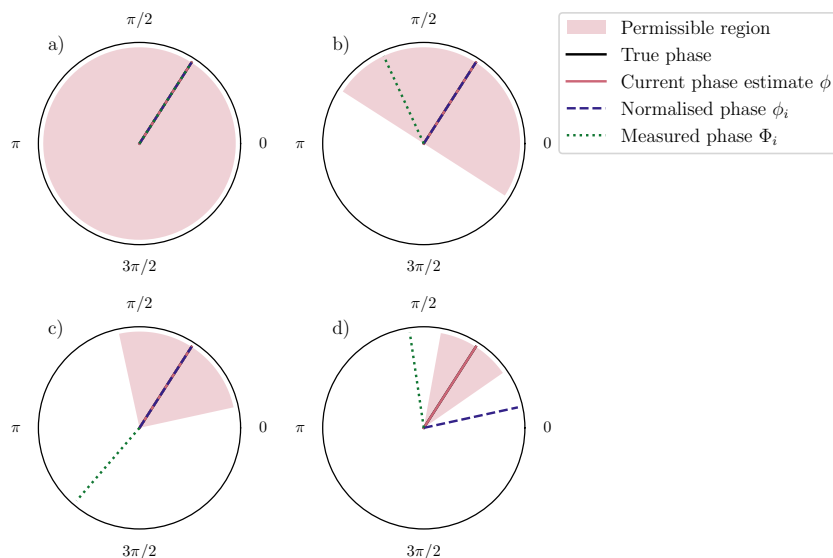


Figure 4.8: A visualisation of the RPE algorithm. The true phase we aim to measure is 1 rad. There is no noise in this simulation and so measurements are perfect. We add noise for illustrative purpose in Fig. 4.9. (a) The permissible region is the entire 2π range as we have no prior assumption of ϕ . After a first measurement the measured phase, normalised phase and current phase estimate all align and are the same as the true value. (b) The permissible region is now centred our phase estimate from the previous iteration with a range of $\pm\pi/2$. The Ramsey interrogation time has doubled, hence, doubling the measured phase. This is normalised and becomes the same as the true phase and current estimate using Eq. 4.35. (c) The same applies again in the next step as no aliasing has occurred. (d) The measured phase has aliased around the 2π region, causing the normalised phase to lie outside of the permissible region. Algorithm 1 then shifts this into the permissible region, giving the current phase estimate which is aligned with the true phase.

region is the entire 2π range, as we have no prior assumption of ϕ . In Fig. 4.8(a), the true phase, current estimate, normalised phase, and measured phase are all the same. In reality this would not be the case and we will explore this in a moment. In the second iteration (Fig. 4.8(b)), the measured phase has doubled, owing to the doubling of the interrogation time. However, no aliasing has yet occurred and so the all other values overlap. This behaviour is repeated in Fig. 4.8(c). Finally, in Fig. 4.8(d), the phase wraps around the 2π boundary, and the normalised phase lies outside of the permissible region from previous measurements. In this instance, Algorithm 1 shifts the phase into the permissible region causing the measured phase to once again overlap with the true value.

In Fig. 4.9, we add some randomness added to the values Φ_i to illustrate a real scenario with shot noise, and prevent all values from overlapping after just a few iterations in the diagram. The true value is again 1 rad, but this time we add uniform noise to each measurement of $\pm\pi/4$ rad¹⁰. We see that the RPE algorithm is resilient to this noise due to the exponential scaling of the interrogation times, allowing the impact of this noise on the final estimate to be greatly reduced. Note that this noise encompasses statistical variations, but not temporal variations. Any temporal fluctuation on the true value of ϕ could have a more adverse impact on the measurement – particularly at long interrogation times – and so a high degree of stability in ϕ is a core assumption of RPE. This visualisation also shows how the algorithm scales so favourably, as the measured phase Φ_i retains access to, and therefore sensitivity across, the entire 2π region, which is mapped to the permissible region with the range $\frac{2\pi}{2^{i-1}}$.

This procedure makes the assumption that the true value of ϕ (or δ) does not change with time or experimental parameters. For a frequency offset error over a short (~ 1 hr) timescale, this is a perfectly valid assumption. Additionally, the procedure can be implemented in the X, Z basis to calibrate pulse area errors. In this scenario, instead of waiting some time τ , we implement an X rotation for a time τ . Scaling this exponentially allows us to calibrate the time required to implement a $\pi/2$ - or π -pulse with Heisenberg scaling¹¹. Instead of estimating ϕ , we are now estimating some parameter α , where

$$P(|\downarrow\rangle) = \frac{1}{2} \left(1 + \cos(\theta(1 + \alpha) + \phi_M) \right), \quad (4.36)$$

and θ is the target rotation angle.

RPE calibration in frequency and pulse duration forms the backbone of our single-qubit calibrations. It is additionally used to measure light shifts, both to

¹⁰Binomial shot noise simulated from individual measurements would provide a more realistic scenario. As this is just for illustrative purposes the uniform noise will suffice.

¹¹Any pulse length could be calibrated within the bound $(0, 2\pi)$.

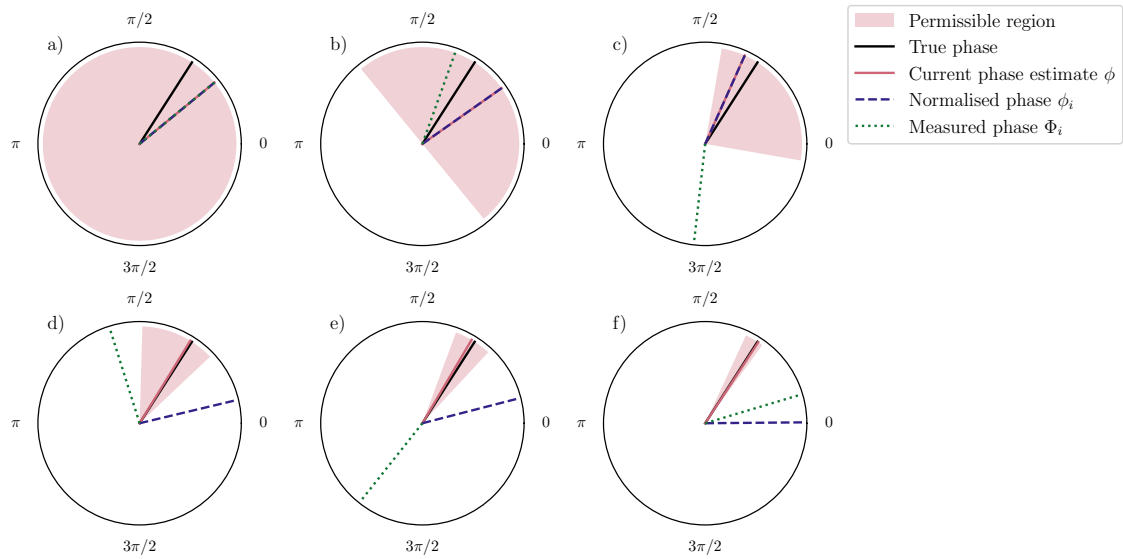


Figure 4.9: Simulation of the RPE protocol with added noise. Each measurement in the 2π range has uniform noise of $\pm\pi/4$ rad added. Each sub-figure represents an iteration similar to Fig. 4.8. The measured phase Φ_i is normalised according to Eq. 4.35 and then shifted into the permissible region using Algorithm 1 giving a phase estimate. We observe that even in the presence of shot noise, an accurate estimate of ϕ (and more importantly δ) can be achieved thanks to the exponential scaling of the interrogation times.

characterise our addressing system, but also calibrate the correct beam intensity to apply for the two-qubit gates in Sec. 7.7.4.

Ramsey example - metastable qubit

The metastable qubit is neither magnetically sensitive (to first order), nor does it suffer from common-mode noise on the driving lasers. As a result, contrast in the Ramsey measurements can be maintained over a very long duration – the qubit has a long coherence time (4.5(7) s; see Fig. 4.16). This coherence time will bound the accuracy with which we can perform RPE while still measuring meaningful results. As an upper bound, the minimum change in frequency that can be measured is $1/\tau$ where τ is the maximum Ramsey duration. An example result set is shown for the metastable clock frequency calibration in Fig. 4.10. The longest Ramsey sequence in this example is 100 ms, allowing for a frequency resolution of ~ 10 Hz. From the data, we estimate a phase offset of 506 Hz. Bootstrapping the data with binomial

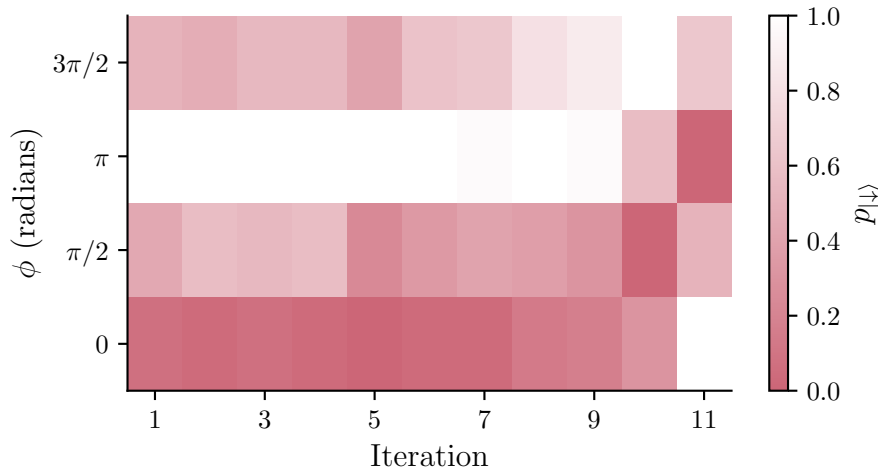


Figure 4.10: An example RPE frequency calibration with the metastable clock qubit. The maximum Ramsey interrogation time is 100 ms. In the ideal case of zero detuning error the pattern would not deviate from that in iterations $1 \rightarrow 9$, i.e. no phase accumulation. The estimated detuning error is 506 Hz. Bootstrapping the data with binomial statistics estimates that the true value is between 496 Hz and 516 Hz.

statistics indicates that the true value lies between 496 Hz and 516 Hz.

Rabi example - optical qubit

We also use RPE to calibrate the optical qubit’s parameters. An example $\pi/2$ -time and π -time calibration is shown in Fig. 4.11 – we refer to this a ‘Rabi’ calibration. Calibrating the π -time could be considered redundant, as we can just implement two $\pi/2$ pulses. However, we choose to perform this extra calibration as it allows us to calculate any dead time in the switching of the lasers by comparing the two results. Knowing this dead time, we can then interpolate the pulse duration to perform any arbitrary rotation on the Bloch sphere. In Fig. 4.11, the current pulse durations are very close to their ideal values. The high sensitivity achieved from the exponential scaling means we are able to measure that the π -pulse duration should increase by a factor of 1.002(2), and the $\pi/2$ -pulse duration should decrease by a factor of 0.985(5). As with the Ramsey experiment, the error is estimated by bootstrapping the data with binomial statistics.




Figure 4.11: An example Rabi RPE calibration with the optical qubit. The current pulse durations are very close to the idea;. (a) $\pi/2$ -pulse calibration, the suggested change is a factor of $0.985(5)$. The difference in structure to the Ramsey RPE is attributed to the phase required to return to $|\downarrow\rangle$. In the ideal case this is the same for all Ramsey iterations, but this is not the case when calibrating pulse durations. (b) Data for a π -pulse calibration. The duration should change by a factor of $1.002(2)$.

It is also worth noting that the structure of the data for a Rabi frequency/ π -time calibration is different from a Ramsey calibration. In the ideal Rabi case, the phase between the first few points changes until the gate implemented has a rotation angle of $2n\pi$. As such, the phase that returns to $|\downarrow\rangle$ on iteration 1 of the $\pi/2$ calibration is $3\pi/2$ (see Fig. 4.11a). On the second iteration, this is π , and then 0 or 2π for all subsequent iterations. The structure is similar for the π -pulse calibration in Fig. 4.11(b), with the first iteration being the equivalent to the $\pi/2$ -pulse's second iteration.

4.4 Randomised benchmarking

Now that we have the ability to calibrate and perform single-qubit operations on our optical and metastable qubits, it is useful to benchmark how well these operations are implemented. Ideally, we would measure the accuracy of each unitary \hat{U} applied to an arbitrary state ρ . This measurement can be made using quantum process tomography (QPT) which provides a full characterisation of a quantum gate. However, this requires a very large number of measurements, which scales exponentially with the number of qubits. We can provide an average measure of the fidelity of single-qubit operations using Randomised Benchmarking (RBM), which requires fewer resources and scales polynomially with system size¹².

The implementation outlined in Ref. [71] utilises the Clifford group as a unitary 2-design to construct a small set of operations, which transforms the average noise process we are interested in (responsible from a reduction in unit fidelity) into a depolarising channel [85].

The protocol is summarised as follows. We generate a sequence of gates \hat{S} from the Clifford group

$$\hat{S} = \prod_{i=1}^N \hat{C}_i, \quad (4.37)$$

where the final Clifford gate \hat{C}_N is chosen such that it inverts the previous gates in the sequence:

$$\hat{C}_N = \left(\prod_{i=1}^{N-1} \hat{C}_i \right)^{-1}. \quad (4.38)$$

As a result, the quantum channel we have ideally constructed is the identity operator such that

$$\hat{S}\rho\hat{S}^\dagger = \rho. \quad (4.39)$$

¹²This scaling is not a concern for our single qubit characterisation, but will become useful as we develop the system further.

However, the channel we implement is never ideal in reality. After the process is “twirled” over the Clifford group, we find an average noise process $\bar{\Lambda}$, which for a system of n qubits is

$$\bar{\Lambda}(\rho) = p(m)\rho + (1 - p(m))\frac{\mathbb{1}}{2^n}, \quad (4.40)$$

where there is some probability $1 - p(m)$ of depolarising ρ into the maximally mixed state [41], and a probability $p(m)$ that the state ρ survives over a sequence of length m . It is this survival probability that we use to characterise the average fidelity of Clifford operations.

With the system initially prepared in the state $|\psi_0\rangle = |\downarrow\rangle^{\otimes n}$, the probability of measuring the state $|\psi_0\rangle$ after a sequence of m Clifford operations is given by

$$P(|\psi_0\rangle) = Ap^m + \frac{1}{2^n}, \quad (4.41)$$

where A is a parameter accounting for state preparation and measurement (SPAM) errors. From this, an average error per Clifford, $\bar{\epsilon}$ (proportional to gate infidelity), can be extracted as

$$\bar{\epsilon} = \left(1 - \frac{1}{2^n}\right)p. \quad (4.42)$$

4.4.1 Metastable qubit RBM

We implement the RBM protocol outlined on the metastable clock qubit. The results are plotted in Fig. 4.12. We measure an average error per Clifford of $1.10(7) \times 10^{-4}$. This represents a reasonable fidelity and well within the expectation for laser-based gates.

To ensure that this result is reproducible, and check for drifts in the system, we repeat this benchmarking over six hours and observe a very small but measurable increasing trend in error. This increase corresponds to a $1.6(4) \times 10^{-6}$ increase in the error per Clifford per hour. There are a few sources of error likely contributing to this error rate, which we will discuss in the following section.




Figure 4.12: Randomised benchmarking of the metastable clock qubit with optimal parameters. Up to 5000 Cliffords are performed, with data taken for 10 random seeds. We measure an error per Clifford of $1.10(7) \times 10^{-4}$.




Figure 4.13: The error from randomised benchmarking of the metastable qubit over a period of six hours. The results show remarkable consistency, illustrating the frequency stability of the qubit. There is a marginal increase in error over time, likely due to slight beam pointing drift adjusting the optimal $\pi/2$ -pulse duration.

Pulse area error

The pulse area may vary in two ways. Firstly, a miscalibration of the $\pi/2$ -pulse duration may result in a systematic error in the pulse duration. Alternatively, there may be timing or amplitude variation between shots. As shown in Fig. 4.14, the ideal $\pi/2$ -pulse duration is $24.8 \mu\text{s}$. This is several orders of magnitude larger than the 8 ns timing resolution of our control system, meaning variations in timing can be considered negligible. By considering the fidelity of the $\pi/2$ -pulses used

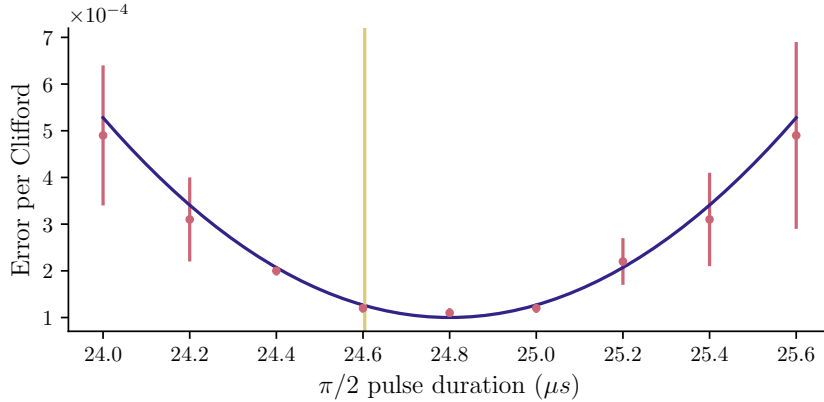


Figure 4.14: Error per Clifford from RBM vs $\pi/2$ -pulse duration. The minimum error measured is $1.10(7) \times 10^{-4}$. The RPE $\pi/2$ -time calibration is marked with the yellow line, indicating that we calibrate to within $0.2 \mu s$ of the ideal value, which contributes an measurable error at the 10^{-5} level.

to implement the Clifford operations as

$$\mathcal{F} = \frac{1}{4} \text{Tr}(|U_{\text{ideal}} U_{\text{impl}}^\dagger|^2), \quad (4.43)$$

we find that the error per Clifford is proportional to $(\Omega \times \delta t)^2$ where Ω is the Rabi frequency and δt is a systematic variation in timing. This can be used to find any systematic errors in the calibration of the $\pi/2$ -pulses. We plot the result as the $\pi/2$ -pulse duration is scanned in Fig. 4.14 and the calibrated pulse duration is marked $-0.2 \mu s$ away from the ideal duration. This contributes an error at the 10^{-5} level. In addition to systematic errors, there are also amplitude fluctuations within and between pulses which contribute an error. We estimate the Raman beams to be stable to within 1% between pulses. This is based upon the $\sim 1\%$ drift we see in the co-propagating beam over 50 shots in Chapter 7 combined with a pessimistic view from a lack of measurement in the metastable beam. As such, the error contribution from amplitude instability could be as high as 5×10^{-5} [24], potentially the leading source of error in our metastable qubit.

Detuning error

Similarly, there may be both systematic and stochastic errors in the qubit frequency compared to the driving frequency. Taking a similar approach to Eq. 4.43, the

error from a systematic detuning error of δ is proportional to $\frac{\delta^2}{\Omega^2}$. In Fig. 4.15, a measurement of this error is shown. Note that the optimal frequency is 550 Hz away from the calibrated frequency. This amounts to a substantial error of $\sim 2 \times 10^{-3}$ if not corrected. The cause of this is likely the light shift induced by the Raman lasers when actively driving the qubit. This light shift is not picked up by our Ramsey/RPE calibration, as the qubit remains idle in these experiments. As a result, further calibration is required of these light shifts to account for the change in frequency between the qubit being active vs idle. Such calibrations are presently being implemented in our system, utilising another variation of RPE.

This issue with idle vs active frequency also highlights another source of error. The beams are not always applied to the ion. There is a $3 \mu\text{s}$ setup time when switching phase with SUServo¹³, as well as an additional $0.5 \mu\text{s}$ to allow our hardware to meet timing restrictions. We directly simulate this effect with RBM sequences and find that there is an average error contribution per Clifford for a 550 Hz offset during qubit dead time is $\sim 3 \times 10^{-5}$. This is not insignificant, and correcting for it would yield a measurable improvement in qubit performance.

Stochastic variations of the qubit frequency are considered negligible, as the low magnetic field sensitivity of the clock qubit ensures that the qubit frequency is highly stable. Noisy light shifts could also contribute to gate-to-gate frequency variations in the qubit, though we have measured the beam amplitude stability to be better than 1% over the short term. This should contribute light shift errors at the 10^{-7} level, though it will more significantly impact the Rabi frequency.

Coherence time

The error per Clifford from decoherence of the qubit is

$$\epsilon \approx \frac{2.2 t_{\pi/2}}{3 T_2^*} \quad (4.44)$$

¹³SUServo is a combination of a Sampler (<https://github.com/sinara-hw/Sampler>) and an Urukul (<https://github.com/sinara-hw/Urukul>) used to create a PID loop for amplitude stabilisation.




Figure 4.15: Error per Clifford from RBM vs detuning. Detuning is measured relative to the frequency calibrated from RPE. The point of minimum error is at 550 Hz, which would mean a 2×10^{-3} error if the RPE calibration values are used directly. The discrepancy likely occurs due to the difference between the active qubit (beams on) frequency when light shifted and the idle qubit (beams off) frequency we calibrate to.

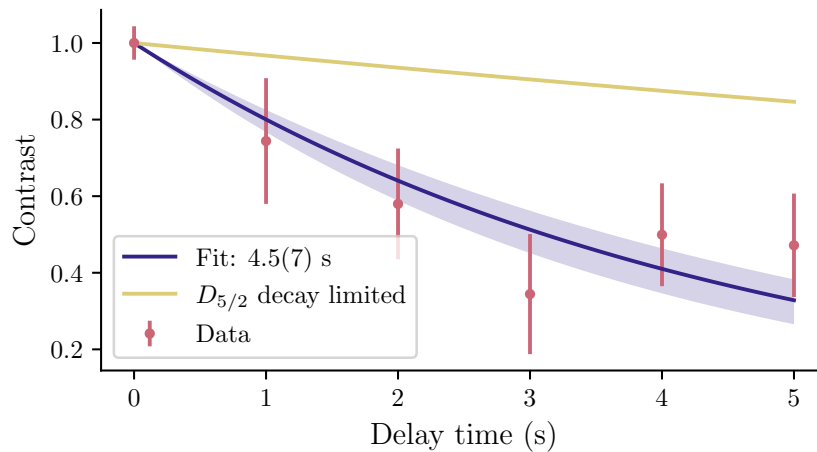


Figure 4.16: Metastable clock qubit coherence time measurement. We measure the contrast of Ramsey experiments at increasing Ramsey durations. The coherence time is 4.5(7) s and fits an exponential decay indicating the noise profile is not coloured [41]. We also plot the contrast limit due to the 30 s lifetime of the $D_{5/2}$ level in yellow. There are improvements to be made to the qubits coherence time before this becomes a limiting factor.

where $t_{\pi/2}$ is the $\pi/2$ -pulse duration and T_2^* is the coherence time [24]. In Ref. [47], the clock coherence time was measured to be 9.1(1.0) s when driving with the Raman lasers. Since rebuilding the apparatus, the coherence time has dropped to 4.5(7) s, as shown in Fig. 4.16. The $\sim 50\%$ drop in coherence time indicates that this is

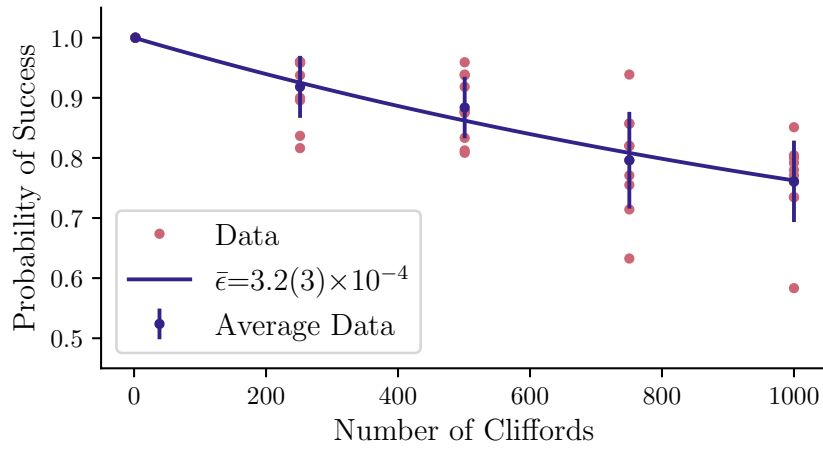


Figure 4.17: Result of RBM with the optical qubit. The error per Clifford $\bar{\epsilon} = 3.2(3) \times 10^{-4}$. Up to 1000 Clifford are performed with 10 Random seeds at each point.

likely limited by magnetic field noise, with the magnetic field having changed and being slightly off the 3.333 G clock point. This translates to an error contribution of $\sim 4 \times 10^{-6}$, which is negligible compared to other sources for the time being.

4.4.2 Optical qubit RBM

We also benchmark the optical qubit using the same methods. The result for a single benchmarking run are in Fig. 4.17. The error is $3.2(3) \times 10^{-4}$, higher than the metastable qubit driven by Raman transitions.

Such an error would still not be a limiting factor in many systems, where 10^{-3} two-qubit gate errors are synonymous with very high performance [76] (though the best achieved is at the 10^{-4} level[25]). This is especially true in the NISQ regime, where we have 10s of qubits and implement 100s of gates.

The error rate is due to a combination of factors, some of which differ from the metastable qubit. This can be seen by the instability of the benchmarking results in Fig. 4.18. The error drifts from $\sim 5 \times 10^{-4}$ down to 2×10^{-4} over an hour, and then returns back to the higher value after a further hour. This is starkly different from the 10^{-5} level stability of the metastable qubit over six hours. The higher instability is likely a result of driving the qubit with a single laser source. This makes it susceptible




Figure 4.18: Error per Clifford from repeated RBM runs over a 2 hour period. There is a significant deviation in error between 2×10^{-4} and 6×10^{-4} . This drift is not consistent with the cavity drift implying shorter term fluctuations in either laser/qubit frequency or beam pointing. The gap in data around 1 hour is due to the 1762 nm laser cavity lock failing.

to phase noise in the laser, as well as frequency drifts from changes in lock parameters. These errors do not occur when driving with Raman lasers as the qubit's frequency and phase is determined by the highly stable RF drive from our control system.

Pulse area error

As with the metastable qubit, both systematic pulse area errors and amplitude noise can contribute to the total error. We again plot the error per Clifford as a function of the $\pi/2$ -pulse duration. The results are shown in Fig. 4.19. There is a small $0.06 \mu\text{s}$ difference between the calibrated duration (marked with the yellow vertical line) and the ideal value measured from benchmarking.

We have measured the stability of the 1762 nm laser across 50 shots using a photodiode separate from the one used in the amplitude stabilisation loop. We plot the result in Fig. 4.20 and find that the pulse area has a standard deviation of 0.23%. The Rabi frequency scales as $\Omega \propto \sqrt{I}$, where I is the laser intensity. Using the error contribution from Ref. [24],

$$\epsilon = \frac{1}{2} \left(\frac{\delta\Omega}{\Omega} \right)^2, \quad (4.45)$$




Figure 4.19: Error per Clifford vs $\pi/2$ -pulse duration in the optical qubit. We fit a quadratic dependence between the pulse duration and error per Clifford. The RPE calibration marked with the yellow line is $\sim 0.5 \mu s$ off from the ideal value. This slight miscalibration adds a 7×10^{-5} error.




Figure 4.20: Amplitude stability of the 1762 nm laser over 50 pulses. The standard deviation of the amplitudes is 0.23%. At this level we can expect 10^{-6} errors from amplitude fluctuations between pulses. We do not yet have data for long pulse sequences, like the 1000 Clifford gates we implement in RBM. If the errors are non-Markovian this may present differently. We do not expect this due to the active stabilisation performed by the our control system while the laser light is applied to the ion.




Figure 4.21: Error per Clifford vs detuning of the 1762 nm laser. The optimal detuning is 50(10) Hz away from the RPE calibration. Additionally, due to the ~ 50 ms coherence time of this qubit (see Fig. 4.22), this accuracy is roughly in line with the limit that is achievable from a Ramsey experiment with this qubit.

then the error from this amplitude noise will be at the 10^{-6} level – a relatively negligible contribution.

Detuning error

Again repeating the analysis with the metastable qubit, we measure the error per Clifford while adjusting the laser detuning. The result is shown in Fig. 4.21. There is a 50(10) Hz offset from the calibrated frequency, which contributes an error of 1.2×10^{-5} .

Coherence time

The coherence time of the optical qubit is significantly lower than that of the metastable qubit. This is because the optical qubit is driven by a single laser source, which is susceptible to common mode noise. To illustrate this, we measure the coherence time of the optical qubit and a metastable qubit with significantly higher magnetic field sensitivity ($|D_{5/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 1, m_F = 0\rangle$). The result is shown in Fig. 4.22. If the dephasing in the optical qubit were limited by magnetic field noise, we would expect the optical qubit to have a coherence

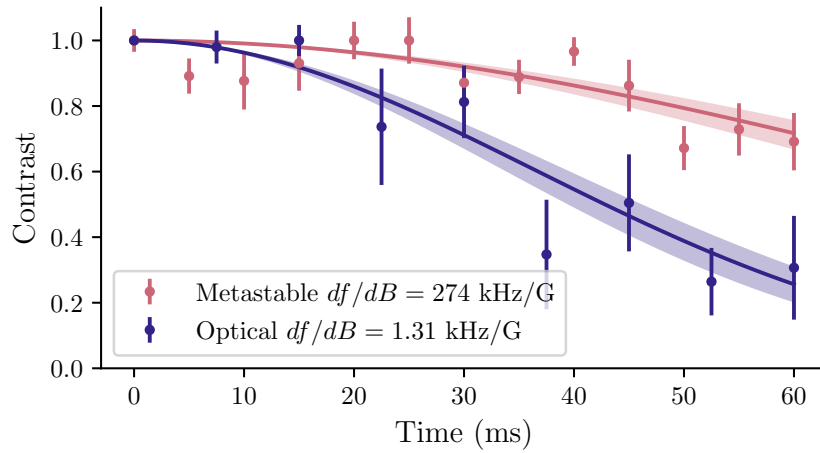


Figure 4.22: Coherence time of the Optical qubit, as well as a metastable qubit (not the Clock) for comparison. The metastable qubit is the $|D_{5/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 1, m_F = 0\rangle$ transition and has a first order magnetic field sensitivity of 274 kHz/G. The optical qubit has a coherence time of 51(4) ms compared to the 104(10) ms coherence time of this metastable qubit. Both fit to an $e^{-(t/\tau)^2}$ decay. Given the $209\times$ higher magnetic field sensitivity of the metastable qubit, we can infer that the optical qubit's coherence time is not limited by magnetic field fluctuations. As such, improvements to the phase noise on the laser are the likely target for improvements to coherence time.

time several orders of magnitude longer than that of the magnetically sensitive metastable qubit. Instead, the optical qubit has a coherence time of 51(4) ms, compared to 104(10) ms in the magnetically sensitive metastable qubit. As such, we are laser-limited, and the error due to decoherence is 1×10^{-4} . This makes coherence time a leading contributor to errors in the optical qubit, meaning that extending the coherence time should substantially reduce these errors. One likely cause of phase noise with the 1762 nm laser is path length jitter in the 10 m fibre that delivers the laser light between rooms. This fibre lies between the AOM and EOM in Fig. 4.3. Phase noise of this form can be mitigated with methods such as fibre-noise-cancellation (FNC), and this work is ongoing as I write¹⁴.

¹⁴Ollie, I hope your summer project has gone well and that we have FNC by the time you read this!




Figure 4.23: Optical qubit frequency tracked over a 40+ hr period. We measure a frequency drift of $-38.4(4)$ mHz/s. This corresponds to a relaxation in the length of the high-finesse cavity of $\sim 10^{-16}$ m/s.

Cavity drift

The instability of the optical qubit RBM highlights a technical challenge with implementing optical qubits. To drive the quadrupole transition we use a single laser beam with the source locked to a high-finesse cavity. This system is susceptible to common-mode drift from fluctuations and drift in the cavity free spectral range (FSR). Over the course of 40+ hrs we track the frequency of an optical qubit, and measure a constant drift of $38.4(4)$ mHz/s. As such, over the course of an hour the driving laser for the qubit will drift by ~ 100 Hz, contributing a small error. As an aside, the measurement of this frequency drift is a rather remarkable result. Estimating the length of the cavity to be ~ 1 m, the frequency drift corresponds to an expansion of the cavity length of 10^{-16} m/s or 0.1 protons/s!

Ion temperature

Another factor in driving the optical qubit with a single laser source is the susceptibility to ion temperature. This is not present in the metastable qubit when driven by co-propagating Raman beams because $\Delta\vec{k} = 0$, and therefore $\eta = 0$. The excitation n in a thermal state with mean excitation \bar{n} is given by

$$P(n, \bar{n}) = \frac{1}{\bar{n} + 1} \left(\frac{\bar{n}}{\bar{n} + 1} \right)^n. \quad (4.46)$$




Figure 4.24: Simulation of the error contribution from ion temperature using Eq. 4.48 for the optical qubit. We use a mode frequency of $2\pi \times 1.6$ MHz and Lamb-Dicke parameter $\eta = 0.012$.

Using the error per Clifford for a small change in Rabi frequency [24]

$$\epsilon = 2.2 \times \frac{1}{2} \left(\frac{\delta\Omega}{\Omega} \right)^2, \quad (4.47)$$

the error from a mean thermal occupation \bar{n} , assuming the qubit is calibrated to the ground-state Rabi frequency, is

$$\epsilon = 2.2 \times \frac{1}{2} \sum_n P(n, \bar{n}) \left(\frac{\Omega_n - \Omega_0}{\Omega_0} \right)^2 \quad (4.48)$$

$$\Omega_n = \Omega_0 e^{-\eta^2/2} \mathcal{L}_n^0(\eta^2)$$

and the Rabi frequency scales as shown in Sec. 4.1.2. We simulate this error using a Lamb-Dicke parameter $\eta = 0.012$ due to the 1762 nm beam coupling to the radial LF mode with frequency $\omega = 2\pi \times 1.6$ MHz and the results are shown in Fig. 4.24. The simulation uses temperatures up to $\bar{n} = 10$ and includes a maximum Fock state of $n = 200$. Results indicate that the error due to finite temperature with only Doppler cooling will stay below the 10^{-5} level. The heating rate on a single ion in our trap on the simulated modes is ~ 50 phonons/s. Over the course of the RBM sequence, we perform up to 5000 Clifford gates. This will take $7.75 \mu\text{s} \times 2.2 \times 5000 = 85 \text{ ms}$ ¹⁵, resulting in a heating of ~ 5 phonons and contributing an error of up to 1.4×10^{-6} .

¹⁵7.75 μs is the duration of a $\pi/2$ -pulse and there are on average 2.2 $\pi/2$ -pulses per Clifford.

Additionally, ion temperature can contribute an error for sufficiently fast pulses. This is due to the ion's motion being at varying points within the full period for different $\pi/2$ -pulses in the sequence [24]. As a result, the effective electric field experienced by the ion in that duration is shifted slightly. In our case the $\pi/2$ -pulses span many motional mode periods, and thus this error contribution is negligible.

4.5 Converting between qubit types

To implement a scheme that utilises multiple types of qubits, it is often required to convert between different qubit types [45]. For example, if the key operational qubit is optical, it will be necessary to convert to a metastable qubit to protect particular qubits from mid-circuit measurements and in-sequence cooling. Not only must the conversion between qubit types faithfully move all population from one type to the other, it must also preserve phase coherence, as the state of the qubit is generally unknown¹⁶.

In Ref. [45], high-fidelity mapping between the $S_{1/2}$ and $F_{7/2}$ is achieved using $< 1 \mu\text{s}$ pulses of 411 nm (first stage to $D_{5/2}$) and 3432 nm (second stage to $F_{7/2}$) light, transferring $|\downarrow\rangle$ and $|\uparrow\rangle$ simultaneously. The fast conversion between qubits with low magnetic field sensitivity allows them to achieve 96.6(4)% fidelity per conversion.

In our system, we have not yet invested heavily into this mapping capability, relying on the flexibility of the 1762 nm laser to map between arbitrary states in the $S_{1/2}$ and $D_{5/2}$. As such, testing our system's capability to convert between qubit types is a critical step in characterising its suitability as a mixed-qubit-type system.

We test this mapping on the two different pairs of qubits outlined in Chapter 2, summarised here as follows:

- Clock qubit

$$\text{O: } |S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$$

$$\text{M: } |D_{5/2}, F = 3, m_F = +2\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = +2\rangle$$

¹⁶Phase coherence isn't required if the qubit is about to be measured in the current basis, since only populations are observed.

- Hybrid qubit

$$\text{O: } |S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = 0\rangle$$

$$\text{M: } |D_{5/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 1, m_F = 0\rangle$$

The choice of qubits greatly impacts the ease with which qubit types can be converted. In this instance, the Hybrid approach has the benefit of a single direct optical transition to map between the two qubit types. This comes with the sacrifice of the low magnetic field sensitivity of the metastable clock transition.

4.5.1 Clock qubit

To transition between the metastable clock qubit and the optical qubit, we use the high flexibility of the 1762 nm laser to drive the following transitions:

1. $|D_{5/2}, F = 3, m_F = +2\rangle \rightarrow |S_{1/2}, F = 1, m_F = 0\rangle$.

2. $|D_{5/2}, F = 2, m_F = +2\rangle \rightarrow |S_{1/2}, F = 2, m_F = 0\rangle$.

3. $|S_{1/2}, F = 1, m_F = 0\rangle \rightarrow |D_{5/2}, F = 2, m_F = 0\rangle$.

A diagram of this procedure is shown in Fig. 4.25. The multiple steps taken in the transfer mean that we inevitably transition through a highly magnetically sensitive qubit. This is the $|S_{1/2}, F = 1, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = +2\rangle$ transition, which has a magnetic field sensitivity of 3.14 MHz/G at 3.333 G. As such this qubit has a coherence time of <1 ms in our system due to high levels of magnetic-field noise. Thus, if the 1762 nm transitions used to transfer the qubits are not sufficiently fast, they will cause a significant error. In our system, the rate at which we can drive quadrupole transitions is limited by the optical power we can pass through the EOM which handles frequency and phase control (see Fig. 4.3), as well as by how small the beam waist can be before coupling to many ions becomes difficult. This limits the duration of our π -pulses to $\sim 10\text{--}100$ μs .




Figure 4.25: Procedure to convert from metastable qubit to optical qubit in both (a) the Clock configuration and (b) the Hybrid configuration. Each arrow represents a π -pulse with the 1762 nm laser. The Clock qubit requires three steps to map to a stable optical qubit. This requires transitions that are highly magnetically sensitive. To map from optical back to metastable, the steps can be applied in reverse. The Hybrid configuration only requires one mapping pulse, improving simplicity and reducing sensitivity to magnetic field fluctuations.

4.5.2 Hybrid qubit

The Hybrid qubit approach is much simpler. We drive the $|D_{5/2}, F = 1, m_F = 0\rangle \rightarrow |S_{1/2}, F = 2, m_F = 0\rangle$ transition to convert from a metastable to an optical qubit. As a result, a relatively low magnetic-field sensitivity is maintained throughout the conversion process, and there are fewer transitions in which a errors can occur.

4.5.3 Maintaining coherence

In general, we require the mapping process to have low error rate for an arbitrary quantum state. As a result, maintaining qubit coherence during the mapping process is a critical element of benchmarking. To test this, we embed the sequence of mapping pulses in a Ramsey-like experiment. We measure the population and phase of the final state by adjusting the phase of the closing $\pi/2$ -pulse in the Ramsey experiment. From this, we infer an error rate by comparing the measured state to the desired state $\frac{1}{\sqrt{2}}(|\downarrow\rangle + e^{i\phi} |\uparrow\rangle)$, where ϕ is a known phase offset induced by the light shift of the 1762 nm laser.

We implement several mapping iterations in both qubit types, measuring the Ramsey contrast and phase at each step. The results for both the Hybrid qubit and Clock qubit are shown in Fig. 4.26 (note “Clock qubit” refers to the named configuration detailed previously, not the concept of a clock qubit). In the Hybrid qubit, the phase accumulated is 2.279(5) rad/map. The contrast takes the form of e^{-x^2} , implying a compounding error as we perform more mapping pulses, and the contrast after just the first pulse is 1.00(1), which is as expected given our 10^{-4} level errors on similar qubits to the one via which we map. We simulate this directly with a unitary which has a small, consistent over rotation. Fitting the over rotation angle to the data via direct simulation of the unitary, we find an over-rotation of 0.0200(4) rad which corresponds to a gate error of 2×10^{-4} . This is again consistent with the RBM results in the previous section. After 10 maps back and forth, the contrast is 0.82(2). A lower contrast is measured when mapping the Clock qubit back and forth just once, 0.69(3). The rate of phase accumulation in the Clock qubit is also much higher at 5.0(7) rad/map, implying that the extra pulses involved cause a further shift on the qubit. After 10 pulses, the Ramsey contrast is almost completely lost: we measure 0.16(5).

Some revivals are notable in the Clock qubit case, indicating that there is either a coherent error that is potentially avoidable, or a particular noise spike at a given frequency. We also simulate the Clock qubit approach directly, and add an exponential decay to the simulation, which accounts for contrast lost due to inconsistent phase shifts. The coherent error in this case comes from a consistent over-rotation of 0.355(6) rad. This corresponds to a 6% error, significantly higher than our Hybrid qubit. The phase noise error has a decay constant of 8(1) maps, meaning that without any coherent errors, the loss in phase will reduce the contrast to $1/e$ after 8 maps back and forth.

This is likely due to the highly magnetically sensitive qubit through which we transition. The transition has a $\sim 200 \mu\text{s}$ coherence time, which poses a limit on the contrast that we are able to maintain. It can be overcome by either reducing

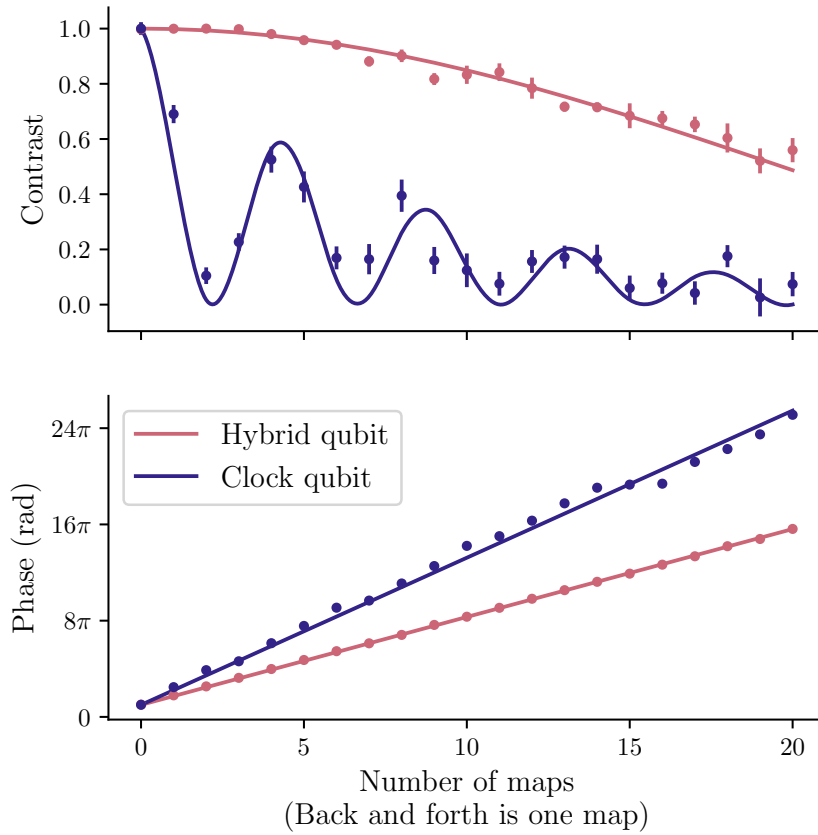


Figure 4.26: Results from mapping the Hybrid qubit and Clock qubit between their optical and metastable configurations. Top: Ramsey contrast for each configuration. The Hybrid qubit maintains contrast over a much longer period and the drop in contrast is dominated by a coherent error. The Clock qubit exhibits oscillations from a much stronger coherent error, as well as an exponential decay due to decoherence. Bottom: Ramsey phase. Both qubits increase in Ramsey phase as the number of maps increases. We measure the phase within the $0 \rightarrow 2\pi$ bound and infer the total phase from previous points.

magnetic-field noise or constructing a set of composite pulses to remove the sensitivity to these fluctuations.

In the above measurements, we have primarily focussed on the contrast. However, it is important that the phase shift be predictable to maintain coherent operations after a mapping sequence. In the bottom part of Fig. 4.26, we plot the measured phase along with a fitted value. We can then use the contrast and phase to calculate the state fidelity by comparing the predicted and measured values. After the first mapping pulse, the Hybrid qubit has a state fidelity of 0.991(7), whereas the Clock




Figure 4.27: State fidelity measurement after mapping the Hybrid and Clock qubit between their optical and metastable configurations. The Hybrid qubit exhibits significantly higher fidelity owed to the consistent low magnetic field sensitivity. The oscillations in the Clock indicate a coherent error which is overtop a decay from decoherence.

qubit has already dropped to 0.74(1). However, in a real computational sequence, we will not always be in an equal superposition state. As such, the susceptibility of the state to these mapping errors may increase or decrease. For example, with the Hybrid qubit, one state is left untouched, so there should be no mapping error. However, given the majority of the error is seemingly from population shifting, acting on the other pole of the Bloch sphere will have an increased error. These two should cancel to give similar behaviour to our equal superposition state.

The Clock qubit exhibits large errors from both decoherence and population loss. When acting on the poles, the decoherence should reduce, but there will still be significant error, as observed above, making this scheme far less viable than the Hybrid qubit at present.

4.6 Summary

We have implemented single qubit operations in both an optical and a metastable qubit informed by the analysis in Chapter 2. The metastable clock qubit boasts lower error rates with better long term stability, achieving errors between $0.9\text{--}1.3 \times 10^{-4}$

over 6 hours. However, there are still improvements to be made, in particular, our RPE calibration has a systematic error of 550 Hz, owed to the light shift changing the qubit frequency while active. Additionally, mapping the Clock qubit to an optical qubit proves to be challenging, making it incompatible in an *omg* scheme where we wish to implement the OTDF gate [4] in Chapter 7. Admittedly, this scenario is rather contrived, the Clock qubit can be used directly with an MS gate scheme instead [47]. In this case no coherent mapping would be required.

The optical qubit has slightly higher error rates, a large fraction of which can be attributed to its lower coherence time (51(4) ms to the metastable clocks 4.5(7) s). With the addition of fibre-noise-cancellation (FNC) we expect this to improve. The qubit is still susceptible to drift in the laser parameters which do not appear with two-photon Raman transitions. The upside to this qubit is that mapping to the metastable level for mid-circuit measurements works very well, making it compatible with the OTDF gate within an *omg* scheme.

*It's murder on the dance-floor
But you better not kill the groove*

— Sophie Ellis-Bextor

5

Individual addressing using a novel integrated photonic chip

Contents

5.1	Methods of individual addressing	77
5.2	Addressing apparatus	80
5.3	System characterisation	82

The *omg* architectures presented thus far address the issue of hiding qubits during global dissipative operations. However, addressing individual qubit operations is still required to perform select unitaries on different qubits, as well as select which qubits are converted between different types. A large-scale system must have many qubits of the same type at any point during operation, and thus require individual addressing to distinguish them.

In this chapter, we discuss a novel photonic chip for addressing individual qubits with very low crosstalk between ions in a chain. The chip is designed to work with 532 nm light, ideal for driving the Raman interactions and two-qubit gates detailed in Chapters 4 and 7, respectively. This makes the addressing system primarily designed to work with ground and metastable qubits, but not with the optical qubits

that have also been a focus. However, there are some options that would enable addressing of the optical qubit with this system. It is possible to use the 532 nm light to apply a light shift that changes the individual qubit transition frequencies. This enables addressing of optical-qubit operations by moving qubits in and out of resonance with the 1762 nm laser [75]. Alternatively, this light shift could be used to introduce ion-selective Z rotations, which can combine with global operations to form an individually addressed gate set.

5.1 Methods of individual addressing

A scalable approach to individual addressing is critical to creating a useful quantum computer. In the ideal case, a quantum system would have a driving field for each qubit with individual phase and amplitude control, as well as minimal cross-talk with other qubits. The driving fields should be independent, meaning that operations can run in parallel. If these requirements are met, along with all-to-all qubit connectivity, then a single layer of quantum gates can be run in constant time irrespective of system size. This is very desirable when designing a large scale quantum system, and could be considered a key part of DiVincenzo's first criterion [12]: *A scalable system of well characterised qubits.*

The approaches taken to addressing vary depending on the fields being used to perform gates: microwaves or lasers. The individual addressing of microwaves often involves using global interactions which vary in strength between qubits. This is favourable for rapid development with little engineering effort, but poses challenges in the time taken to optimise and implement a given pulse sequence [86]. Global microwave fields do not provide each qubit with individual phase and amplitude control, thus parallel operations are not possible. It is possible to achieve this in more complex architectures, such as the QCCD architecture [87], but this comes with engineering requirements that are beyond the scope of research environments. As such, from this point onwards we will focus our discussion to laser based approaches, which compare most naturally to our system. Within this, we will consider the individual

addressing of ions in a long chain. Addressing individual ions within a chain is very challenging. With the ions spaced by $\sim 5 \mu\text{m}$ we must focus the addressing beam to much smaller than this to minimise cross talk, and do so at each ion-site. This means that we are working near the diffraction limit, where standard optics will not suffice.

Before going into detail on our method of individual addressing, I would like to summarise a few other approaches which also focus down laser radiation, along with their pros/cons. Some addressing systems may trade-off individual phase and amplitude control to instead prioritise flexibility, reconfigurability, and low cross-talk. An example of this is to use a pair of crossed acousto-optic deflectors (AOD) to steer lasers onto each individual ion [88–90]. This achieves very low cross-talk and is highly flexible, as the beam steering can be recalibrated on the fly. The downside is that crossed AODs are best suited to serialised ion interactions. Adding extra tones to produce more spots results in unwanted cross-terms which produce erroneous laser spots reducing the total laser power delivered to the ions. This means driving fields are not independent, and operations are not fully parallelisable.

Alternatively, micro-mirror beam steering can be used in a similar manner [91–93]. This offers a highly reconfigurable and low cross-talk method of addressing. However, the reconfiguration time is typically on a timescale similar to the gate operation itself, adding significant overhead. Additionally, the approach in Ref. [91] and Ref. [92] relies on sequential operations, returning us to the parallelisation issue exhibited in AODs.

Micro-lens arrays can steer multiple spots and address ions in parallel with low cross-talk, making them also an excellent candidate for addressing long chains of ions [94]. Conceptually, this approach is very similar to ours, simply with different optical components to handle the splitting of light and delivery to each ion. We did not follow the approach in Ref. [94] as it has a very high insertion loss, limiting the available intensity at the ions for performing quantum operations. Additionally, it makes use of both a waveguide and micro-lens array, increasing the fabrication complexity of the system.

Finally, our approach is to image the output of a waveguide directly onto the ions [2]. In our system, combined with a network of fibre AOMs for individual phase and amplitude control, the waveguides allow for fast and parallel addressing with very low cross-talk. There is little overhead to scaling the ion chain, as adding more beams is as simple as writing a new chip and connecting more fibres to the network. The system is fully parallelisable with the complexity shifted to maintaining a large number of microwave control channels for each AOM – classical bookkeeping for each ion’s phase, as is required in all quantum devices aiming to scale. This classical control problem is one that modern hardware is very capable of tackling.

Waveguides have since been used for the direct delivery and collection of light for dissipative operations, such as the 650 nm laser light delivered to barium ions in Ref. [95]. Such an implementation could allow for the parallelisation of quantum networking operations, making waveguides a promising technology with several applications to trapped ions.

One downside of this approach is the lack of live reconfigurability. When developing an ion system, it can be useful to test different trapping potentials – for example when choosing an ideal mode structure for gates or to avoid noise spikes. If an addressing system cannot be adjusted on the fly, this may slow the development cycle of the system. As such, this makes waveguides and micro-lens arrays less flexible compared to AODs, though the boost to parallelisability is one we consider a worthwhile trade-off. Additionally, to be useful for ion chains of differing lengths, advanced trapping potentials are required to shift the ions into alignment with each beam. Without this, using a single waveguide for a variable number of ions in a harmonic potential causes some ions to sit between beams. In our system this is not a problem, as our 3D segmented microfabricated trap allows us to produce anharmonic potentials, providing the degree of freedom needed to control the ion spacing.

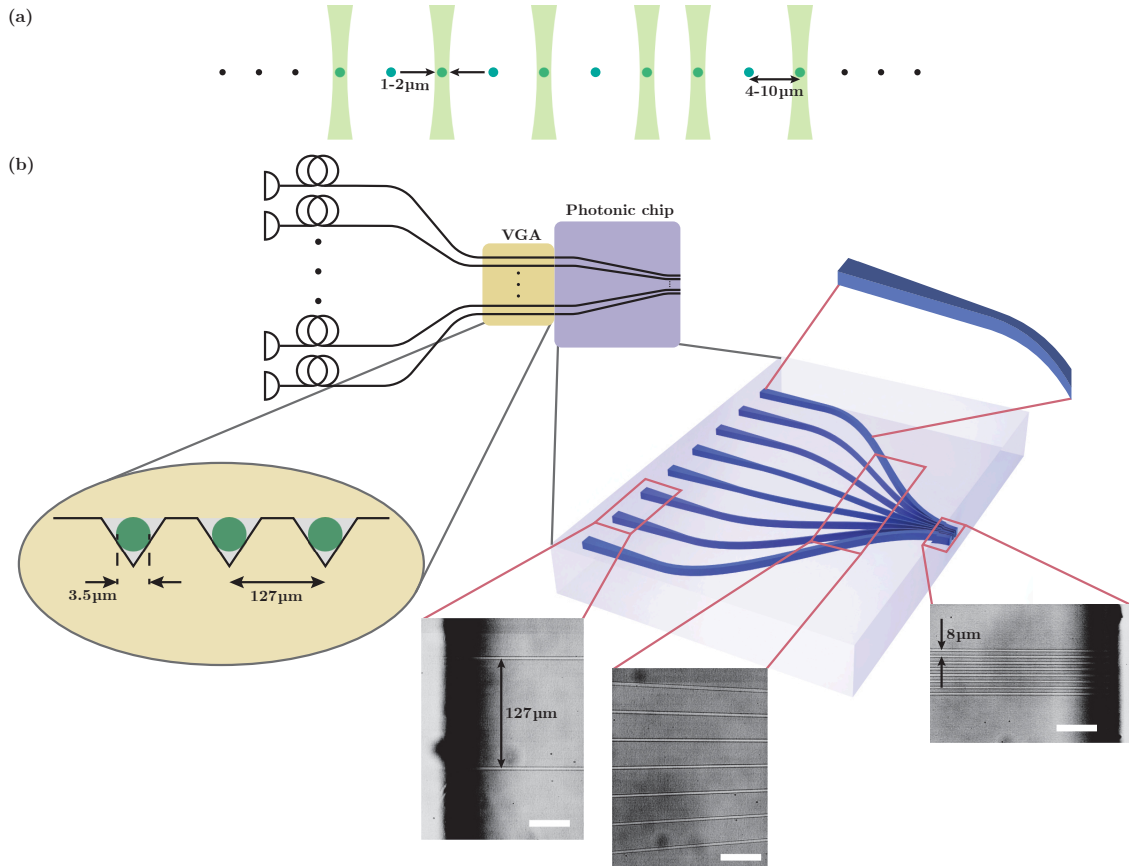


Figure 5.1: (a) The ideal output of a laser based addressing system for trapped ions. Ions are typically spaced by $\sim 5\mu\text{m}$. We address the ions with individual beams with a waist of $1-2\mu\text{m}$, much smaller than the ion-ion spacing. Each beam has its own independent phase and amplitude control allowing for parallel operation of each channel. (b) Diagram of the V-groove array (VGA) and photonic chip. The VGA and chip are aligned and glued together on a 6 axis stage to maximise coupling efficiency. The photonic chip brings the output from a $127\mu\text{m}$ spacing down to $8\mu\text{m}$ for direct imaging onto the ions.

5.2 Addressing apparatus

Here, I provide a brief overview of the addressing apparatus. For further details, see Refs [2, 47]. The 532nm addressing system can be broken down into three key parts. First, a single fibre input is split into eight channels using off the shelf fibre splitters¹ The light then passes through individual fibre-AOMs² for phase and amplitude control. We also provide each fibre channel with a polarisation controller³,

¹Thorlabs TWQ560HF & TW560R5A1

²Gooch & Housego, SFO5564-S-M200-0.4C2C-3-F2PS-01

³Thorlabs CBC900




Figure 5.2: Telecentric imaging system which maps the SPIM-WG output onto the ions. The waveguide output is a series of 8 spots with mean separations of $7.84(3) \mu\text{m}$ and waist size of $1.25(1) \mu\text{m}$. The laser light passes through the main EFL 4 mm objective we use to image the ions with the two paths separated by a dichroic mirror (DM). The 1762 nm laser enters parallel to the B-field direction, this has since changed and so is different to the geometry used in the rest of this thesis.

giving us the degree of freedom to ensure all beams output from the waveguide have the same polarisation. This is required as we do not use polarisation-maintaining fibres in the splitters or fibre AOMs.

This fibre network is coupled to a set of fibres that connect to the centre-piece of the system, the V-groove array (VGA) and photonic chip. A diagram of the V-groove array and photonic chip is shown in Fig 5.1. The VGA on its own turns the bundle of fibres into a photonic chip with an acceptable spot size but the wrong spacing between channels. When coupled to the photonic chip, we can also achieve the desired spacing for our addressing system. As such, the input of the photonic chip is designed to match the spot size and spacing of the VGA to provide maximum coupling. From there, the photonic chip brings the laser light from a spot size and spacing of $3.5 \mu\text{m}$ and $127 \mu\text{m}$ to $1.25(1) \mu\text{m}$ and $7.84(3) \mu\text{m}$ respectively. Finally, a telecentric 2:1 imaging system shown in Fig 5.2, brings the spots to a $3.95(3) \mu\text{m}$ separation and $0.67(6) \mu\text{m}$ spot size to match the ions.

5.3 System characterisation

To ensure the suitability of our addressing system to ion-trap quantum computing, we characterise the device in-situ on a single ion. By utilising the segmented electrodes of our ion trap, we can move a single ion through each of the eight incident beams to characterise their profile. To measure the beam intensity at the ion, we measure the light shift on an optical qubit. We can then calculate a normalised intensity very easily, as the light shift is proportional to the intensity⁴.

We wish to measure the light shift across the beam profile, from the centre to the next ions location. This measurement must have a large dynamic range spanning several orders of magnitude to resolve the crosstalk at the next ion's position. In practice, this means resolving from the MHz level down to the 10 Hz level of light shifts in the same measurement. This poses a similar challenge to our calibration in Section 4.3.2. Choosing a short Ramsey duration ($<1 \mu\text{s}$) will limit us to measuring only very large light shifts ($>1 \text{ MHz}$), meaning we cannot measure the low cross-talk. Choosing a long Ramsey duration will cause aliasing at the centre of the beam where the light shift is at the $\sim 1 \text{ MHz}$ level, again preventing a proper characterisation of the device. Without prior knowledge of the light shift, any single choice of Ramsey duration would not suffice to characterise the system. As such, we turn to robust phase estimation (RPE) once again; for details, see Section 4.3.2.

The coherence time of the optical qubit we used only allows Ramsey sequences with $\sim 1 \text{ ms}$ of interrogation time⁵. This only allows us to measure light shifts down to $\sim 1 \text{ kHz}$, which is insufficient for measuring the waveguide cross-talk. To overcome this limitation, when using interrogation times beyond the qubit coherence time, we embed the Ramsey sequence in a Knill-dynamical-decoupling (KDD) sequence [96], greatly extending the effective coherence time of the transition, thus allowing much

⁴It is also possible to measure an absolute intensity by comparing to the theoretical light shift on the transition, though this isn't required to characterise cross-talk, which is our primary interest.

⁵These measurements were taken before the 1762 nm laser was moved in our system to allow $\Delta m = 0$ transitions. As such, we settled on a magnetically sensitive qubit $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = -1\rangle$.




Figure 5.3: (a) The Ramsey sequence used to measure the light shift from the 532 nm beams for short durations. (b) The KDD sequence used to embed 532 nm pulses while maintaining coherence on the optical qubit for long probe durations (>1 ms). The extended coherence time allows Ramsey durations of up to 100 ms (200 ms total duration) to be implemented allowing ~ 10 Hz level measurement of the light shift from the waveguide output.

greater resolution in the measurement. This sequence is depicted in Fig. 5.3. Using RPE, we measure the intensity profile of all eight channels in our waveguide, with results shown in Fig. 5.4. Across the chip we measure cross-talk (where a second ion would be placed in practice) below the $\sim 10^{-3}$ level, sufficient for high-fidelity quantum computation. At its lowest we measure cross-talk of $\sim 10^{-5}$.

This result is very promising for our short-term use cases, but there are also options available to push this cross-talk even lower. For example, it is possible to implement composite pulse schemes to reduce the effective waist of an applied beam as seen by the ions [97], or to perform active cancellation in neighbouring channels to directly reduce cross-talk levels further [98]. Furthermore, the $\sim 10^{-3}$ cross-talk would only contribute to a $\sim 10^{-3}$ level of error if one of the two Raman beams were individually addressed. By addressing both, the error rate from cross-talk is suppressed exponentially, pushing this far below the typical error rate in most trapped ion systems.

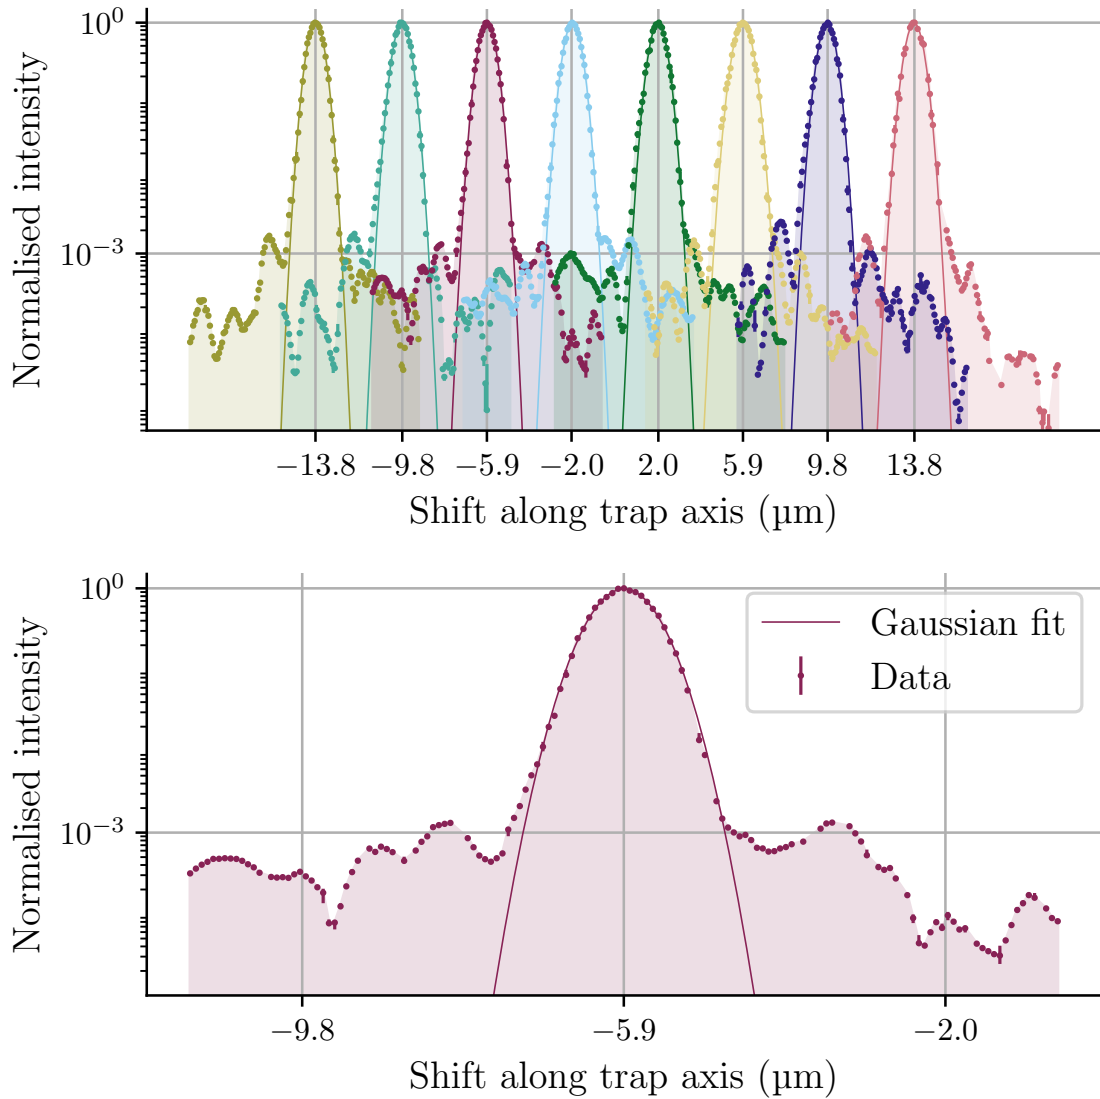


Figure 5.4: (a) Cross-talk results from scanning a single ion through the output of the SPIM waveguide array. Each channel was enabled and scanned separately. We measure cross-talk below 10^{-3} across the chip and at the 10^{-5} level in the best cases. Variations in cross-talk across the chip arise from a combination of optical aberrations and light leakage from neighbouring waveguide channels. The detection limit of the measurement is 10^{-6} , which is well below the lowest measured cross-talk level of 10^{-5} . The intensity of each channel is normalised to its own maximum value. Each x-tick is located at the centre of a waveguide channel where an ion would be placed during full operation. (b) An enhanced view of the cross-talk measurement of channel 3.

6

High-fidelity heralded state preparation and measurement

Contents

6.1	Introduction	86
6.2	Taxonomy of errors	87
6.3	The protocol	89
6.3.1	State preparation	89
6.3.2	Measurement	90
6.4	Scaling of data rates and errors	91
6.4.1	State preparation	91
6.4.2	Measurement	94
6.5	Observable Statistics	95
6.6	Implementation in $^{137}\text{Ba}^+$	98
6.6.1	Metastable level qubit	99
6.6.2	Optical qubit	100
6.6.3	Ground qubit	101
6.7	Results	103
6.7.1	Optical detection	103
6.7.2	Metastable level lifetime	104
6.7.3	Protocol results	106
6.7.4	Observable bias	107

In this chapter, we discuss the work in Ref. [3], a protocol for high-fidelity quantum state preparation and measurement. Parts of this work, such as details of

the $^{137}\text{Ba}^+$ ion and imaging system, can be found in my co-author's thesis (Ref. [47]).

6.1 Introduction

When performing a quantum experiment, both initialising the system and performing a measurement with high-fidelity is of vital importance. If these crucial steps are error prone, then the total fidelity of the experiment is limited before any action is taken on the quantum system. In this chapter, we describe a new protocol to detect errors in state preparation and measurement (SPAM), and convert them into erasures [99, 100]. By integrating erasure conversion into our SPAM methodology, there is an additional benefit over simply reducing error rates. Not only does our protocol work with systems that contain technical errors, but it also enables us to overcome errors that would otherwise require considerable effort. For example, in $^{137}\text{Ba}^+$, the error when performing optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$ is $\sim 10^{-3}$, even in low magnetic fields, and it cannot be improved by laser engineering alone.

At small scales, a trade-off can be made between error rates and runtime, allowing for flexibility. On the other hand, large scale devices can embed erasure codes into the error correcting cycle to eliminate SPAM and qubit leakage errors. Our protocol has minimal requirements, which are typically met by today's quantum devices, allowing it to be implemented with only software changes.

In atomic platforms, state preparation is usually done via optical pumping [50, 101–104], and can be combined with coherent processes to reduce errors ($\leq 10^{-4}$). This is particularly useful in trapped ions with nuclear spin $I > 1/2$ [105–107], where there are a large number of states that cannot be pumped out of using frequency selectivity.

The most common method of performing a measurement in trapped ion systems is to detect state-dependent fluorescence [37, 108–110]. By shelving one qubit state to a metastable level, such as the $D_{5/2}$ level in Ca^+ , Sr^+ , or Ba^+ , detecting fluorescence allows us to discriminate between qubit population in $S_{1/2}$ and $D_{5/2}$. Since this method only discriminates between $S_{1/2}$ and $D_{5/2}$, it can only give a measurement of $|\downarrow\rangle$ or $|\uparrow\rangle$. However, these are not the only possibilities in a trapped

ion system; population could be lost in error states outside of the qubit manifold. For example, leakage into additional states in $S_{1/2}$ would lead to a measurement of $|\downarrow\rangle$ irrespective of the qubit state, and thus an error has occurred. Such errors are referred to as leakage errors [111–113].

There has been significant recent work on *erasure conversion* [45, 99, 100, 113, 114], which aims to detect errors in quantum systems with a known location and to either discard or correct them. This contrasts with typical error correction, where the error’s location is inferred through what is effectively a majority vote. Our protocol uses a series of diagnostic measurements during the state preparation and measurement processes to detect leakage errors and convert them into erasures. Such erasure conversion can be done using post-selection [115], or repeat-until-success (RUS) [116]. The protocol is lightweight, as the diagnostic measurements add little overhead while exposing and correcting errors that would have otherwise gone undetected.

In the following chapter, we demonstrate the use of the protocol for three different qubit encodings in a single trapped $^{137}\text{Ba}^+$ ion. For all three, we achieve the lowest reported SPAM infidelities of $7(4) \times 10^{-6}$ (optical qubit), $5(4) \times 10^{-6}$ (metastable-level qubit), and $8(4) \times 10^{-6}$ (ground-level qubit). In general, the protocol can be applied to any quantum system which meets the basic requirements outlined in Sec. 6.3.

6.2 Taxonomy of errors

Our SPAM protocol accounts for errors where the qubit population is outside of the qubit manifold. These are typically known as leakage errors, but for the purpose of understanding the protocol we can break these down further. As we go through the protocol in section 6.3, we will highlight how each error can be eliminated.

The different leakage errors that can occur in our system during SPAM are as follows:

1. Optical pumping error. Due to polarisation impurity, population may be left in the ground level, but not in the target state after optical pumping [50, 101–105].
2. Population transfer error (aka mapping error). When mapping between manifolds for quantum non-demolition (QND) measurements or initialising qubits via optical pumping [1], it is possible that population is left behind. If left in the ground level when mapping to the metastable level, this would constitute a shelving error, its inverse a deshelling error, and within a manifold, a general mapping error. In many systems this error rate is the same order of magnitude as the infidelity of single-qubit operations on the optical qubit via which mapping occurs, $\sim 10^{-3}$ [117, 118].
3. Scattering [119, 120] & decay [121] error. Throughout the computational and measurement process, it is possible that population leaves the computational manifold through an off-resonant laser scattering event e.g. during Raman operations or decay from the metastable level to the ground level.
4. Ion loss error. The ion may be lost entirely due to background gas collisions in the trap.

Aside from leakage error, measurement error is also present in our protocol. Due to noise in our optical detection system, it is possible we infer the wrong state when performing a measurement. We cannot fully eliminate this, and it sets the bound for our protocol’s capabilities. Measurement errors can also arise from decay during the measurement process. If a decay occurs during a measurement that would have read *dark* (from which we infer $|\uparrow\rangle$), then the outcome will be $|\downarrow\rangle$ even though the correct result was $|\uparrow\rangle$. Since such events are undetectable within the protocol, they are treated separately from typical decay errors.

6.3 The protocol

Our protocol has three requirements for successful implementation. The following description is specific to trapped ions, but can be applied to any quantum system which meets the requirements.

1. A subspace for quantum non-demolition (QND) measurements which does not couple to the $S_{1/2} \rightarrow P_{1/2}$ fluorescence cycle. In Ca^+ , Sr^+ , and Ba^+ systems this is the metastable $D_{5/2}$ level (B in Fig. 6.1(a)).
2. The ability to map between the $S_{1/2}$ level (A in Fig. 6.1) and B . These are the *population transfer pulses* in Fig. 6.1.
3. The ability to randomise population in A and flush population from B . This is required to enable the reset of qubits irrespective of the ion's state. Cooling and optical pumping require an ion in the $S_{1/2}$ (A) manifold, hence we must be able to clear population from $D_{5/2}$ (B). Ideally this randomisation also allows us to prepare into some initial state $|SP\rangle_A$ with high probability ($> 90\%$) to increase data rates. In atomic platforms, this is done through optical pumping.

With this, our protocol can utilise high-fidelity fluorescence detection to determine the qubit state, or if population is outside of the qubit manifold. Fluorescence detection allows us to distinguish if population is in A or B . If population is in A , the ion will scatter photons via the $S_{1/2}$ to $P_{1/2}$ cycle and we call this a *bright* result. Conversely, projection into B does not scatter photons and so is *dark*. The full SPAM sequence is shown in Fig. 6.1(b). We split the protocol into a *state preparation* part and a *measurement* part, in between which a quantum computation can be implemented.

6.3.1 State preparation

State preparation begins by flushing population out of B into A , and optical pumping into a given state within A (or population randomisation within A). In trapped ions this is done by optical pumping, where light of a pure polarisation is used to prepare

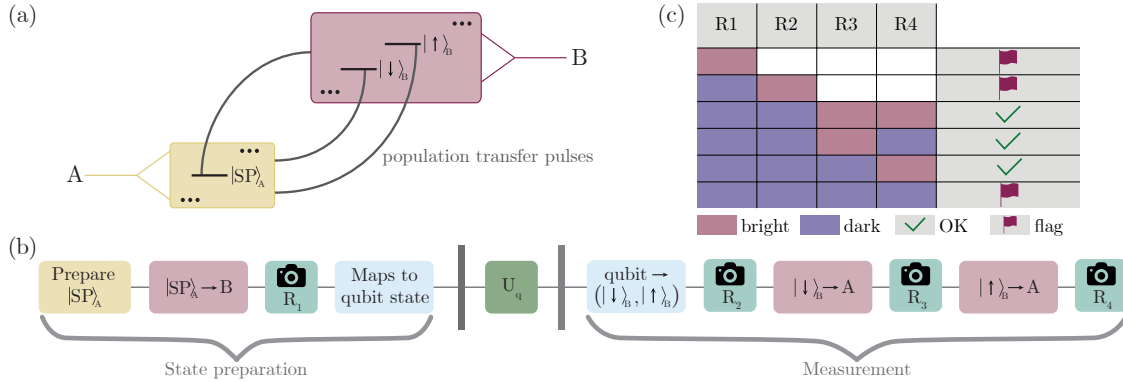


Figure 6.1: SPAM protocol. (a) Basic requirements of the quantum system performing the protocol. Two manifolds A and B must exist with B capable of protecting qubit population in QND measurements [1, 122, 123]. (b) The full SPAM sequence. Camera icons represent population detection steps. The light blue steps are only required if the qubit encoding is not $|\downarrow / \uparrow\rangle_B$. (c) Truth table for possible outcomes in the SPAM protocol.

into $|\text{SP}\rangle_A$, typically with $\sim 10^{-3}$ error [50]. Next, a transfer pulse carries population from $|\text{SP}\rangle_A$ to B . After this mapping step, we perform a population detection step (R_1) to check if any population has been left in A which would lead to a *bright* measurement. If a *bright* result is found, this indicates either an optical pumping error, or a mapping (shelving) error. The *bright* measurement triggers an error flag and the process can be repeated, or this shot can be rejected in post-selection.

In the absence of any error flag, transfer pulses can move population to a state in the qubit manifold if this is not already the case.

6.3.2 Measurement

The measurement process begins by coherently mapping qubit population to B , if the qubit is not already fully encoded in B . Following this is a population detection step (R_2) to check for decay errors, mapping errors, or leakage errors from scattering [119, 120]. If an error is detected, a flag is raised. Next, each qubit state is mapped from B to A , followed immediately by population detection each time (R_3 and R_4). The step R_3 acts as the discriminator measurement, determining which state the qubit is in. A *bright* outcome corresponds to $|\downarrow\rangle$, and a *dark* outcome corresponds to $|\uparrow\rangle$. In the event that R_3 reads *dark*, R_4 checks that population can be found in the $|\uparrow\rangle$

state. If all measurements are *dark*, an error has occurred, leaving population in B either outside of the qubit manifold or imperfectly mapped to A , and a flag is raised. The truth table for all of these population detection steps is given in Fig. 6.1(c).

Before and after the core protocol, two additional measurements can be added R_0 and R_5 . These measurements are made with population flushed from B into A . They should always be *bright* and serve to check for ion loss.

6.4 Scaling of data rates and errors

Throughout our protocol, errors are flagged via diagnostic measurements. It is important to consider how these flags affect an experiment's data rates as the system size scales, enabling a balanced trade-off between the engineering effort required to reduce error rates and the data rate impact of our protocol. The general result for our protocol is that if data rates for a given system are prohibitive to run a quantum algorithm, then the corresponding error rates would also prohibit achieving a meaningful result. This means that to scale a system effectively, one must start from a place with low error rates; our protocol can then be used to mitigate errors below any engineering limits encountered.

In this section, we explore the optimal number of state preparation cycles to perform, and how data rates in state preparation and measurement are each impacted by scaling the system size. The scaling we observe should be viewed as converting a scaling in error rates to a scaling data rates in the worst case scenario. In the state preparation case, the ability to individually reset each qubit enables better scaling of data rates than the native exponential scaling of error rates a system would otherwise exhibit.

6.4.1 State preparation

We start by considering the state preparation stage of the protocol. To prepare a register of N qubits, the optimal strategy is to only re-prepare qubits which have raised an error flag in the state preparation stage. In our protocol, we assume that

the probability of failing to prepare a qubit is much higher than the optical detection error, since the latter can be easily reduced with high quality imaging and long detection durations. With finite preparation error and a large qubit register, needing to re-prepare qubits which were not prepared successfully is inevitable. During a cycle of state preparation, there is the possibility that a qubit which was successfully prepared decays from B to A . We label the probability of this decay q , and the probability of successfully preparing a qubit p . The change in population for a state preparation cycle is driven by the following Markov matrix

$$M = \begin{pmatrix} 1-p & q \\ p & 1-q \end{pmatrix}. \quad (6.1)$$

The protocol starts with all population in A , represented by the vector $s = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$. After a single state preparation cycle, the system will be in the state

$$s = \begin{pmatrix} 1-p \\ p \end{pmatrix}, \quad (6.2)$$

where the top/bottom row of the vector represents unprepared/prepared population respectively. This leaves the fraction $1-p$ of the qubits unprepared, and so we repeat the process. After n cycles of state preparation, the Markov matrix is

$$M^n = \frac{1}{p+q} \begin{pmatrix} q+p(1-p-q)^n & q-q(1-p-q)^n \\ p-p(1-p-q)^n & p+q(1-p-q)^n \end{pmatrix}, \quad (6.3)$$

which puts the system in the state

$$s = \frac{1}{p+q} \begin{pmatrix} q+p(1-p-q)^n \\ p-p(1-p-q)^n \end{pmatrix}. \quad (6.4)$$

We see that there is a balance between the preparation of previously unprepared qubits, and the decay of qubits which have already been successfully prepared. With an infinite number of attempts, $s_\infty = \frac{1}{p+q} \begin{pmatrix} q \\ p \end{pmatrix}$. This sets the error for an infinitely large qubit register to be $\epsilon = \frac{q}{p+q}$. For a register of finite size N , the expectation is that the populations S will reach $S = \frac{N}{p+q} \begin{pmatrix} q \\ p \end{pmatrix}$.

In our system, the probability of success $p \sim 0.9$ and a single cycle of state preparation takes around $400 \mu\text{s}$ (see section 6.6 for details). Using the ~ 30 s lifetime

of the $D_{5/2}$ state of Ba^+ , we calculate $q = 1 - e^{-0.0004/30} = 1.33 \times 10^{-5}$. Using these parameters in Eq. 6.1 we simulate the error rate by repeatedly applying M to the initial state $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and plot the result in Fig. 6.2. After six steps the system reaches a steady state, where the rate of qubits being re-prepared is balanced by the decay of previously prepared qubits. The protocol should be truncated at or shortly after this stage, as further state preparation cycles are not expected to improve error rates.




Figure 6.2: Simulation of the Markov matrix in Eq. 6.1 with $p = 0.9$, $q = 1.33 \times 10^{-5}$ being applied to $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$. The blue dotted line represents the rate of improving preparation in the absence of decay. The yellow dashed line shows the asymptotic value, which is reached after approximately six iterations.

Even though a steady state is reached with error ϵ , it is still possible to achieve error rates below ϵ by terminating the state preparation stage our protocol immediately after a successful population detection step. After some number of iterations, we assume the system to be in the state $S = \begin{pmatrix} N_0 \\ N_1 \end{pmatrix}$, where $N_0 = \lceil \frac{Nq}{p+q} \rceil$, and $N_1 = \lfloor \frac{Np}{p+q} \rfloor$. Each cycle of state preparation is probabilistic, so there is a non-zero probability that the next cycle achieves a fractional error less than ϵ . For this to occur, there must be N_s qubits successfully prepared out of the N_0 available and N_d decay events such that,

$$N_\delta = N_0 - N\epsilon' \leq N_s - N_d, \quad (6.5)$$




Figure 6.3: Simulation of Eq. 6.6 using $\epsilon = 5 \times 10^{-6}$ and $p = 0.9$ with three different values of q : 10^{-4} , 10^{-5} , 10^{-6} . With $N < 10,000$ it is feasible for platforms to perform state preparation observing zero errors. The probability of such an event is > 0.1 in all three cases.

where N_δ is the number of qubits by which the protocol must randomly outperform the steady state. The probability of Eq. 6.5 is given by a joint binomial distribution,

$$\mathbf{P} = \sum_{i=\lceil N_\delta \rceil}^{N_0} \left[\binom{N_0}{i} p^i (1-p)^{N_0-i} \sum_{j=0}^{\lfloor i-N_\delta \rfloor} \left(\binom{N_1}{j} q^j (1-q)^{N_1-j} \right) \right]. \quad (6.6)$$

Equation 6.6 highlights how we achieve such a low SPAM error with a single qubit in Sec. 6.7. In our case $N = 1$, so no sum is required, $i = 1$, $j = 0$, $N_0 = 1$, and $N_1 = 0$. As such, we recover that the probability of successfully preparing the qubit is p . With only one qubit the fidelity of this is to within measurement error, which is less than ϵ .

Using $p = 0.9$, we compute the probability of preparing to 5×10^{-6} error as the number of qubits increases. The result is shown in Fig. 6.3 for three different values of q . For present day devices, where $N < 10,000$, and with small q ($< 10^{-4}$), it is still possible to prepare the register to within measurement error. This is because for a small number of qubits, $N_0 = 1$ and $N_1 = N - 1$, meaning only one successful preparation is required with zero decays.

6.4.2 Measurement

Scaling of the measurement aspect of the protocol is simpler than the state preparation stage. Before a measurement can be repeated, the entire system must be re-prepared

and the quantum algorithm must be re-run. As a result, there is a large overhead in data rates to performing repeat-until-success. If the probability of successfully measuring a qubit state without the protocol is $p_{\mathbf{M}}$, then the probability of a qubit register of size N being measured with a fractional error less than $p_{\mathbf{M}}$ scales as $\sim p_{\mathbf{M}}^N$. As with state preparation, it is important to note that prohibitive data rates when using our protocol are a strong indication of prohibitive error rates when not.

6.5 Observable Statistics

In our protocol, the rate at which data is discarded can be dependent upon the qubit state. Since each qubit state requires different mapping pulses which have different error rates, this state dependence is inevitable. In this section, we will investigate the consequences of state dependent post-selection, and see how it impacts the statistics of the observable \mathcal{Z} . Such an impact is particularly relevant for quantum computing applications which are dependent upon measuring observable averages [124–126]. Later in section 6.7.4, we perform measurements to verify the statistical framework we form in this section.

We adopt the notation R_3 is *bright/dark* as b/d , and a shot (that has already passed the state preparation stage) is accepted/rejected as a/r . For brevity we adopt the notation $P(|i\rangle) \rightarrow P(i)$.

Starting with the sum of probabilities,

$$P(b, a) = \sum_i P(b, a, i), \quad (6.7)$$

we can express conditional probabilities,

$$P(b, a) = \sum_i P(b, a | i)P(i). \quad (6.8)$$

For a qubit, $|i\rangle \in \{|\downarrow\rangle, |\uparrow\rangle\}$, we use $P(\downarrow) + P(\uparrow) = 1$, substitute into eq. 6.8 and rearrange to find

$$\begin{aligned}
P(\downarrow) &= \frac{P(b, a) - P(b, a | \uparrow)}{P(b, a | \downarrow) - P(b, a | \uparrow)}, \\
P(\uparrow) &= \frac{P(b, a) - P(b, a | \downarrow)}{P(b, a | \uparrow) - P(b, a | \downarrow)}, \\
\langle \mathcal{Z}_{\text{real}} \rangle &= \frac{2P(b, a) - P(b, a | \uparrow) - P(b, a | \downarrow)}{P(b, a | \downarrow) - P(b, a | \uparrow)}.
\end{aligned} \tag{6.9}$$

The property $P(b, a | i)$ can be characterised for each qubit state $|i\rangle$ allowing corrections to be made to the measured observable. Using conditional probability, this property can be broken down into two components, $P(b, a | i) = P(b | a, i)P(a | i)$. This highlights the two error sources that may bias the observable measured.

$P(b | a, i)$ characterises the single shot accuracy of the implementation of the protocol. The ideal and theoretical error rates are described in Tab. 6.1. For our

Outcome	Ideal	Theory
$P(b a, \downarrow)$	1	$1 - \epsilon_M$
$P(d a, \downarrow)$	0	ϵ_M
$P(b a, \uparrow)$	0	$e^{-t_d/\tau}$
$P(d a, \uparrow)$	1	$1 - e^{-t_d/\tau}$

Table 6.1: Ideal and theoretical error rates in the measurement process of the SPAM protocol. ϵ_M is the measurement or optical detection error. The readout duration is t_d and τ is the natural lifetime of B . In Ba^+ this is 30.14 s [121].

system the qubit state $|\downarrow\rangle$ is measured first and so is unaffected by decay error, however measurement error is still possible. This is negligible compared to decay error when characterising our SPAM fidelity and so it is not included in the calculation. The dominant source of error for our protocol is decay error which only impacts qubit state $|\uparrow\rangle$ as it is measured second. This means there is a chance that $|\uparrow\rangle$ decays during the measurement of $|\downarrow\rangle$ leading to a false *bright*.

On the other hand, $P(a | i)$ characterises the probability of a shot measuring state $|i\rangle$ to be accepted. In our system, this can be at the $\sim 1\%$ level due to imperfect

mapping between A and B - details can be found in ref. [47]. In comparison, the errors from decay are negligible when considering the \mathcal{Z} observable. As such, we only consider mapping errors and obtain,

$$\begin{aligned} P(\downarrow) &= \frac{P(b, a)}{P(a | \downarrow)} \\ &= \frac{P(b | a)P(a)}{P(a | \downarrow)}, \\ P(\uparrow) &= \frac{P(d | a)P(a)}{P(a | \uparrow)}, \end{aligned} \tag{6.10}$$

where we have used $P(b | a) + P(d | a) = 1$ and substituted. In the event of perfect mapping $P(\downarrow / \uparrow) = P(b/d | a)$.

With imperfect but uncharacterised population transfer pulse fidelities, the observable $\langle \mathcal{Z}_{\text{meas}} \rangle$ is given by

$$\begin{aligned} \langle \mathcal{Z}_{\text{meas}} \rangle &= P(b | a) - P(d | a) \\ &= \frac{\gamma P(\downarrow) - P(\uparrow)}{\gamma P(\downarrow) + P(\uparrow)} \end{aligned} \tag{6.11}$$

where $\gamma = \frac{P(a|\downarrow)}{P(a|\uparrow)}$. If $P(a | \downarrow) = P(a | \uparrow)$, we recover $\langle \mathcal{Z}_{\text{real}} \rangle$. In Fig. 6.4

For each qubit type, the value of $P(a | i)$ differs. If the qubit state is in the metastable level, then only a single mapping error causes data to be discarded and so $P(a | i) = \epsilon_{\text{Map}}$. For qubit states in the ground level, the error applies both on shelving and deshelving, giving $P(a | i) = \epsilon_{\text{Map}}^2$.

Even with this bias in $\langle \mathcal{Z}_{\text{meas}} \rangle$, it is possible to recover $\langle \mathcal{Z}_{\text{real}} \rangle$ with several methods.

1. If the mapping errors are well known, one could adjust each mapping pulse to give the same error. However, this will impact data rates and is subject to drifts in the experimental parameters.
2. Again with well characterised maps, we can apply a correction to the statistics.

If γ is well known, then

$$\langle \mathcal{Z}_{\text{real}} \rangle = \frac{\langle \mathcal{Z}_{\text{meas}} \rangle (\gamma + 1) - (\gamma - 1)}{(\gamma + 1) - \langle \mathcal{Z}_{\text{meas}} \rangle (\gamma - 1)}.$$




Figure 6.4: A plot of $\langle \mathcal{Z}_{\text{meas}} \rangle - \langle \mathcal{Z}_{\text{real}} \rangle$ with $\gamma \in [\frac{1}{2}, 2]$ and $P(|\downarrow\rangle) \in [0, 1]$. $\langle \mathcal{Z}_{\text{meas}} \rangle$ is using the form in Eq. 6.11. The bias tends to zero in the cases where $\gamma \rightarrow 1$, $P(|\downarrow\rangle) \rightarrow 0$ or $P(|\downarrow\rangle) \rightarrow 1$.

3. We can also alternate the definition of $|\downarrow\rangle$ and $|\uparrow\rangle$. This ensures that error contributions from imperfect mapping pulses affect both qubit states equally, removing the bias that would otherwise arise from rejecting more shots in one state than the other. However, this does not address the bias introduced by the decay channel, which cannot be corrected and still impacts our single-shot accuracy.

Of these three methods, number 3 is the most convenient and reliable choice. It does not require accurate or frequency calibration to work and is robust to system drifts, for example changes in the fidelity of the mapping pulses which 1 and 2 are susceptible to.

6.6 Implementation in $^{137}\text{Ba}^+$

We benchmark the protocol on three qubit types in $^{137}\text{Ba}^+$, optical (O), metastable (M) and ground (G) [1]. Fluorescence from the $^{137}\text{Ba}^+$ ion is detected using a scientific

CMOS (sCMOS) camera¹. For details of the $^{137}\text{Ba}^+$ ion, imaging system, and sCMOS camera region of interest see ref. [47].

The following encodings are used O: $|D_{5/2}, F = 2, m_F = -1\rangle \leftrightarrow |S_{1/2}, F = 2, m_F = 0\rangle$, M: $|D_{5/2}, F = 2, m_F = -1\rangle \leftrightarrow |D_{5/2}, F = 1, m_F = -1\rangle$, and G: $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |S_{1/2}, F = 1, m_F = 0\rangle$. We identify the manifolds A and B with the levels $S_{1/2}$ and $D_{5/2}$ respectively and the state $|\text{SP}\rangle_A$ is $|S_{1/2}, F = 2, m_F = 0\rangle$.

Each experimental sequence begins with Doppler cooling of the trapped $^{137}\text{Ba}^+$ ion [127–129]. For details of the system parameters see ref. [47] or ref. [2]. This is followed by optical pumping into $|\text{SP}\rangle_A$, using three lasers. The 614 nm beam flushes population from $D_{5/2}$ to $S_{1/2}$ via $P_{3/2}$. Additionally, π -polarised 493 nm light drives the $S_{1/2}$ to $P_{1/2}$ cycling transition. The 493 nm laser is resonant with $|S_{1/2}, F = 2\rangle \rightarrow |P_{1/2}, F = 2\rangle$ and $|S_{1/2}, F = 1\rangle \rightarrow |P_{1/2}, F = 2\rangle$. Under π -polarisation $\Delta m = 0$ transitions are driven, causing population to be pumped into $|S_{1/2}, F = 2, m_F = 0\rangle$ due to the forbidden $|S_{1/2}, F = 2, m_F = 0\rangle \rightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ transition. The final laser in optical pumping is 650 nm light to drive population that has decayed from $P_{1/2}$ to $D_{3/2}$ back to $S_{1/2}$.

All three qubit states begin with this optical pumping followed by a 1762 nm mapping pulse driving population into $D_{5/2}$. However, to track population effectively, the exact states used differ.

6.6.1 Metastable level qubit

We begin by describing the process used for the M qubit. This is best suited to the protocol as the qubit manifold is within B meaning extra mapping pulses after state preparation and before measurement are not required.

1. State preparation
 - (a) Flush population from $D_{5/2}$ using the 614 nm laser
 - (b) Population detection step R_0 .

¹Hamamatsu Orca Fusion BT

- (c) Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (d) 1762 nm transfer pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to either $|D_{5/2}, F = 2, m_F = -1\rangle$ ($|\downarrow\rangle$) or $|D_{5/2}, F = 1, m_F = -1\rangle$ ($|\uparrow\rangle$).
- (e) Population detection step R_1 .

2. Measurement

- (a) Population detection step R_2 .
- (b) 1762 nm transfer pulse from $|D_{5/2}, F = 2, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (c) Population detection step R_3 .
- (d) 1762 nm transfer pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (e) Population detection step R_4 .
- (f) Flush population from $D_{5/2}$ using the 614 nm laser
- (g) Population detection step R_5 .

6.6.2 Optical qubit

The optical qubit requires some extra mapping pulses to transfer between $S_{1/2}$ and $D_{5/2}$ states after R_1 and before R_2 . This is because one qubit state is in the ground level, which is susceptible to the population detection steps. The procedure runs as follows.

1. State preparation

- (a) Flush population from $D_{5/2}$ using the 614 nm laser
- (b) Population detection step R_0 .
- (c) Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$.

- (d) 1762 nm transfer pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to either $|D_{5/2}, F = 2, m_F = -1\rangle$ ($|\downarrow\rangle$) or $|D_{5/2}, F = 1, m_F = -1\rangle$ ($|\uparrow\rangle$).
- (e) Population detection step R_1 .
- (f) If preparing $|S_{1/2}, F = 2, m_F = 0\rangle$ ($|\uparrow\rangle$), do a 1762 nm transfer pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.

2. Measurement

- (a) 1762 nm transfer pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{5/2}, F = 1, m_F = -1\rangle$ ².
- (b) Population detection step R_2 .
- (c) 1762 nm transfer pulse from $|D_{5/2}, F = 2, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (d) Population detection step R_3 .
- (e) 1762 nm transfer pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ ³ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (f) Population detection step R_4 .
- (g) Flush population from $D_{5/2}$ using the 614 nm laser.
- (h) Population detection step R_5 .

6.6.3 Ground qubit

As with the optical qubit, the ground qubit requires extra mapping pulses after R_1 and before R_2 between $S_{1/2}$ and $D_{5/2}$. In this qubit, we also utilise different $D_{5/2}$ states for state preparation and measurement to avoid second order mapping errors from the additional mapping pulses.

²Mapping to $|D_{5/2}, F = 1, m_F = -1\rangle$ could introduce second order mapping errors. If population is left in $|D_{5/2}, F = 1, m_F = -1\rangle$ after state preparation due to a mapping error, and then is not mapped down to $S_{1/2}$ before R_2 we will observe the square of the mapping error. To work around this, a different state in $D_{5/2}$ can be used for R_1 and R_2 with the ground qubit state. We do this in the ground qubit implementation for demonstration purposes. In this case, without an algorithm running, there is no difference made to the SPAM result.

³See footnote 2.

1. State preparation

- (a) Flush population from $D_{5/2}$ using the 614 nm laser
- (b) Population detection step R_0 .
- (c) Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (d) 1762 nm transfer pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{5/2}, F = 2, m_F = -1\rangle$.
- (e) Population detection step R_1 .
- (f) 1762 nm transfer pulse from $|D_{5/2}, F = 2, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$ ($|\downarrow\rangle$) or $|S_{1/2}, F = 1, m_F = 0\rangle$ ($|\uparrow\rangle$).

2. Measurement

- (a) 1762 nm transfer pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{5/2}, F = 2, m_F = 1\rangle$.
- (b) 1762 nm transfer pulse from $|S_{1/2}, F = 1, m_F = 0\rangle$ to $|D_{5/2}, F = 1, m_F = -1\rangle$.
- (c) Population detection step R_2 .
- (d) 1762 nm transfer pulse from $|D_{5/2}, F = 2, m_F = 1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (e) Population detection step R_3 .
- (f) 1762 nm transfer pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
- (g) Population detection step R_4 .
- (h) Flush population from $D_{5/2}$ using the 614 nm laser
- (i) Population detection step R_5 .

As previously mentioned, the metastable state used in the state preparation stage ($|D_{5/2}, F = 2, m_F = -1\rangle$) is different to those used in the measurement stage ($|D_{5/2}, F = 2, m_F = 1\rangle$ and $|D_{5/2}, F = 1, m_F = -1\rangle$). This is to avoid population remaining in the $D_{5/2}$ state after state preparation, but being categorised as within the qubit manifold upon performing the measurement process. Consider a 1% mapping error in the transfer $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |D_{5/2}, F = 2, m_F = -1\rangle$. After successful state preparation, we expect 1% of the qubit population to remain in $|D_{5/2}, F = 2, m_F = -1\rangle$. After an algorithm, mapping up from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{5/2}, F = 2, m_F = -1\rangle$ would guarantee $(1\%)^2 = 10^{-4}$ of the qubit population to be in $|D_{5/2}, F = 2, m_F = -1\rangle$ and having passed readout steps R_0 , R_1 , and R_2 . This is a 10^{-4} error which can be avoided by using different metastable states across state preparation and measurement stages to trigger a flag on R_4 as population is deemed lost but in $D_{5/2}$.

6.7 Results

6.7.1 Optical detection

By measuring the optical detection error independently of the metastable decay error, we can characterise our error sources and ensure that we are pushing our system to achieve the lowest possible SPAM error within the protocol. There are two forms of optical detection error we must measure, false *bright*, and false *dark*. To distinguish each, we perform two different measurements. In the ideal case, the probability of measuring *bright* when population is in the $S_{1/2}$ manifold is $P(b | S_{1/2}) = 1$, where we use b/d to represent *bright/dark* as in section 6.5. Similarly, $P(d | D_{5/2}) = 1$ in the ideal case.

To measure the optical detection error, we prepare our $^{137}\text{Ba}^+$ ion into either $S_{1/2}$ or $D_{5/2}$ and perform three consecutive population detection steps. To prepare in the ground state we flush population from $D_{3/2}/D_{5/2}$ using 650 nm/614 nm light respectively. Preparing in the $D_{5/2}$ state requires optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$

and a 1762 nm transfer pulse. If the first and last measurements are in agreement, then we can conclude that population has stayed within the chosen manifold to within our optical detection capabilities. Therefore, any inconsistency in the middle measurement must be the result of an optical detection error. A combination of two errors is possible, but is several orders of magnitude less likely than a single optical detection error.

We perform measurements for both the *bright* and *dark* states and construct the histogram shown in Fig. 6.9. We measure a false *dark* error ($P(d | S_{1/2})$) of $3.2(2.8) \times 10^{-6}$ and a false *bright* error ($P(b | D_{5/2})$) of $0.0(1.1) \times 10^{-6}$.

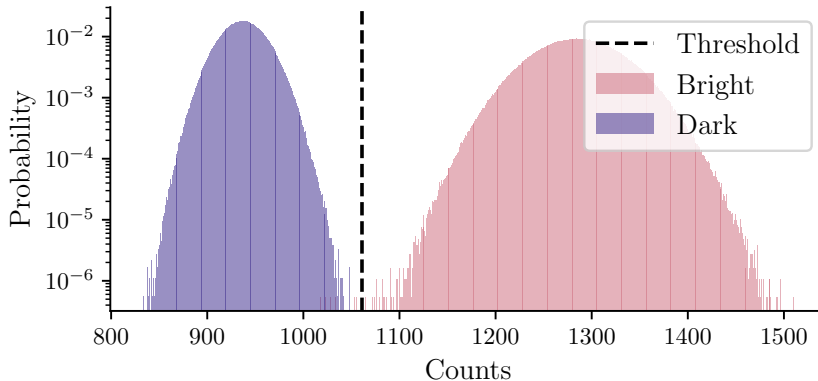


Figure 6.5: Histogram of the counts collected on camera for the second readout in three sequential population detection steps. The data is post-selected such that the first and last readout show the same result. The blue indicates the ion was prepared *dark* in $D_{5/2}$ and red the ion was prepared *bright* in $S_{1/2}$. We measure a false *dark* error ($P(d | S_{1/2})$) of $3.2(2.8) \times 10^{-6}$ and a false *bright* error ($P(b | D_{5/2})$) of $0.0(1.1) \times 10^{-6}$. This measurement consists of 1.8×10^6 experimental runs.

Both error rates are lower than the contribution from decay and so are not leading contributors to the error rate in our protocol.

6.7.2 Metastable level lifetime

The fidelity of our protocols implementation is limited by the decay error between readout steps R_2 and R_3 . This could cause a qubit which should read *dark* on R_3 to read *bright* instead. This decay will only miss the flag from R_2 if it occurs within a certain window before R_2 ends and R_3 begins. This error is given by $1 - e^{-t/\tau}$

where t is the duration to map a qubit state from $D_{5/2}$ to $S_{1/2}$ and readout, and τ is the $D_{5/2}$ decay time. We can reduce the error by either increasing the lifetime, which has a fixed upper limit of 30.14 s [121] or reducing the readout duration and mapping duration. The latter may come at the cost of increased optical detection error in the absence of improved optics and imaging.

To measure the metastable level lifetime, we begin by preparing into $|S_{1/2}, F = 2, m_F = 0\rangle$ with optical pumping. A 1762 nm mapping pulse then transfers population to $|D_{5/2}, F = 2, m_F = -1\rangle$. Next a series of population detection steps are performed with 20 ms exposure time. These detection steps are repeated until a *bright* measurement is recorded, at which point the duration for which the ion was in the $D_{5/2}$ state is recorded.

The results of this procedure are shown in Fig. 6.6. We measure a lifetime $\tau = 27.2(2)$ s, approximately 10% lower than the natural lifetime. This discrepancy can be attributed to leakage of the 614 nm beam onto the ion. Likely causes are RF leakage driving the double pass AOM which acts as a switch for this laser, and back-reflections from optics in the beam path. Some optimisations have already been made on the latter increasing lifetime from ~ 100 ms to 27.2(2) s which are detailed in Ref. [47].




Figure 6.6: Probability of decay of the metastable level as a function of the delay after shelving from $S_{1/2}$ to $D_{5/2}$. The solid line is a fit to the exponential decay $1 - e^{-t_d/\tau}$ and we measure a decay time of $\tau = 27.2(2)$ s.

6.7.3 Protocol results

In Fig. 6.7 we plot the post-selected histogram for the metastable level qubit. Over the 1,000,000 shots taken per state, 10 false *bright* measurements were detected and not removed by post-selection. There were and 0 false *dark* measurements which were not post-selected. On the $|D_{5/2}, F = 2, m_F = -1\rangle$ state 3.42% of the shots were rejected and we measure an error of $0.0(2.0) \times 10^{-6}$. For the $|D_{5/2}, F = 1, m_F = -1\rangle$ state 8.75% of the shots were rejected and we measure an error of $1.1(0.7) \times 10^{-5}$. The average SPAM error for the metastable level qubit is $5(4) \times 10^{-6}$. There is a lower error on the $|\downarrow\rangle$ state ($|D_{5/2}, F = 2, m_F = -1\rangle$) than the $|\uparrow\rangle$ state ($|D_{5/2}, F = 1, m_F = -1\rangle$) as the $|\downarrow\rangle$ state is measured first. This means any decay events between R_2 and R_3 do not contribute to errors, whereas decays of $|\uparrow\rangle$ during this time do. This asymmetry in error rate between the two qubit states may be useful in applications such as quantum error correction (QEC) [99, 130, 131] where the measurement of one qubit state is significantly more likely than the other.

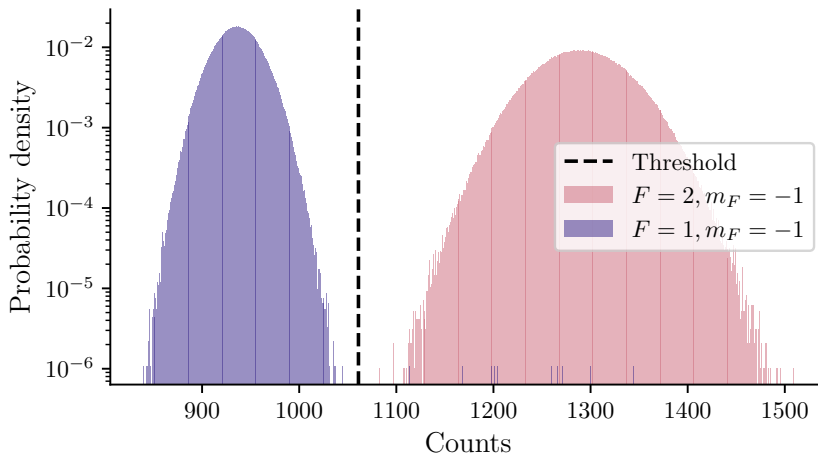


Figure 6.7: Histogram of counts from the camera when performing the SPAM protocol on the M qubit. The data shown is the remaining data after 1,000,000 shots are post-selected on $R_0, R_1, R_2, R_4,$ and R_5 for each qubit state. We observe 10 false *bright* measurements which were not removed by post-selection. The dashed vertical line shows the threshold for classifying results as *bright* or *dark*.

In Fig. 6.8 we plot the SPAM error (a) and success probability (b) after each post-selection step. We see that the change in SPAM error differs for each qubit state,

which is due to differences in mapping errors. Details of these errors can be found in ref. [47]. The $|D_{5/2}, F = 2, m_F = -1\rangle$ drops several orders of magnitude after R_3 and R_4 are factored in as this removes mapping errors for this state. The success probability in Fig. 6.8 takes the two largest drops at R_1 and the R_3 and R_4 stage as these remove mapping errors, which are our largest sources of error. The results for the optical and ground qubits are summarised in Fig. 6.9 and Fig. 6.10 respectively.




Figure 6.8: (a) The SPAM error after each post-selection step for the M qubit. The final error for $|D_{5/2}, F = 2, m_F = -1\rangle$ is $0(2) \times 10^{-6}$ and for $|D_{5/2}, F = 1, m_F = -1\rangle$ is $1.1(0.7) \times 10^{-5}$. (b) The probability of a shot being accepted after each post-selection step for the M qubit. The two qubit states have different probabilities of success due to differences in their mapping errors. With $\sim 10\%$ of the shots being rejected, we are able to reduce the SPAM error by four orders of magnitude.

6.7.4 Observable bias

We experimentally measure the observable bias as follows. We start by preparing into $|\downarrow\rangle$ or $|\uparrow\rangle$ as outlined in Fig. 6.1. For each shot we alternate the state we

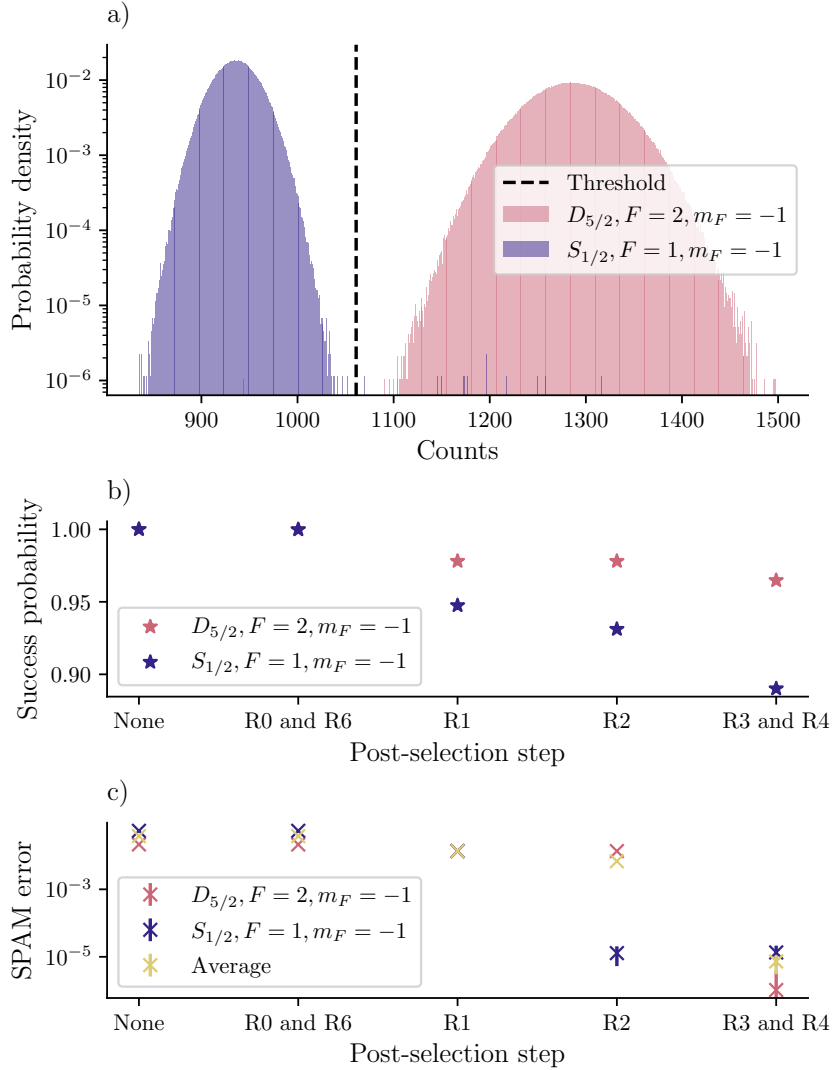


Figure 6.9: (a) Histogram of counts from the camera when performing the SPAM protocol on the O qubit. The measured SPAM error is $7(4) \times 10^{-6}$ after post-selection of 1,000,000 shots on each qubit state. (b) The SPAM error after each post-selection step for the O qubit. (c) The probability of a shot being accepted after each post-selection step for the O qubit.

prepare into to guarantee 50/50 statistics across $|\downarrow\rangle$ and $|\uparrow\rangle$. Next, we drive a $\pi/2$ rotation in the qubit to produce an equal superposition state of $|\downarrow\rangle$ and $|\uparrow\rangle$ ⁴, where we expect $\langle \mathcal{Z}_{\text{real}} \rangle = 0$. We choose this state as it makes the measurement

⁴This may seem redundant given the alternating $|\downarrow\rangle$ and $|\uparrow\rangle$ preparation. However, the state preparation stage also has its own data rate and this will be dependent upon the qubit state we prepare. By both alternating state prep, and creating an equal superposition we generate 50/50 statistics with an error lower than our single qubit infidelity, or that introduced by unequal state preparation rates.

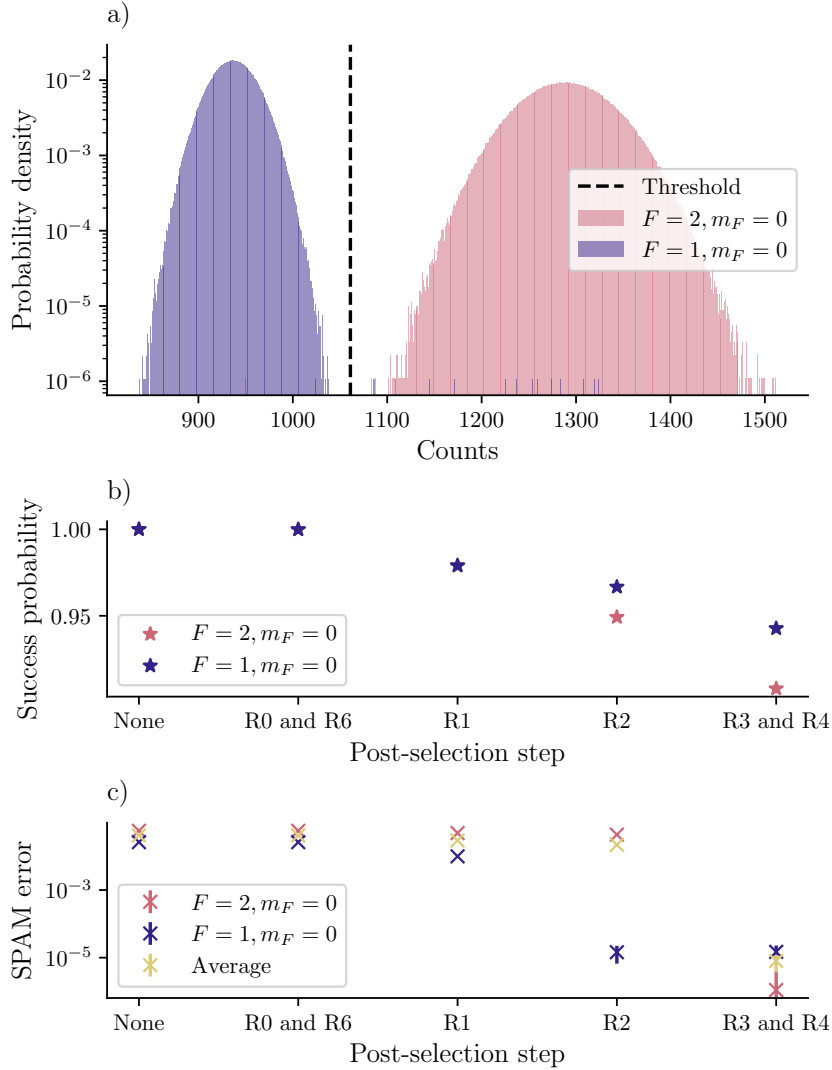


Figure 6.10: (a) Histogram of counts from the camera when performing the SPAM protocol on the G qubit. The measured SPAM error is $8(4) \times 10^{-6}$ after post-selection of 1,000,000 shots on each qubit state. (b) The SPAM error after each post-selection step for the G qubit. (c) The probability of a shot being accepted after each post-selection step for the G qubit.

the most sensitive to bias in the observable statistics. The qubit state is then measured as in Fig. 6.1. We change the duration t of individual population transfer pulses between A and B from its pre-calibrated value of t_π to simulate an error in the pulses. The measurement outcome $\langle \mathcal{Z}_{\text{meas}} \rangle = P(b | a) - P(d | a)$ is then calculated and plotted in Fig. 6.11. In the same figure, we also plot the expected functional form of the bias based on Eq. 6.11 and the expected dependence of $P(a | 0)$

and $P(a | 1)$ as tabulated in Tab. 6.2. We find that the measured bias closely matches the model, indicating that data post-processing using Eq. 6.11 can be used to eliminate bias in observable measurements.




Figure 6.11: A plot of bias $\langle \mathcal{Z}_{\text{meas}} \rangle - \langle \mathcal{Z}_{\text{real}} \rangle$ against the duration t of the population transfer pulse. Each curve corresponds to varying the duration of a pulse mapping the corresponding qubit state between the $S_{1/2}$ and $D_{5/2}$ manifolds. For the optical qubit, $|\downarrow\rangle$ is in the $D_{5/2}$ manifold and $|\uparrow\rangle$ is in the $S_{1/2}$ manifold. For the ground and metastable qubits, the qubit states are in the same manifold, so the curves for varying the duration of the mapping pulse for $|\downarrow\rangle$ and $|\uparrow\rangle$ overlap, and we only show the results for $|\downarrow\rangle$ for clarity. Measurements are plotted against calculations. To reduce the error in statistics from not faithfully preparing an equal superposition, we alternate the initial qubit state $|\downarrow\rangle$ or $|\uparrow\rangle$ from which a superposition is created for each shot.

Name	$P(a 0)$	$P(a 1)$
Optical $ \downarrow\rangle$	$\sin^2(t/2t_\pi)$	1
Optical $ \uparrow\rangle$	1	$\sin^4(t/2t_\pi)$
Metastable $ \downarrow\rangle$	$\sin^2(t/2t_\pi)$	1
Ground $ \downarrow\rangle$	$\sin^4(t/2t_\pi)$	1

Table 6.2: The parameters used to simulate the bias from scans in Fig. 6.11. The entries in Optical $|\downarrow\rangle$ and Metastable $|\downarrow\rangle$ use $\sin^2(t/2t_\pi)$ since the imperfect mapping will result from just a single pulse that maps B ($D_{5/2}$) \rightarrow A ($S_{1/2}$). In the case of Optical $|\uparrow\rangle$ and Ground $|\downarrow\rangle$, there are two pulses, leading to $\sin^4(t/2t_\pi)$. One mapping $A \rightarrow B$ before $R2$, and another $B \rightarrow A$ before $R3$.

7

OTDF Gates

Contents

7.1	Motivating an alternative two-qubit gate	113
7.2	Spin-dependent forces	115
7.2.1	Background	115
7.2.2	Generating a $\hat{\sigma}_z$ force	116
7.3	Entering the Lamb-Dicke regime	120
7.3.1	Sideband cooling	121
7.3.2	Sideband thermometry	122
7.4	Spin-dependent forces with a single ion	124
7.4.1	Phase tracking the motion	125
7.5	Choice of polarisation	128
7.6	Entanglement generation	129
7.7	Gate setup & implementation	131
7.7.1	Heating rate and motional coherence time	132
7.7.2	Mode calibration	134
7.7.3	The spectator mode	135
7.7.4	Balancing light shifts	137
7.8	Results and sources of error	137
7.8.1	Results	137
7.8.2	Sources of error	138
7.9	Summary	152

So far, we have focussed on single-qubit operations within our ion-trap quantum computer. To build a universal quantum processor, however, we must also implement

a two-qubit entangling gate. In this chapter, we apply a variant of the geometric phase gate that is particularly well suited to our optical qubit. In ion-trap systems, two-qubit entangling operations are typically implemented using geometric phase gates, which exploit the shared motional states of the ions to mediate interactions.

7.1 Motivating an alternative two-qubit gate

Two widely used variants are the Mølmer–Sørensen (MS) gate [132, 133] and the $\sigma_Z\sigma_Z$ [134, 135], implemented in laser-based systems via the light shift (LS), which will be our focus. Both are capable of high-fidelity performance, and each has its own set of requirements, advantages, and drawbacks, which we briefly review below.

The MS gate acts in the X basis and can be readily applied to clock qubits. It can even be implemented with co-propagating beams when using optical qubits, reducing the alignment challenges of Raman laser configurations. However, the MS gate has several drawbacks. Firstly, there is a speed limit with which the sidebands can be addressed before off-resonant excitation of the carrier transition becomes overwhelmingly large. This speed limit can be overcome by utilising standing waves at the cost of further experimental complexity [118]. Secondly, the interaction in the X basis means that the driving force must maintain phase coherence with the qubit across interactions. In many systems this is not an issue, but when working with mixed qubit types it adds an additional requirement to drive single-qubit rotations with the two-qubit gate lasers. If these lasers are the same, such as in a purely ground or purely metastable qubit implementation, this is not an issue. However, in a purely optical qubit implementation, the 1762 nm driving laser has a small Lamb-Dicke parameter of 0.012. This weak motional coupling would limit gate performance and so we must look to an alternative. Finally, being in the X basis, the gate interaction does not commute with magnetic field and laser phase noise, meaning that a spin-echo sequence cannot be applied to mitigate such noise.

By contrast, the LS gate acts in the Z basis, which allows the use of spin-echo sequences to mitigate noise. However, LS gates cannot be implemented effectively

in ground or metastable clock qubits, since the differential light shift from dipole coupling is too weak to generate a strong spin-dependent force. Quadrupole couplings can be used [136], but in $^{137}\text{Ba}^+$ this would again require long-wavelength radiation, bringing us back to the issue of a small Lamb–Dicke parameter. In addition, an $\sigma_Z\sigma_Z$ gate typically cannot be implemented with co-propagating beams, as a difference in k -vector is required to generate the motional coupling. One possible workaround is to use a bi-chromatic field, but this requires radiation near resonance with an optical qubit [137]. In our case, this again involves the 1762 nm laser, which is not feasible due to its small Lamb–Dicke parameter.

In this chapter, we explore the optical transition dipole force (OTDF) gate, an alternative implementation of the LS gate that is compatible with optical qubits and works with a driving laser of any wavelength [4]. By working with an optical qubit, we circumvent the small differential light shift in clock qubits and create an LS-type gate compatible with optical clock qubits. As outlined in Chapters 2 and 4, our optical qubit system is limited by laser coherence, so employing a metastable clock qubit would not improve performance. Instead, we exploit the hyperfine structure of $^{137}\text{Ba}^+$ to realise an optical qubit with low first-order magnetic-field sensitivity, retaining the metastable clock qubit for use elsewhere in the architecture. This enables us to demonstrate the OTDF interaction in a system with naturally long coherence times, that also works in optical clock qubits, reducing the need for extensive engineering to suppress dephasing. The first implementation of this gate also achieved an excellent error of $6(3)\times 10^{-4}$ in $35.2\ \mu\text{s}$ using 532 nm light in $^{40}\text{Ca}^+$ which was primarily limited by spin dephasing and $D_{5/2}$ decay [138]. When combined with our SPAM protocol in Chapter 6, errors from qubit loss can be further mitigated allowing this gate to integrate nicely with our existing framework.

7.2 Spin-dependent forces

7.2.1 Background

The basis of all geometric phase gates is applying a spin-dependent force (SDF) to a set of ions. In this section, we will study the force in isolation and apply it to a single qubit, allowing us to gain intuition for the following section where we construct a two qubit gate from this force.

Under the Lamb-Dicke and rotating-wave approximations, a spin-dependent force on a quantum harmonic oscillator coupled to a two-level system with resonant motional frequency ω_0 and driving frequency ω can be described by the following Hamiltonian [139],

$$\hat{H}_F = -\frac{F_0 z_0}{2} \hat{\sigma}_\phi (\hat{a} e^{i\delta t + i\phi_0} + \hat{a}^\dagger e^{-i\delta t - i\phi_0}), \quad (7.1)$$

where $\delta = (\omega - \omega_0)$, $z_0 = \sqrt{\frac{\hbar}{2m\omega_0}}$ and m is the mass of the ion. The spin-dependence comes from the $\hat{\sigma}_\phi$ term, and ϕ determines the basis in which the force acts on the ion. MS gates are performed in the x or y basis (or some linear combination), and LS gates in the z basis. Since we will focus on LS gates later in this chapter, we will assume the z basis from now onwards.

From \hat{H}_F we can apply the Magnus expansion [140, 141] to construct the propagator \hat{U}_F ,

$$\hat{U}_F(t) = \hat{\mathcal{D}}(\alpha(t) \hat{\sigma}_z) e^{-i\Phi(t) \hat{\sigma}_z^2}, \quad (7.2)$$

where $\hat{\mathcal{D}}(\alpha) = e^{\alpha \hat{a}^\dagger - \alpha^* \hat{a}}$ is the displacement operator. The parameters $\alpha(t)$ and $\Phi(t)$ are given by,

$$\begin{aligned} \alpha(t) &= i \frac{z_0 F_0 t}{2\hbar} \text{sinc}\left(\frac{\delta t}{2}\right) e^{i\phi_0} e^{-i\delta t/2} \\ \Phi(t) &= \left(\frac{z_0 F_0}{\hbar}\right)^2 \frac{1}{4\delta^2} [\delta t - \sin(\delta t)]. \end{aligned} \quad (7.3)$$

Taking the eigenstates and eigenvalues of σ_z for a single spin as $|\uparrow\rangle / |\downarrow\rangle$ and $+1 / -1$ respectively, we can write the effect of the propagator more naturally

in the z eigenbasis,

$$\hat{U}_F(t) = \hat{\mathcal{D}}(\alpha(t))e^{-i\Phi(t)} |\uparrow\rangle \langle\uparrow| + \hat{\mathcal{D}}(-\alpha(t))e^{-i\Phi(t)} |\downarrow\rangle \langle\downarrow|. \quad (7.4)$$

The action of this propagator is to displace the ion in phase space by some amount α , entangling the spin and motion. With some finite detuning δ , $\alpha(t)$ is periodic with period $4\pi/\delta$, so the qubit states trace closed loops in phase space that return to the origin every $t = 2\pi/\delta$, disentangling the spin and motion. During this evolution the spins acquire a phase Φ . For a single spin this phase is global, but in multi-qubit systems it becomes relative and provides the mechanism for generating entanglement.

7.2.2 Generating a $\hat{\sigma}_z$ force

As seen in Chapter 4, the light shift serves as a $\hat{\sigma}_z$ operation on a single qubit. While this is often an unwanted term in our single qubit operations, we can use this to construct a spin-dependent force, and thus, a two-qubit gate. The light shift from a single beam can produce a spin-dependent force. However, the field oscillations are in the THz regime, many orders of magnitude faster than the ions motion (MHz). As such, the motional coupling from a single beam can be considered a fast rotating term which is averaged over the course of a gate. To produce a meaningful light shift force, we must use the beating of two lasers, similar to the Raman transition in Sec. 4.1.4.

Two-laser light shifts

We begin by following the treatment in Ref. [4]. Consider the application of two lasers $\hat{E}_1(\vec{r}, t)$ and $\hat{E}_2(\vec{r}, t)$

$$\hat{E}_i(\vec{r}, t) = \vec{\hat{e}}_i E_i \cos(\vec{k}_i \cdot \vec{r} - \omega_i t - \phi_i), \quad (7.5)$$

where $\vec{\hat{e}}$ is a polarisation vector which we drop in further discussion, considering only two beams of matching polarisation¹.

¹Differing polarisations can be included by considering each polarisation term as a separate electric field and summing the light shifts.

When investigating the gate dynamics, we are interested in time-scales around the motional mode period (\sim MHz), which is also approximately the frequency difference between the beams. However, this MHz scale is much slower than the \sim 100 THz optical frequencies that each individual electric field oscillates with. Thus, we time average over optical frequencies and the light shift due to dipole interactions on a state $|i\rangle$ is [4],

$$\begin{aligned}
\Delta_i(\vec{r}, t) &= \gamma_i \langle E(\vec{r}, t)^2 \rangle \\
&= \gamma_i \langle E_1^2 \cos(\vec{k}_1 \cdot \vec{r} - \omega_1 t - \phi_1)^2 \\
&\quad + 2E_1 E_2 \cos(\vec{k}_1 \cdot \vec{r} - \omega_1 t - \phi_1) \cos(\vec{k}_2 \cdot \vec{r} - \omega_2 t - \phi_2) \\
&\quad + E_2^2 \cos(\vec{k}_2 \cdot \vec{r} - \omega_2 t - \phi_2)^2 \rangle, \\
\gamma_i &= \sum_k \frac{|\langle i | \vec{d} \cdot \vec{\epsilon} | k \rangle|^2}{4\hbar} \left(\frac{1}{\omega_{ki} - \omega_{\text{laser}}} + \frac{1}{\omega_{ki} + \omega_{\text{laser}}} \right).
\end{aligned} \tag{7.6}$$

This can be rearranged using $\cos(a) \cos(b) = \frac{1}{2}[\cos(a-b) + \cos(a+b)]$,

$$\begin{aligned}
\Delta_i(\vec{r}, t) &= \gamma_i \langle \frac{E_1^2}{2} [1 + \cos(2\vec{k}_1 \cdot \vec{r} - 2\omega_1 t - 2\phi_1)] \\
&\quad + E_1 E_2 [\cos(\vec{k}_\Sigma \cdot \vec{r} - \omega_\Sigma t - \phi_\Sigma) + \cos(\vec{k}_\Delta \cdot \vec{r} - \omega_\Delta t - \phi_\Delta)] \\
&\quad + \frac{E_2^2}{2} [1 + \cos(2\vec{k}_2 \cdot \vec{r} - 2\omega_2 t - 2\phi_2)] \rangle,
\end{aligned} \tag{7.7}$$

where X_Σ denotes the sum of two variables X_1, X_2 , and X_Δ denotes the difference. The terms with ω_1, ω_2 , and ω_Σ oscillate with optical frequencies and so are averaged over leaving,

$$\Delta_i(\vec{r}, t) = \gamma_i \left(\frac{E_1^2 + E_2^2}{2} + E_1 E_2 \cos(\vec{k}_\Delta \cdot \vec{r} - \omega_\Delta t - \phi_\Delta) \right). \tag{7.8}$$

The first term in Eq. 7.8 is the arithmetic mean of the two individual laser light shifts and provides a single qubit light shift. This term contains no motional coupling and therefore does not participate in the spin-dependent force. Thus, we will not consider it from here onwards. This leaves the second term, which will soon become the spin-dependent force. The amplitude of this term is the geometric mean of the two light shifts and allows us to gain an intuition of how much the laser intensity

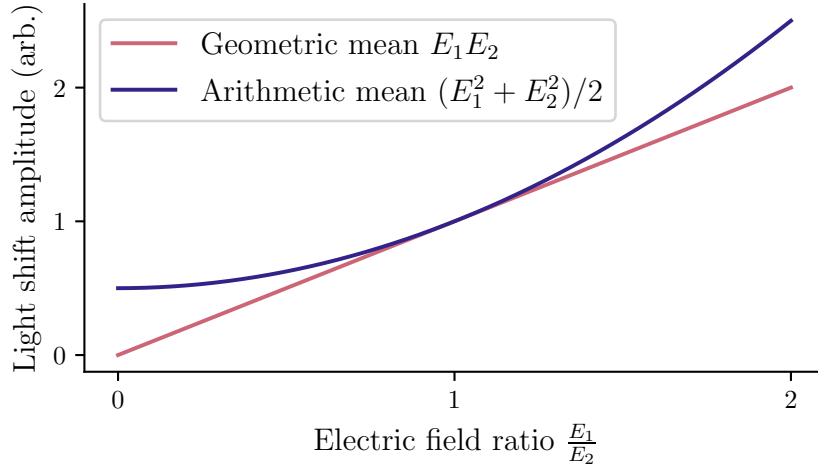


Figure 7.1: Illustration of the divergence between the arithmetic mean and the geometric mean of two variables. The two are equal when the variables are the same. As such, the maximum efficiency in the Rabi frequency of the OTDF gate for a given power occurs when the two beams are of equal intensity.

contributes to the spin-dependent force. As illustrated in Fig. 7.1, the arithmetic mean is greater than the geometric mean for all ratios of electric fields except for $E_1 = E_2$. This means that the optimal light shift gate will have a balanced light shift contribution from each beam, and any imbalance is wasted laser power.

Two-beam light shift Hamiltonian

Each state in our ion will have a different coefficient γ_i^2 and therefore a different light shift. When considering many ions, we relabel this $\gamma_{j,i}$ where j denotes the ion and i the state. Only taking the second term from Eq. 7.8, we use this to construct the light shift Hamiltonian for N ions.

$$H_{LS} = \sum_{j=1}^N E_1 E_2 \left(\gamma_{j,\downarrow} |\downarrow\rangle \langle \downarrow|_j + \gamma_{j,\uparrow} |\uparrow\rangle \langle \uparrow|_j \right) \cos(\vec{k}_\Delta \cdot \vec{r} - \omega_\Delta t - \phi_\Delta), \quad (7.9)$$

where $|x\rangle \langle x|_j$ implicitly contains the tensor product with the identity for other qubits. Next, we write the light shift in terms of the sum and difference terms

²With the exception of ground clock qubits, where there is no differential light shift.

$\Delta_{j,\Delta} = \frac{E_1 E_2}{2}(\gamma_{j,\uparrow} - \gamma_{j,\downarrow})$ and $\Delta_{j,\Sigma} = \frac{E_1 E_2}{2}(\gamma_{j,\uparrow} + \gamma_{j,\downarrow})$ giving,

$$\begin{aligned} H_{LS} &= \sum_{j=1}^N \left([\Delta_{j,\Sigma} + \Delta_{j,\Delta}] |\uparrow\rangle \langle\uparrow|_j + [\Delta_{j,\Sigma} - \Delta_{j,\Delta}] |\downarrow\rangle \langle\downarrow|_j \right) \\ &\quad \times \cos(\vec{k}_\Delta \cdot \vec{r} - \omega_\Delta t - \phi_\Delta) \\ &= \sum_{j=1}^N (\Delta_{j,\Sigma} \hat{\mathbf{1}}_j + \Delta_{j,\Delta} \hat{\sigma}_{z,j}) \cos(\vec{k}_\Delta \cdot \vec{r} - \omega_\Delta t - \phi_\Delta). \end{aligned} \quad (7.10)$$

From this we obtain two contributions. The first arises from the common light shift, which generates a motional coupling but no spin–motion coupling. The second is the differential light shift, which produces the spin–motion coupling and constitutes the key term in the Hamiltonian.

Writing $\vec{k}_\Delta \cdot \vec{r}$ in terms of the raising and lowering operators for each mode k ³ and exponentiating the cosine we find

$$H_{LS} = \sum_{j=1}^N \frac{1}{2} \left([\Delta_{j,\Sigma} \hat{\mathbf{1}}_j + \Delta_{j,\Delta} \hat{\sigma}_{z,j}] e^{i \sum_k \eta_{k,j} (\hat{a}_k + \hat{a}_k^\dagger)} e^{-i\omega_\Delta t - i\phi_j} \right) + \text{H.c.} \quad (7.11)$$

where $\eta_{k,j} = |\vec{k}_\Delta| \lambda_{k,j} \sqrt{\frac{\hbar}{2m_j \omega_k}}$ is the Lamb-Dicke parameter. The value $\lambda_{k,j}$ is the eigenvalue of ion j 's participation in motional mode k with mass m_j and angular mode frequency ω_k . We have also written the phase ϕ_Δ in terms of individual ion phases ϕ_j .

Finally, we take H_{LS} into the Lamb-Dicke regime, keeping only first order terms in $\eta_{k,j}$, and move into the interaction picture with respect to the motion, defining the detuning $\delta_k = \omega_k - \omega_\Delta$,

$$H_{LS} \approx \frac{1}{2} \sum_{j=1}^N \left([\Delta_{j,\Sigma} \hat{\mathbf{1}}_j + \Delta_{j,\Delta} \hat{\sigma}_{z,j}] [\hat{\mathbf{1}}_k + i \sum_k (\hat{a}_k^\dagger \eta_{k,j} e^{i\delta_k t - i\phi_j} + \hat{a}_k \eta_{k,j} e^{-i\delta_k t + i\phi_j})] \right). \quad (7.12)$$

The zeroth order term in a_k (or a_k^\dagger) is a single ion light shift, and does not participate in the spin-dependent force, though it must be tracked or echoed out to maintain phase coherence. As such, in all experimental data shown we split the SDF into two stages with an echo in between to compensate this light shift. The first order term has the same form as the spin-dependent force in Eq. 7.1 and is the main point of interest.

³No relation to the wave-vector \vec{k} .

7.3 Entering the Lamb-Dicke regime

When taking the final step to reach Eq. 7.12, we make the approximation that the ion is in the Lamb-Dicke regime and so terms of order $\eta_{k,j}^2$ can be ignored. This approximation is valid in the case $\eta^2(2n + 1) \ll 1$ [76], where n is the Fock state occupation number of the ions. The Lamb-Dicke parameter $\eta_{k,j} = |\vec{k}_\Delta| \lambda_{k,j} \sqrt{\frac{\hbar}{2m_j\omega_k}}$ for a $^{137}\text{Ba}^+$ ion with motional mode frequency of 1.51 MHz (the lower radial rocking mode in our setup) with two 532 nm driving lasers is 0.0823. As such, in the best case scenario when $n = 0$ we have $\eta^2([2 \times 0] + 1) = 6.78 \times 10^{-3}$. However, any thermal occupation will move us away from this best case scenario. From Sec. 3.3, we saw that the mean thermal occupation from Doppler cooling is $\bar{n} \approx 13$, and thus, $\mathbb{E}[\eta^2(2n + 1)] \approx 0.18$. At this level, the Lamb-Dicke approximation is a poor one.

Assuming a thermal distribution,

$$P(n, \bar{n}) = \frac{1}{\bar{n} + 1} \left(\frac{\bar{n}}{\bar{n} + 1} \right)^n, \quad (7.13)$$

we plot the mean, 80th percentile, and 95th percentile of $\eta^2(2n + 1)$ in Fig. 7.2, illustrating that for even small values of $\bar{n} \sim 1$, there is significant population far from the Lamb-Dicke regime. Additionally, we see that the 95th percentile of Fock states diverges from the Lamb-Dicke regime much more quickly than the bulk of the distribution, meaning some finite temperature will introduce very strong out of Lamb-Dicke effects for the upper part of the thermal distribution.

To model the gate dynamics accurately, we must keep our ions within the Lamb-Dicke regime and therefore as cold as possible. We do this with two stages of sub-Doppler cooling. Firstly, electromagnetically induced transparency cooling (EIT) provides fast broadband cooling bringing all motional modes to $\bar{n} \approx 0.5$. This is cold enough to reduce most spectator mode effects, however, we aim to make the gate mode as cold as possible by finishing the cooling sequence with sideband cooling (SBC).




Figure 7.2: Simulation of the mean value of $\eta^2(2n + 1)$ considering a full thermal distribution as well as the 80th and 95th percentile of the distribution. We see that for even small values of $\bar{n} \sim 1$ there is a significant deviation from the Lamb-Dicke regime.

7.3.1 Sideband cooling

Typically sideband cooling involves repeated cycles of a sideband pulse on a ground qubit driven with Raman lasers, and state preparation to reset the qubit [67]. Unfortunately, in our optical-/metastable-qubit system we do not have access to a ground qubit. Another common approach would be to sideband cool with the optical qubit [142]. For species such as Ca^+ or Sr^+ , the optical qubits have favourable wavelengths for driving a motional sideband (larger Lamb-Dicke parameter), 729 nm and 674 nm respectively. However, the optical qubit in Ba^+ is driven with 1762 nm light, a considerably longer wavelength. When working with a 1.6 MHz radial mode, the optical qubit will have a Lamb-Dicke parameter of 0.012. As shown in Sec. 4.1.2, for small n the sideband Rabi frequency scales as $\Omega_{\text{SB}} \approx \eta\Omega_0\sqrt{n+1}$, meaning that the optical qubit sideband will be $\sim 80\times$ slower than the carrier. With a carrier π -time of $\sim 15 \mu\text{s}$, we can expect a sideband duration of up to 1.2 ms! When battling heating from trap imperfections, this is a considerable amount of time and will significantly reduce the effectiveness of sideband cooling.

This leaves the metastable qubit as our other option for sideband cooling. Fortunately, the Lamb-Dicke parameter of this qubit is $\eta \approx 0.08$ which is much

more favourable. Additionally, our system has a very large amount of 532 nm laser power with which to drive this qubit. As such, we can achieve carrier π -times of $\sim 10 \mu\text{s}$ and therefore sidebands with π -times of $\sim 120 \mu\text{s}$. The only downside to this is the preparation steps involved for the metastable qubit are more involved than the optical qubit, which only requires optical pumping.

The steps to perform sideband cooling on the metastable qubit are shown in Fig. 7.3. We begin with the typical optical pumping for state preparation described in Sec. 3.4, preparing into $|\downarrow\rangle_G = |S_{1/2}, F = 2, m_F = 0\rangle$. Next the 1762 nm laser drives population into $|\downarrow\rangle_{M,n} = |D_{5/2}, F = 3, m_F = +2\rangle$. We then apply a red sideband with the 532 nm lasers, ideally removing one quanta of motion such that $n \rightarrow n - 1$, driving into the state $|\uparrow\rangle_{M,n-1} = |D_{5/2}, F = 2, m_F = +2\rangle$. Before clearing population from the $D_{5/2}$ manifold with the 614 nm laser, we coherently deshelve population in $|\downarrow\rangle_M$ to $|\downarrow\rangle_G$. This is to protect population in the motional ground state $n = 0$ from dissipative processes which may excite the motion. Finally, the 614 nm laser clears all population from $D_{5/2}$ to $S_{1/2}$ via $P_{3/2}$, emitting a 455 nm photon.

7.3.2 Sideband thermometry

Given how critically temperature impacts the dynamics of our SDF, it is important to measure the temperature of our ions after cooling. To do this, we measure the relative amplitude of the red (motion subtracting) and blue (motion adding) sidebands. As the ion gets colder, the strength of the red sideband reduces to zero and the blue sideband increases. In the limit of low thermal occupation this gives us an estimate of \bar{n} assuming a thermal distribution⁴.

To see how the relative amplitude of the red and blue sideband gives a measure of temperature we follow the explanation in Appendix A of Ref. [68]⁵.

⁴Sideband cooling does not in general leave behind a thermal distribution. However, the low starting thermal occupation from EIT cooling means that this is a good approximation.

⁵This paper also illustrates a more sophisticated approach to temperature measurement which is yet to be implemented in our system.

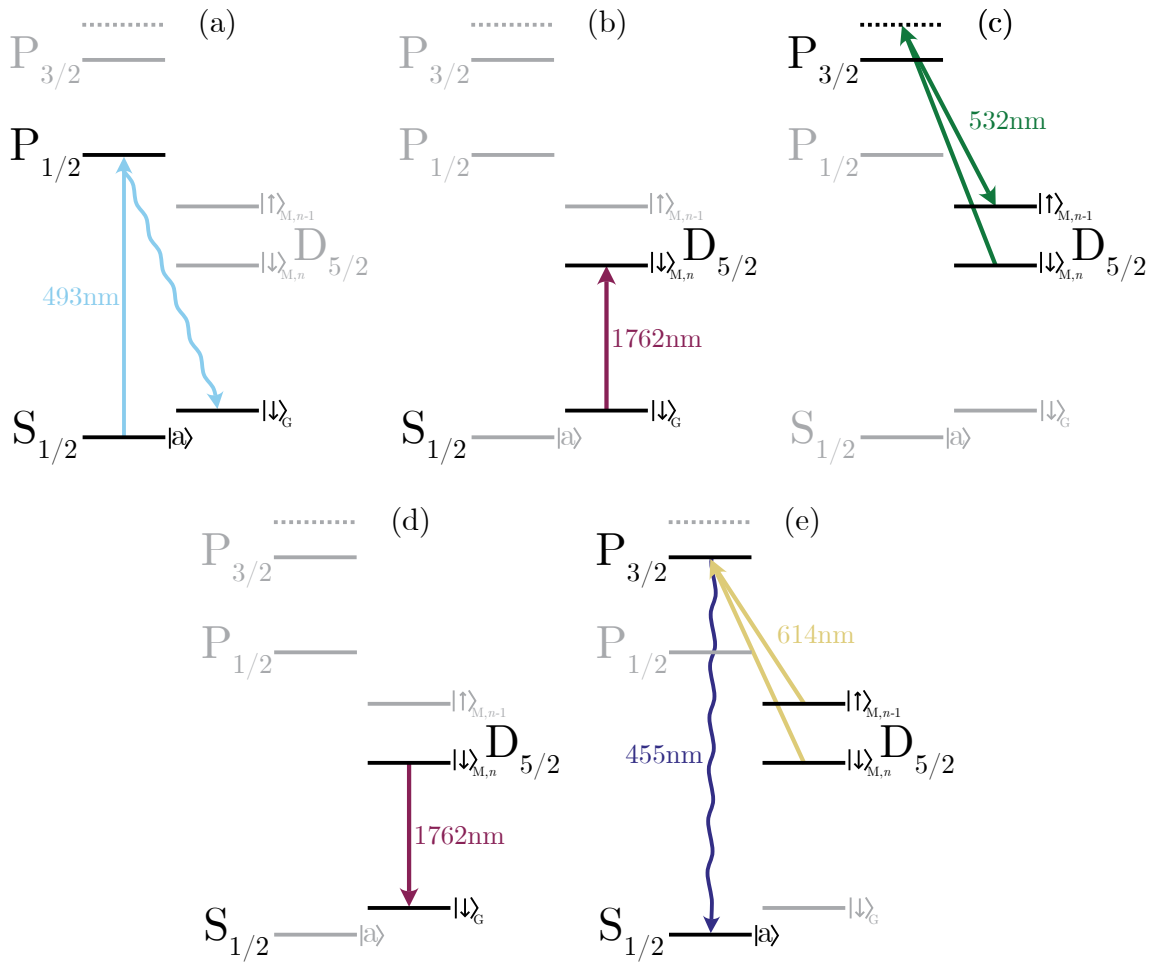


Figure 7.3: Step by step diagram of sideband cooling with the metastable qubit. (a) Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$ as described in Sec. 3.4. (b) Shelving from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{5/2}, F = 3, m_F = +2\rangle$. (c) A π -pulse with the Raman lasers on the red sideband of $|D_{5/2}, F = 3, m_F = +2\rangle \rightarrow |D_{5/2}, F = 2, m_F = +2\rangle$. (d) Deshelving with the 1762 nm to protect population in $n = 0$. (e) Deshelving the $D_{5/2}$ with the 614 nm laser.

Starting in $|\downarrow\rangle$ and driving a blue/red sideband, the probability of measuring




Figure 7.4: Sideband thermometry data for the lower frequency radial mode taken on a single ion. The temperature is estimated by fitting to Eq. 7.15 and we measure $\bar{n} = 0.40(5)$ phonons.

the state $|\uparrow\rangle$ is given by,

$$\begin{aligned}
 P_{\uparrow}^{\text{BSB}}(\bar{n}, t) &= \sum_{n=0}^{\infty} P(n, \bar{n}) \sin^2\left(\frac{\Omega_{n+1, n} t}{2}\right), \\
 P_{\uparrow}^{\text{RSB}}(\bar{n}, t) &= \sum_{n=1}^{\infty} P(n, \bar{n}) \sin^2\left(\frac{\Omega_{n, n-1} t}{2}\right), \\
 &= \frac{\bar{n}}{\bar{n} + 1} \sum_{n=0}^{\infty} P(n, \bar{n}) \sin^2\left(\frac{\Omega_{n+1, n} t}{2}\right), \\
 &= \frac{\bar{n}}{\bar{n} + 1} P_{\uparrow}^{\text{BSB}}(\bar{n}, t).
 \end{aligned} \tag{7.14}$$

Therefore, the mean occupation assuming a thermal state is

$$\bar{n} = \frac{P_{\uparrow}^{\text{RSB}}}{P_{\uparrow}^{\text{BSB}} - P_{\uparrow}^{\text{RSB}}}. \tag{7.15}$$

In Fig. 7.4, we plot the red and blue sideband on the lower frequency radial mode of a single ion after 2 ms of EIT cooling and 10 SBC cycles. By fitting Ω_0 and \bar{n} we estimate a temperature of $\bar{n} = 0.40(5)$ phonons.

7.4 Spin-dependent forces with a single ion

Now that the ion is in a well-defined thermal state, we can characterise the spin-dependent force that we apply. This is a useful step to take before performing a full




Figure 7.5: OTDF gate sequence. We begin by performing a $\pi/2$ -pulse to change basis such that the σ_z SDF we create also induces a relative phase. The spin-dependent force is then applied for a duration $\tau/2$. This is followed by an echo pulse in the Y basis and then another spin-dependent force. The Ramsey-like sequence is then closed with a final $\pi/2$ -pulse in the X basis. When measuring the parity of the final state, there is a further $\pi/2$ -pulse with variable phase ϕ . This is not depicted in the sequence.

two-qubit gates as we can characterise the force being applied to the ion without the complex dynamics of spin-spin entanglement.

We implement the sequence shown in Fig. 7.5. We begin by performing a $\pi/2$ -pulse in the σ_x basis, preparing the superposition $\frac{1}{\sqrt{2}}(|\downarrow\rangle - i|\uparrow\rangle)$. This places the qubit in a basis in which the σ_z SDF will cause a relative phase to be acquired. The SDF is then applied for some duration $\tau/2$. After this, we apply a spin-echo pulse in the σ_y basis. By flipping the spin states, the light shifts which do not couple to the motion are undone by the second SDF which follows for another $\tau/2$. Finally, the sequence is closed by returning to the σ_z basis with another $\pi/2$ -pulse in the σ_x basis. After this sequence the probability of measuring $|\downarrow\rangle$ is given by Eq. 7.16 [139],

$$P_{|\downarrow\rangle}(\tau) = \frac{1}{2} \left(1 - e^{-\langle \bar{n} \tau | 4\alpha_0 |^2 / 2 \rangle - \langle [\bar{n} + 1/2] | 2\alpha(\tau) |^2 \rangle} \right), \quad (7.16)$$

where $\alpha_0 = \Delta_\Delta / 2\delta$ and $\alpha(\tau) = \alpha_0 \text{sinc}\left(\frac{\delta\tau}{2}\right) (e^{i\phi_0} + e^{i\phi_1}) e^{-i\delta\tau/2}$ similar to Eq. 7.3, but with the addition of a second phase ϕ_1 which is the phase of the second SDF pulse. In Fig. 7.6, we plot this probability and compare to simulation using Eq. 7.16. In the simulation we use $\bar{n} = 0.4$ and $\dot{\bar{n}} = 40$ (see Sec. 7.7.1 for details) and $\Delta_\Delta = 31.4$ kHz which is measured from single Raman beam light shifts. We float the centre frequency as a free parameter in the simulation.

7.4.1 Phase tracking the motion

From the parameter $\alpha(\tau)$ in Eq. 7.16 arises a new variable, $\delta\phi = \phi_0 - \phi_1$. This is the relative phase of the two forces. We see that for $\delta\phi = \pi$ then $\alpha(\tau) = 0$ for




Figure 7.6: Data from a spin-dependent force applied to a single ion. As the detuning is adjusted points are found where loops in phase space are closed, leading to no spin-motion entanglement. We compare to simulation of Eq. 7.16 using $\bar{n} = 0.4$, $\dot{\bar{n}} = 40$, and $\Delta_{\Delta} = 31.4$ kHz which have all been measured independently.

all values of τ . As a result, we must ensure this phase is well known, else the two forces may combine in unpredictable ways. The detuning of the force relative to the gate mode, combined with the echo pulse adds an additional layer of complexity to this. For a detuning δ and echo pulse of duration t_{echo} , the phase of the Raman LS force advances by $\delta t_{\text{echo}} + \pi$ after the echo, where the π accounts for the flip in spin states [40]. This means that we must know δ and t_{echo} well to maintain coherence between the two forces.

When tracking $\delta\phi$, there are two options that give different dynamics. Each of these are shown in Fig. 7.7. If we choose $\phi_1 = \phi_0 + \pi + \delta\tau/2$, the phase space trajectory will continue between the two SDF pulses, giving the “continuous” trajectory shown in the figure. If the pulses are timed such that $\tau = 2\pi/\delta$, then two full loops in phase space will be performed as shown in Fig. 7.7(a). However, Fig. 7.7(b) shows the case where $\tau < 2\pi/\delta$ and the “continuous” pulse does not result in a closed trajectory. This leaves the spin and motion entangled, which is undesirable for the two-qubit gate as we will see later.

The alternative, which we call “discrete”, sets $\phi_1 = \phi_0 + \pi$. This results in a flip in direction in phase space after the first pulse. For two sequences of equal length

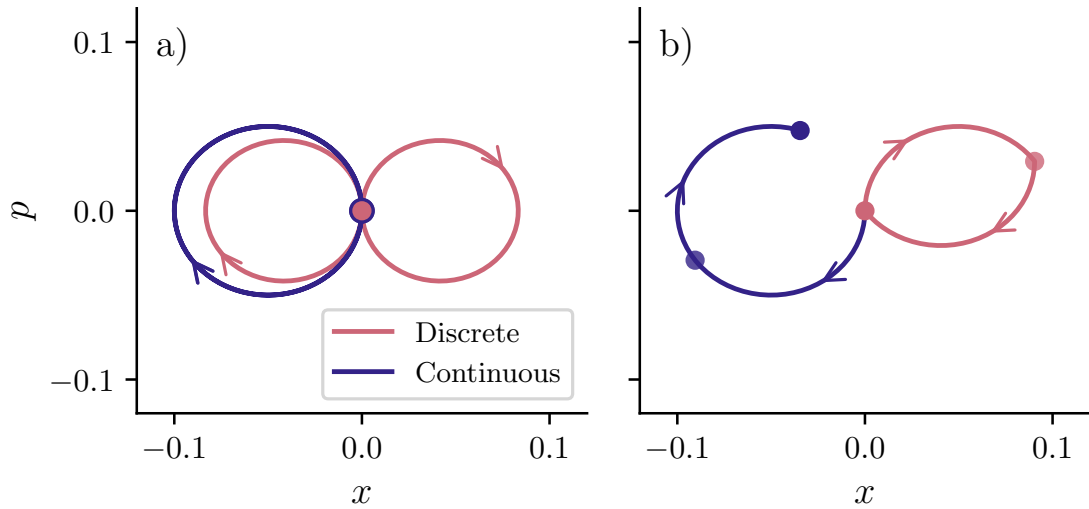


Figure 7.7: Simulation of the phase space trajectories for two applications of a spin-dependent force with a π -pulse on the qubit separating them. There are two cases, “continuous” where the initial phase (relative to the ions motion) of the second force continues from the final phase of the first force, and “discrete” where the two forces have the same initial phase. (a) When the timing condition $t = 2\pi/\delta$ is met, then the loops are closed in both cases (the continuous plot has two loops on top of each other). If there is some error in timing, only the “discrete” case closes loops due to the π phase-flip of the force due to the π -pulse.

the phase space trajectory is always closed, similar to Walsh modulation [141].

In Fig. 7.8 we implement each of the “discrete” and “continuous” phase tracking modes. We use a detuning $\delta = 2\pi \times 10$ kHz. For both methods when two full loops in phase space are performed (at $\tau = 200 \mu\text{s}$), $P(|\downarrow\rangle) \rightarrow 0$. This means that in Fig. 7.8(a), the trajectories are closed for all force durations. However, in (b) there is a particular phase which closes the trajectory at each time, but any error in timing will cause the spin and motion to remain entangled at the end of the sequence. For this reason, we choose to implement the “discrete” phase tracking in our gate sequences. In both cases a miscalibration in δ will result in the data in Fig. 7.8 being shifted vertically by an offset in $\delta\phi$ which will also leave spin motion entanglement.




Figure 7.8: Experimental data illustrating the difference between the “discrete” and “continuous” applications of the SDF with detuning $\delta = 2\pi \times 10$ kHz on a single ion. A phase $\delta\phi$ is added to the second force relative to the first. Points where $p(|\downarrow\rangle) \rightarrow 0$ indicate closed loops in phase space. We observe that for any time where a phase of 0 or (equivalently) 1 turns is used in the “discrete” case, then loops in phase space are closed. However, a timing error in the “continuous” case will leave residual spin motion entanglement. As such, the “discrete” case is desirable when implementing two-qubit gates to add robustness to timing or detuning errors.

7.5 Choice of polarisation

In Eq. 7.5, we define the electric field as the sum of polarisation components \hat{e} . From this point onwards, we assume a single polarisation which is the same for the two beams which produce the SDF. But what polarisation should these beams take? From Chapter 3, Fig. 3.1, we have a 45° angle between the B-field and the counter-propagating Raman beam, and a -135° angle to the co-propagating Raman beam. This will restrict our choice of polarisation in the Raman beams which can produce a travelling wave at the ion.

In the spherical basis, a laser beam with angle ϕ to the magnetic field, with

polarisation angle γ and phase θ has electric field components,

$$\begin{aligned} E_{\sigma+} &= \frac{E}{\sqrt{2}}(-\cos(\gamma)\cos(\phi) - i\sin(\gamma)e^{i\theta}), \\ E_{\sigma-} &= \frac{E}{\sqrt{2}}(\cos(\gamma)\cos(\phi) - i\sin(\gamma)e^{i\theta}), \\ E_{\pi} &= -E\cos(\gamma)\sin(\phi). \end{aligned} \tag{7.17}$$

Plugging in $\phi^{\text{co}} = 3\pi/4$ and $\phi^{\text{counter}} = \pi/4$ and solving for $E_i^{\text{co}} = E_i^{\text{counter}}$ for $i \in \{\sigma+, \sigma-, \pi\}$ we find $\gamma^{\text{counter}} = \pi - \gamma^{\text{co}}$. This constrains us to using linear polarisation in both beams, with equal and opposite angles. We choose to use vertical polarisation for the ease of setup and verification with a polarising beam-splitter.

7.6 Entanglement generation

Now that we understand the effect of the spin-dependent force on a single ion, we can use it to generate entanglement. With the light shift gate we do this to implement a unitary of the form,

$$U_{\text{CZ}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{7.18}$$

This can be combined with Hadamards on either side to produce the maximally entangled state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ starting from $|00\rangle$. This entangling gate combined with access to arbitrary single qubit unitaries produces the universal gate set required for quantum computing [12].

In Eq. 7.4, we wrote the effect of the spin-dependent force on a single qubit. More generally, writing the state of a many qubit system in the eigenbasis of the SDF (conveniently Z in this case, though we keep things general at first), the propagator can be written as,

$$\hat{U}_F(t) = \sum_{\lambda} \hat{D}(\alpha(t) |\lambda\rangle \langle \lambda|) e^{-i\lambda^2 \Phi(t)} |\lambda\rangle \langle \lambda|, \tag{7.19}$$

where in the Z basis the eigenstates are $|\lambda\rangle \in \{|\downarrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\uparrow\uparrow\rangle\}$. Putting this into context with the OTDF force Hamiltonian in Eq. 7.12, for two ions, we can use the following components,

$$\begin{aligned} F_{k,\Delta,+} &= \frac{i}{2} \left(\eta_{k1} e^{-i\phi_1} \Delta_{1,\Delta} + \eta_{k2} e^{-i\phi_2} \Delta_{2,\Delta} \right), \\ F_{k,\Delta,-} &= \frac{i}{2} \left(\eta_{k1} e^{-i\phi_1} \Delta_{1,\Delta} - \eta_{k2} e^{-i\phi_2} \Delta_{2,\Delta} \right), \\ F_{k,\Sigma} &= \frac{i}{2} \left(\eta_{k1} e^{-i\phi_1} \Delta_{1,\Sigma} - \eta_{k2} e^{-i\phi_2} \Delta_{2,\Sigma} \right), \end{aligned} \quad (7.20)$$

to write out the Hamiltonian in the Z basis.

$$\begin{aligned} \hat{H}_{LS,k}(|\downarrow\downarrow\rangle) &= (F_{k,\Delta,+} + F_{k,\Sigma}) e^{i\delta_k t} \hat{a}_k^\dagger + \text{H.c.}, \\ \hat{H}_{LS,k}(|\downarrow\uparrow\rangle) &= (F_{k,\Delta,-} + F_{k,\Sigma}) e^{i\delta_k t} \hat{a}_k^\dagger + \text{H.c.}, \\ \hat{H}_{LS,k}(|\uparrow\downarrow\rangle) &= (-F_{k,\Delta,-} + F_{k,\Sigma}) e^{i\delta_k t} \hat{a}_k^\dagger + \text{H.c.}, \\ \hat{H}_{LS,k}(|\uparrow\uparrow\rangle) &= (-F_{k,\Delta,+} + F_{k,\Sigma}) e^{i\delta_k t} \hat{a}_k^\dagger + \text{H.c.} \end{aligned} \quad (7.21)$$

Labelling the coefficient in each term of the Hamiltonian as $F_k^{[i]}$ where $|i\rangle \in \{|\downarrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\uparrow\uparrow\rangle\}$, the geometric phase acquired after a single pulse is,

$$\Phi_k^{[i]}(t) = \frac{|F_k^{[i]}|^2}{\delta_k^2} [\delta_k t - \sin(\delta_k t)]. \quad (7.22)$$

We apply the force for a time $\tau = 2n\pi/\delta$ to close the phase space loops traversed, leading to net displacement $\alpha(\tau) = 0$. This single application of the SDF has two problems. Firstly, the phase $\Phi_k^{[i]}$ is different for all $|i\rangle$, meaning we cannot implement U_{CZ} without making a very specific choice of the amplitude and detuning. Secondly, there is the single-qubit light shifts which must either be tracked or echoed out. Fortunately, the spin-echo solves both of these issues. After the spin-echo pulse we see $|\downarrow\downarrow\rangle \leftrightarrow |\uparrow\uparrow\rangle$ and $|\uparrow\downarrow\rangle \leftrightarrow |\downarrow\uparrow\rangle$. After the second application of the SDF, an additional phase is acquired making the total geometric phase,

$$\begin{aligned} \Phi_k^{|\text{Even}\rangle}(t) &= \Phi_k^{|\uparrow\uparrow\rangle}(t/2) + \Phi_k^{|\downarrow\downarrow\rangle}(t/2) \\ \Phi_k^{|\text{Odd}\rangle}(t) &= \Phi_k^{|\downarrow\uparrow\rangle}(t/2) + \Phi_k^{|\uparrow\downarrow\rangle}(t/2), \end{aligned} \quad (7.23)$$

where $|\text{Even}\rangle \in \{|\uparrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$, and $|\text{Odd}\rangle \in \{|\downarrow\uparrow\rangle, |\uparrow\downarrow\rangle\}$. We are only interested in the relative phase between states, so this reduces three phases down to one,

$$\Delta\Phi^{\text{Total}}(t) = 2\Delta_{1,\Delta}\Delta_{2,\Delta} \cos(\phi_2 - \phi_1) \sum_k \frac{\eta_{k1}\eta_{k2}}{\delta_k^2} [\delta_k t - \sin(\delta_k t)], \quad (7.24)$$

which produces U_{CZ} when $\Delta\Phi^{\text{Total}} = \left| \sum_k (\Phi_k^{\text{Even}} - \Phi_k^{\text{Odd}}) \right| = \pi/2$. It is worth noting that the geometric phase has no dependence on the common light shift Δ_Σ but only on Δ_Δ . This is as expected as the common light shift does not contribute to qubit dynamics and is immeasurable without comparing to external states.

For a single mode, when $|\eta_{k,1}| = |\eta_{k,2}|$ and $\Delta_{\Delta,1} = \Delta_{\Delta,2}$, performing two loops in phase space (one per pulse), the optimal Δ_Δ to implement U_{CZ} is

$$\Delta_\Delta = \sqrt{\frac{\delta_k^2}{|8\eta_k^2 \cos(\phi_2 - \phi_1)|}}. \quad (7.25)$$

Note that this will not close loops for spectator modes, which will also contribute some geometric phase which may reduce the fidelity of the gate. This is covered in the following description of the gate implementation.

7.7 Gate setup & implementation

We choose the geometry of the Raman lasers with respect to the trap axes such that only the radial modes participate in motional interactions with the Raman lasers. This makes the gate dynamics insensitive to the potentially high axial mode temperature (with the exception of beam amplitude modulation and Kerr cross coupling) which our cooling system is not optimised to reduce. With two ions in the trap with a $2\pi \times 0.5$ MHz axial centre of mass (CoM) mode, the radial modes have frequencies $\omega_{\text{Rocking}}^{LF} = 2\pi \times 1.51$ MHz, $\omega_{\text{CoM}}^{LF} = 2\pi \times 1.6$ MHz, $\omega_{\text{Rocking}}^{HF} = 2\pi \times 1.79$ MHz, $\omega_{\text{CoM}}^{HF} = 2\pi \times 1.86$ MHz. We choose to perform the gate on the lower frequency (LF) rocking mode for two key reasons. Firstly, the rocking modes are not susceptible to global electric field noise, making their heating rates much lower and the motional coherence times much longer. Secondly, when applying force with a negative detuning, the force is furthest from all spectator modes, which will reduce undesired coupling. The 90 kHz detuning from the centre of mass (CoM) mode is still very close, but we reduce the residual spin-motion entanglement by amplitude shaping the SDF [141]. It is possible to push the CoM and Rocking mode splitting further apart by

increasing the axial confinement in our trap, though this currently causes heating of the ions with unknown origin. Additionally, the LF and HF mode each have 45° projection on to the Raman beams. Our 3D micro-fabricated trap has the degrees of freedom to rotate the radial modes, bringing one into full projection of the Raman beams, and the other completely out. This has the benefit of removing another set of spectator modes, a particular challenge when scaling to long chains of ions, though not overly helpful in the two-ion scenario. We do not rotate the modes in this gate implementation, but plan to do so in the future.

7.7.1 Heating rate and motional coherence time

To see the benefit in choosing the rocking mode over the centre of mass mode we compare the heating rate and motional coherence time of each. To measure the heating rate, we cool to the ground state of motion in the target mode, followed by a wait for some duration. After this time, we probe on both the red-sideband and blue-sideband (in separate shots) to measure temperature as in Sec. 7.3.2. By measuring the mean thermal occupation \bar{n} at multiple delay times, we measure the rate at which the mode heats. The results of this measurement is shown in Fig. 7.9. The two CoM modes both exhibit considerable heating of >100 q/s. This limits the motional coherence time to ~ 10 ms. On the other hand, both rocking modes show no measurable heating over a 100 ms period. Fitting the low frequency rocking mode we measure $0.00(4)$ q/s heating and the high frequency rocking mode has $0.10(3)$ q/s heating. This makes the low frequency rocking mode the best choice for gates, since the gate force can be detuned the furthest from all other modes, and the gate mode exhibits the lowest heating rate.

However, this does not strictly mean that the motional coherence is >1 s, since motional dephasing can occur without heating. Variations in the trap electrode voltages can occur on time scales which are slow compared to the motional frequency, and so do not heat. However, this can cause dephasing on the millisecond time scale over which we aim to maintain coherence. To measure this we also perform a




Figure 7.9: Heating rate measurement of the four radial modes on two ions. The CoM modes exhibit high heating rates of over 100 q/s. The LF Rocking mode exhibits no measurable heating over 100 ms and the HF rocking mode measures $\dot{n} = 0.10(3)$ q/s. Measurement beyond this duration is difficult due to heating on the CoM modes leading to de-crystallisation.

Ramsey experiment on the blue-sideband to measure the motional coherence time. Typically a Ramsey experiment on the motion contains carrier pulses to decouple the spin coherence from the motion. Fortunately, working with a clock-qubit with a several second coherence time this is not a limiting factor, and so we work with just the blue-sideband. The result for our mode of interest, the LF Rocking mode, is shown in Fig. 7.10. The coherence time without an echo pulse is 1.9(2) ms, which is 40 \times lower than the 80(20) ms when measured with an echo pulse. This points to a consistent slow phase and frequency drift on the rocking mode which makes shot to shot phases differ, but remain coherent on a per shot basis. This will show up as a small shot-to-shot frequency error in our gate, leaving residual spin-motion




Figure 7.10: Coherence time measurement of the LF Rocking mode. Without an echo the coherence is 1.9(2) ms fit using an exponential decay model. With an echo this increases 40 \times to 80(20) ms. Some oscillations are seen with a 20 ms period indicating coherent 50 Hz noise. The large increase in coherence time with the echo indicates that the noise profile is dominated by a slow phase/frequency drift which can be echoed out within each shot, but causes a variable phase accumulation between shots. For details of the other modes see Appendix A.

entanglement. Such an error can be reduced by using our “discrete” phase tracking mode and employing pulse shaping.

7.7.2 Mode calibration

In the heating rate and coherence time experiments, we implement red/blue sidebands on the metastable clock qubit which are calibrated by changing the relative beam detuning on each application of the sideband and looking for a sinc² resonance. Unfortunately, this method may result in state dependent light shifts which mean the mode is actually a few kHz away from where we measure it to be. To give the best possible measure of the action of the spin-dependent force, we calibrate the mode frequency using the spin-dependent force. To do this, we perform the same sequence of force-echo-force, but this time do not track the phase, only adding π to account for the spin-flip. In cases near zero detuning, the two forces will cancel and the net phase-space displacement is zero, tracing a trajectory similar to the “discrete” case in Fig. 7.7(b). However, as the detuning increases, if full loops are




Figure 7.11: Calibration of the Radial LF Rocking mode frequency using the SDF. The SDF is applied twice, with a spin-echo in between. There is no update in phase of the force after the first pulse. As a result, the spin-motion entanglement is removed only when the force is on resonance. We take the average excitation of the two ions and fit Eq. 7.16. In this instance the motional mode frequency is updated by 290(20) Hz.

not performed, then the two forces do not cancel and the remaining spin-motion entanglement is detected as excited qubit population. We can use this to construct a very sharp resonance with which we calibrate the frequency of the mode of interest.

We illustrate this in Fig. 7.11, fitting Eq. 7.16 where $\phi_0 = 0$, $\phi_1 = \delta t_\pi$, and t_π is the duration of the π -pulse between the two SDFs. This is performed on Doppler cooled ions (no EIT or SBC) to give increased sensitivity to the force. In the fit we use $\bar{n} = 10$ and $\Delta_\Delta = 50$ kHz.

7.7.3 The spectator mode

In our system, the spectator mode (LF CoM) is only 90 kHz away from the target mode. As a result, there is an unavoidable coupling to this mode when we perform the gate. This coupling contributes a geometric phase which must be accounted for in the Rabi frequency, and can also leave spin-motion entanglement, leading to spin dephasing via the motional mode heating or dephasing.

Adjusted Rabi frequency

The spectator mode contributes a geometric phase according to Eq. 7.22. In the case of a CoM spectator on a rocking mode, the ions oscillate predominantly out of phase, with a small in-phase component from the CoM contribution. Consequently, for the CoM mode we have $\eta_1 = \eta_2$, whereas for the rocking mode $\eta_1 = -\eta_2$. This alters the states on which a geometric phase is acquired and can be interpreted as the two modes accumulating relative phases of opposite sign. If the two had equal coupling strength, the effect of the gate would be negligible. As a result the ideal Rabi frequency for a two loop gate with detuning from the rocking mode δ_R and CoM mode δ_{CoM} is,

$$\begin{aligned}\Delta_\Delta &= \sqrt{\frac{\pi}{4 \cos(\phi_2 - \phi_1) (\tilde{\Phi}_R + \tilde{\Phi}_{\text{CoM}})}} \\ \tilde{\Phi}_R &= \frac{2\pi\eta_R^2}{\delta_R^2} \\ \tilde{\Phi}_{\text{CoM}} &= -\frac{\eta_{\text{CoM}}^2}{\delta_{\text{CoM}}^2} \left[\frac{2\pi\delta_{\text{CoM}}}{\delta_R} - \sin\left(\frac{2\pi\delta_{\text{CoM}}}{\delta_R}\right) \right]\end{aligned}\tag{7.26}$$

where η_i^2 is shorthand for $|\eta_i|^2$ and we assume that $\Delta_{1,\Delta} = \Delta_{2,\Delta}$ ⁶.

Loop closure and pulse shaping

The ideal $\tilde{\Phi}$ for the CoM mode indicates the second change made when considering a spectator mode, closing both sets of loops in phase space. If the gate duration $\tau = 2\pi/\delta_R$ does not satisfy $2\pi\delta_{\text{CoM}}/\delta_R = n\pi$, then the spin and motion will remain entangled at the end of the gate. This can be mitigated in two ways. Firstly, choosing a detuning and therefore gate duration such that loops close in both modes, we can remove the issue altogether. Secondly, we can shape the amplitude of the SDF to ramp in and out of the pulse. This reduces the off-resonant coupling to spectator modes, leaving the position and momentum displacement of these modes at a lower amplitude. We choose a ramp with shape \sin^2 at the start and end of each SDF, and a ramp duration $t_{\text{ramp}} > 1/\delta_{\text{CoM}}$ to suppress coupling to the CoM mode.

⁶The difference in sign between $\tilde{\Phi}_R$ and $\tilde{\Phi}_{\text{CoM}}$ is due to $\eta_{1,R} = -\eta_{2,R}$ and $\eta_{1,\text{CoM}} = \eta_{2,\text{CoM}}$.

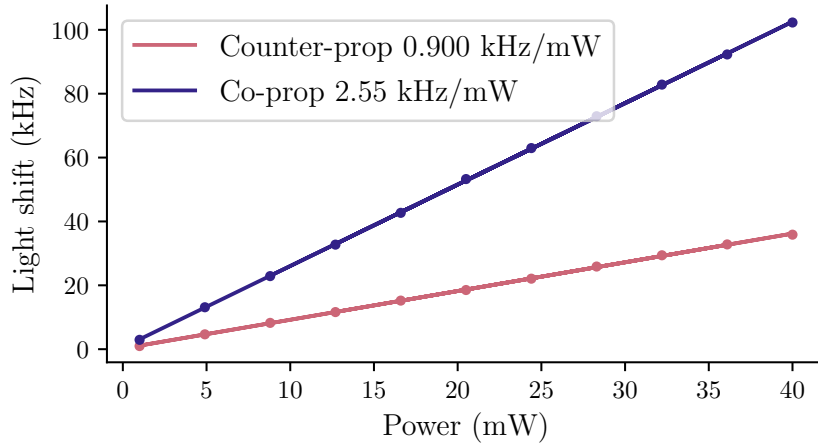


Figure 7.12: Average light shift vs optical power on two optical qubits when applying either the counter-propagating beam or co-propagating beam. The co-propagating beam has a slightly tighter beam waist, and therefore for a given optical power induces a greater light shift. We use these measurements to balance the light shift from each beam in the OTDF gate.

7.7.4 Balancing light shifts

We perform RPE (see Sec. 4.3.2) applying an increasing intensity on the co-propagating and counter-propagating beams to measure the light shift per mW of power. The result of this is shown in Fig. 7.12. Due to the slightly tighter waist, the co-propagating beam induces $2.83\times$ the light shift of the counter-propagating beam. To balance the light shift from each beam, as is desired from Eq. 7.8, we apply 2.83 mW of the counter-propagating beam for every 1 mW of co-propagating to the ions. In Fig. 7.12, we take the mean light shift across the two ions. Ideally this light shift is equal across the ions, however the echo pulse in the gate sequence also removes small imbalances.

7.8 Results and sources of error

7.8.1 Results

We implement the OTDF gate on the LF rocking mode with a detuning $\delta = -2\pi\times 10$ kHz, performing two loops. Each pulse is ramped in and out over $30\ \mu\text{s}$, to reduce coupling to the LF CoM mode from which the gate force is detuned by

$-2\pi \times 90$ kHz. The duration of each pulse is increased by $30 \mu\text{s}$ from the ideal to compensate for the lost pulse area. As a result, we observe the loop closure at $260 \mu\text{s}$. For pulses with a duration less than $2 \times 30 \mu\text{s}$, we reduce the ramp duration to half of the pulse length, allowing for points with duration shorter than $60 \mu\text{s}$ to be taken. We plot the result in Fig. 7.13(a).

Taking the populations at the ideal shaped gate duration of $260 \mu\text{s}$, we measure $99.2(6)\%$ of the population to be in the $|\downarrow\downarrow\rangle$ and $|\uparrow\uparrow\rangle$ states, leaving a $0.8(6)\%$ error from population in $|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle$. We additionally perform a parity scan analysing the result of an additional $\pi/2$ -pulse after the gate with varying phase. The parity scan in Fig. 7.13(b) shows a contrast of $98.8(1.2)\%$. Combining these two results we estimate that the current gate fidelity is $99.0(7)\%$. This is well below the state of the art and the fundamental limit from qubit decay and Raman scattering ($\sim 5 \times 10^{-5}$), indicating some large sources of technical error. We additionally scan the power in the co-propagating beam to verify that the correct phase is acquired. The counter-propagating power is scaled accordingly using $2.83\times$ that of the co-propagating beam. The result is shown in Fig. 7.14. We use a co-propagating power of 36.2 mW to optimise the geometric phase.

In the parity scan in Fig. 7.13(b), we also observe a phase offset which amounts to an over-rotation in single qubit phase, a potential indicator of our source of error. We will now go through the potential sources of error, giving their likely contribution based upon auxiliary measurements such as heating rates.

7.8.2 Sources of error

From the estimate of the current gate fidelity, we can see that there is much room for improvement and some large sources of error to grapple with. At the time of writing, no single standout error source has become clear, so in this section we consider the possibility of different error sources, ruling out those which are inconsistent with the data we have collected.

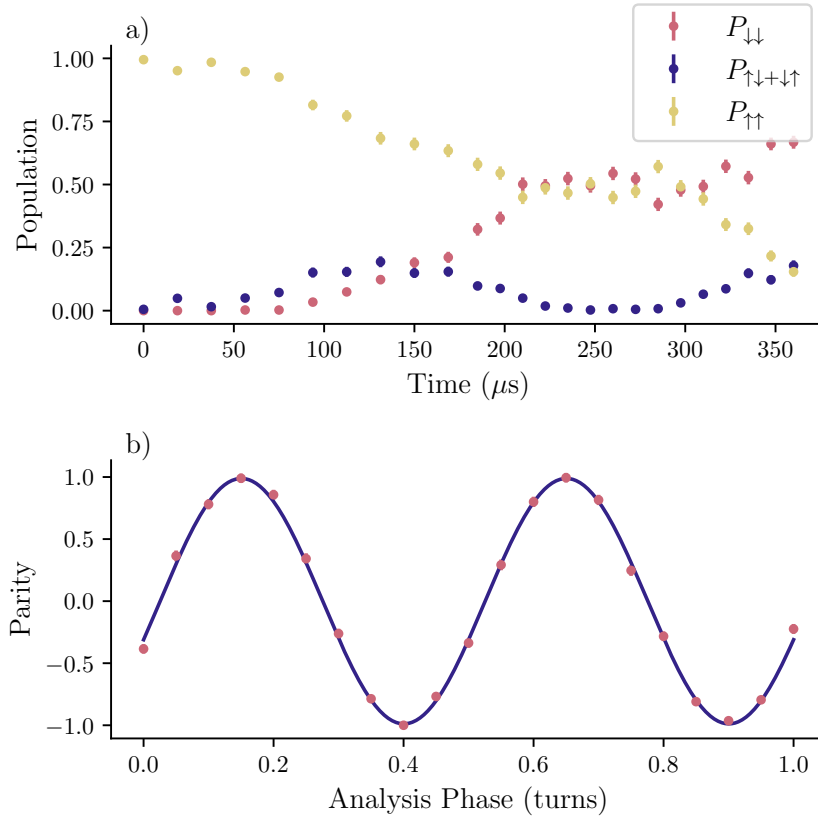


Figure 7.13: OTDF gate results. These results are used for the fidelity estimate of our gate implementation. The gate is performed with a $-2\pi \times 10$ kHz detuning and $30 \mu\text{s}$ ramps. (a) Populations vs time, the phase space loops close at $260 \mu\text{s}$ which is expected due to the $30 \mu\text{s}$ ramps. $0.8(6)\%$ of the population is in the state $|\downarrow\uparrow\rangle$ and $|\uparrow\downarrow\rangle$. (b) Parity scan with total gate duration of $260 \mu\text{s}$. The parity scan has a contrast of $98.8(1.2)\%$. Combining these results we measure a fidelity of $99.0(7)\%$.

In each of the following cases, we consider the error as the deviation from an ideal gate. For most cases errors will add in quadrature, and therefore independently, although it is possible errors could compound, as is the case for heating and Kerr-cross coupling. Since we are looking for some large sources of error, many approximations made for small errors will not be valid, as such we take the approach of directly simulating the Hamiltonian in Eq. 7.12. When computing an error, we use the same metric in simulation as in experiment, taking the mean of the error due to population in $|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle$ and the parity contrast of the resulting state. While for certain errors it is possible that we have created a high-fidelity Bell state, if it is not a known and




Figure 7.14: Qubit populations vs power under the OTDF gate implementation. This is used to calibrate the power used for the OTDF gate. We work where the populations cross at 36.2 mW. This data was taken with a 260 μs gate duration, with 30 μs ramps inclusive and a $-2\pi \times 10$ kHz detuning from the lower radial rocking mode.

desired state, this is an error. In one case we do compare to the *fully entangled fraction*, which calculates the distance to the nearest maximally entangled state. This illustrates how a coherent error can be compensated with single qubit rotations.

The ideal gate

We begin by considering the technical limit of our equipment as things stand. If we were to perfectly implement the OTDF gate, but have limits to the qubit and motional coherence time, what would the error be? The dynamics of the ideal gate are shown in Fig. 7.15. Decoherence is implemented through a series of Lindblad operators: $\hat{\sigma}_{z,j}$ for decoherence on qubit j , $\hat{a}_k + \hat{a}_k^\dagger$ captures heating on mode k , and $\hat{N}_k = \hat{a}_k^\dagger \hat{a}_k$ expresses motional dephasing [84]. Note that in Fig. 7.15 the x-axis is not the time within a single evolution of the Hamiltonian, but rather the duration of the SDF for a full evolution with two pulses, making it directly comparable to the results in Fig. 7.13. The most notable difference is the $P_{\uparrow\downarrow+\downarrow\uparrow}$ population, which is not excited. This is expected if the “discrete” phase tracking is correctly implemented, similar to Fig. 7.8. The parameters used in simulating the ideal gate are as follows. We implement a qubit with coherence time 10s with a -10 kHz




Figure 7.15: Simulation of a perfectly calibrated gate in our system with no added noise from the 532 nm lasers. This simulation uses measured parameters in our system such as the heating rate on the rocking (CoM) mode with a heating rate of 0.1 q/s (110 q/s) and motional coherence time of 80 ms (2 ms). We assume a spin coherence time of 200 ms with the spin-echo pulse and the gate detuning is -10 kHz with a $30 \mu s \sin^2$ ramp.

detuning to a 1.51 MHz rocking mode with heating rate 0.1 q/s and coherence time 80 ms. This qubit coherence time is far beyond the 51(4) ms we measured in Sec. 4.4.2. This is because the Lindblad operators we employ effectively assume a white-noise model, however, our coherence time was estimated with a pink-noise model. We convert between the two by taking a coherence time that would give equivalent Ramsey contrast at the gate duration. This comes with the assumption that any change in dynamics from the switch in noise profiles is a small, second-order effect. Consequently, the noise model is only valid for a single gate, since in the pink-noise model the qubit coherence would rapidly decay over repeated applications of the OTDF gate, consistent with the 51(4) ms experimental coherence time. There is a spectator mode with frequency 1.591 MHz, a heating rate of 110 q/s and a coherence time of 2 ms. There is a $30 \mu s \sin^2$ ramp in and out of each SDF pulse. The single qubit operations are taken to be ideal, a fair assumption given the 10^{-4} level errors we see in RBM compounded with the XYX pulse scheme that eliminates first order Rabi frequency errors in the single-qubit 1762 nm pulses. The error from this ideal gate is 1.8×10^{-3} , meaning we have some large sources of additional error.




Figure 7.16: Error of the OTDF gate vs increasing detuning error. With zero detuning error the gate is the same as shown in Fig. 7.15 with an error of 4.3×10^{-3} . We anticipate a maximum detuning offset of 200 Hz which translates to an additional error of $\sim 3 \times 10^{-3}$.

Detuning error

There are two possible ways a detuning error could be present in our system. The first is a straight forward and absolute detuning error. This would change the ideal gate duration by some amount, which if not accounted for could leave spin-motion entanglement. However, such an error can be calibrated out by simply adjusting the gate duration. This is something we would observe in Fig. 7.13, but the loop closure is at the expected $260 \mu\text{s}$ ($2 \times [100 + 30] \mu\text{s}$ to include the \sin^2 ramps). Alternatively, the gate timing and detuning could be well compensated but the exact detuning from the gate mode not known. In this instance, the phase compensation between the two SDF pulses will be wrong, resulting in trajectories which do not always close loops. If the detuning and gate duration are well matched, this does not contribute an error, but also does little to mitigate timing/detuning errors.

In Fig. 7.16 we plot the error when there is an absolute detuning error, not compensated in time. We find that to reach a 1% level error would require a detuning error of ~ 500 Hz. This is possible, but unlikely, for two reasons. The first is the ability to dynamically compensate for such an error by adjusting the gate duration. Additionally, we measure the mode drift for both the LF Rocking and LF CoM mode



Figure 7.17: (a) Lower frequency radial mode stability over a 13 hour period. The Rocking and CoM modes show a high degree of correlation implying drift is primarily due to RF amplitude instability. Data points are taken every three minutes and measured using the SDF. (b) Histogram of the frequency change between each measurement in the Radial Rocking mode. We use this data to estimate that a ~ 200 Hz frequency offset is possible over the period the data in Fig. 7.13 was taken.

using the SDF over a 13 hour period. The result of this is shown in Fig. 7.17(a). The two modes primarily move in step with each other, implying the majority of drift is due to RF amplitude noise. The modes are measured every 3 minutes and over the 13 hour duration there is a total change of ~ 2 kHz in both modes. The absolute frequency change between measurements is plotted in Fig. 7.17(b) as a histogram, revealing that the mean change between measurements is 120(90) Hz. While there are some periods of directional frequency change, for example at the very start of the measurement process, this is not always the case. Even in the most extreme case, there is a 1 kHz change over a 30 minute period. Given that the data collection

process takes only a few minutes, we can assume that the frequency is good to within a few hundred Hertz, thus introducing an error of at *most* 1%.

Rabi frequency error

The only other parameter available to us which may be miscalibrated is the Rabi frequency of the SDF (Δ_{Δ}). As with the detuning error, there are two ways in which this can go wrong. There could be a consistent offset in the amplitude relative to the desired value. This would cause a reduction in the parity contrast as the states $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ do not have equal populations, but the timing of the gate will be correct, resulting in loop closure and minimal $|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle$ population. The error for a given fractional Rabi frequency offset is shown in Fig. 7.18. Given the manual calibration method for the Rabi frequency, using Fig. 7.14, an offset in Rabi frequency of ± 1 mW is possible. This would translate to a $\sim 3\%$ fractional Rabi frequency offset which we can see from Fig. 7.18 gives a 0.3% gate error. Additionally, we characterise the shot-to-shot deviation in the beam intensity of both the co-propagating and counter-propagating beams in Fig. 7.19(a). This is done by measuring the beam intensities on a photodiode when performing the SDF sequence. The co-prop shows an increasing drift in power with a $\pm 1\%$ maximum deviation in power, with standard deviation of 0.5%. The counter-prop is more randomised, though shows a larger maximum range of 1.7% and a standard deviation of 0.9%. These fluctuations are sufficiently small that we can expect them to contribute an error below 0.1% and thus are not a leading cause of error.

An alternative type of error arises from an imbalance in Rabi frequency between the two SDF pulses in the gate sequence. This imbalance produces a single-qubit σ_z rotation, which affects both qubit populations and the parity contrast. The resulting error is shown in Fig. 7.20, where we assume the first pulse has the desired Rabi frequency while the second deviates by some fraction. In Fig. 7.20, we plot both the error from population and parity and the *fully entangled fraction*. In this case, the




Figure 7.18: Gate error vs a constant offset error in Rabi frequency. We anticipate that at an absolute maximum the Rabi frequency could be miscalibrated by 5%. A more realistic estimate combining the relatively balanced $P_{\downarrow\downarrow}$ and $P_{\uparrow\uparrow}$, accuracy of our calibration and data in Fig. 7.19, is an offset of 1–2%. A 2% error in Rabi frequency translates to a 0.2% gate error.

fully entangled fraction is the more appropriate metric, since the coherent error can be compensated by adjusting the closing phase of the Ramsey-like sequence.

It is worth noting that we simulate only the Hamiltonian in Eq. 7.12, which neglects the light shift from the arithmetic mean of the two beams in Eq. 7.8. For small Rabi frequency asymmetries, reintroducing this light shift is approximately equivalent to doubling the light-shift asymmetry. Accordingly, when comparing the intensity asymmetry in Fig. 7.19(b), we consider the sum of the counter-propagating and co-propagating asymmetries rather than their mean. The counter-propagating beam contributes the larger error, with a fractional pulse asymmetry of 0.2(2)%. By contrast, the co-propagating beam is significantly more stable, with a systematic offset of $-0.07(4)\%$ between the first and second pulse. Taking the sum of these contributions, we expect an overall Rabi frequency imbalance of $\sim 0.1\%$. Comparison with the simulation in Fig. 7.20 shows that such an imbalance can quickly generate substantial error—up to 1.5% if uncompensated. However, Fig. 7.20 also plots the fidelity as determined from the *fully entangled fraction*, which characterises entanglement by measuring the distance to the nearest Bell state. Importantly,




Figure 7.19: Measurement of the counter- and co-propagating beam amplitudes across 50 shots measured on a photodiode. (a) The average pulse area across both arms of the SDF. We observe up to 1% level fluctuations in the counter-propagating beam and a slow drift in the co-propagating beam. This type of offset corresponds to the error in Fig. 7.18. (b) Pulse area asymmetry between the two arms of the SDF. We see a consistent offset of 0.2(2)% in the counter-propagating beam. Again, the co-propagating beam is much more stable, and the second pulse slightly undershoots the first by just 0.07(4)% on average. This corresponds to a 1.9% error in Fig. 7.20.

the single-beam light shift does not reduce this fidelity, since there remains a basis in which a high-fidelity Bell state can be found. This illustrates that purely coherent errors can be corrected with single-qubit rotations. In Fig. 7.21 we plot the average phase acquired across the two qubits after the gate sequence, with the lasers far detuned to avoid motional coupling. From this, we measure a phase shift of 0.22(2) rad, which we compensate by incorporating it into the closing $\pi/2$ pulse of the Ramsey interferometer. This correction removes the coherent error from Rabi-frequency imbalance between the two SDF pulses, though it does not




Figure 7.20: Gate error vs Rabi frequency asymmetry in the two SDF’s which make up the OTDF gate. We see that the error when calculated using population and parity measurements rapidly grows with asymmetry. Just a 0.2% asymmetry is sufficient to cause a $\sim 5\%$ error. From data in Fig. 7.19, we anticipate a 0.13% asymmetry in Rabi frequency. This is the greatest source of error we have with a contribution of 1.9%. Note that the error in the *fully entangled fraction* does not increase with this asymmetry. This is because the error is coherent and equivalent to a single qubit rotation in each ion. As such, a basis can still be found in which we have a maximally entangled state.

address shot-to-shot variations. We instead model these as a form of spin dephasing, discussed later in this section.

Motional coherence

In the ideal gate simulation, we have assumed an 80 ms motional coherence time in the LF rocking mode. This is likely an optimistic estimate given the data in Fig. 7.10. The 20 ms revivals observed with an echo pulse suggest the presence of coherent modulation. Since the gate samples only the first 200 μs of this noise rather than averaging over several oscillations, a more realistic estimate is obtained by fitting just the first two points, yielding a coherence time of 20 ms. In Fig. 7.22, we simulate the gate error for varying motional coherence times. For a 20 ms coherence time, the error contribution is 0.3% in total, representing a small but avoidable source of error.




Figure 7.21: Ramsey phase measurement with the OTDF gate force far detuned. We measure a phase offset of $0.22(2)$ rad from the gate sequence implying there is a consistent amplitude offset between the two pulses. We compensate this single-qubit rotation in the final gate implementation in Fig. 7.13.




Figure 7.22: Gate error vs coherence time of the LF Rocking mode. In the best case, the coherence time is 80 ms where the error from dephasing and heating is $\sim 2 \times 10^{-3}$. From the 50 Hz modulation visible in Fig. 7.10, a more realistic estimate for the purposes of the gate is 20 ms which introduces a further 0.1% error.




Figure 7.23: Coherence time measurement when implementing the OTDF sequence. The 532 nm lasers are far detuned allowing us to measure just spin dephasing. Under a pink-noise model we measure a coherence time of 3.6(4) ms. This is equivalent to a 50 ms white-noise coherence time which translates to a 1% gate error.

Qubit coherence

We have measured our qubit coherence time in Section 4.4.2 to be 51(4) ms under a pink-noise spectrum⁷. Over the first 200 μs of coherence, this is equivalent to a 10 s coherence time under a white-noise model, which is the coherence time we have used in our Lindblad implementation of dephasing. However, we have seen that there are shot-to-shot variations in the amplitude of the 532 nm lasers driving the gate, which induce qubit dephasing equivalent to decoherence. To characterise this we measure the optical qubit coherence time when implementing the OTDF gate, but with the lasers far detuned to avoid motional coupling. The result of this is shown in Fig. 7.23 where we measure a significantly reduced 3.6(4) ms coherence time again under a pink-noise model. This is equivalent to a white-noise coherence time of 50 ms. In Fig. 7.24 we plot the gate error vs qubit coherence time and find that at with a 50 ms white-noise coherence time (under white-noise), the error from dephasing is 1%, accounting almost entirely for the error measured experimentally. It is worth reiterating that this conversion between pink-noise and white-noise is only applicable under a single implementation of the gate. Further iterations within the

⁷This was measured without a spin-echo, so is also a pessimistic estimate.




Figure 7.24: Gate error vs spin coherence time. We know the qubit coherence to be 51(4) ms under a pink-noise model. This is equivalent to 10 s in a white-noise model over the gate duration. However, amplitude noise on the 532 nm lasers introduce further spin-dephasing as measured in Fig. 7.23. A more realistic estimate is a 50 ms coherence time under a white-noise model. At this reduced coherence time, simulation reports a 1% error in the OTDF gate.

same experimental sequence will sample a greater range of the pink-noise spectrum, greatly increasing the spin-dephasing over many gate sequences. This simulation also illustrates how the hyperfine structure and choice of qubit massively improves the performance of our system. When working with magnetically sensitive transitions, we see coherence times of at best 1 ms, but often 200 μ s. This would bring about errors at the 10 % level which would require a significant system redesign to eliminate.

Fundamental errors

There are two fundamental errors associated with the OTDF gate: qubit decay, and spontaneous photon scattering [4]. Firstly, the metastable level has a 30 s lifetime, which will cause the $D_{5/2}$ qubit population to decay to the $S_{1/2}$ level. Given the 260 μ s gate duration, this will introduce an error of 8×10^{-6} , well below our technical limitations and the present state of the art [138] for this type of gate. The photon

scattering rate between states $|i\rangle$ and $|f\rangle$ is given by [4],

$$\Gamma^{i \rightarrow f} = \frac{3\pi c^2 I}{2\hbar} (\omega_{\text{Laser}} - \omega_{fi})^3 \sum_{s=\{0,\pm 1\}} \tilde{\Gamma}_s^{i \rightarrow f}$$

$$\tilde{\Gamma}_s^{i \rightarrow f} = \left| \sum_k (2J_k + 1) \sqrt{\frac{A_{J_k J_i} A_{J_k J_f}}{\omega_{ki}^3 \omega_{kf}^3}} \left(\frac{\langle f | T_s^1 | k \rangle \langle k | \hat{\epsilon} \cdot T | i \rangle}{\omega_{ki} - \omega_{\text{Laser}}} + \frac{\langle f | \hat{\epsilon} \cdot T | k \rangle \langle k | T_s^1 | i \rangle}{\omega_{kf} + \omega_{\text{Laser}}} \right) \right|^2, \quad (7.27)$$

where the electric dipole operators have been replaced with unitless spherical tensors T^1 and the coefficients $A_{J_k J_i}$ and $A_{J_k J_f}$ are the spontaneous emission rates for states $|k\rangle \rightarrow |i\rangle$ and $|k\rangle \rightarrow |f\rangle$ with orbital angular momentum J_x . Evaluating Eq. 7.27 over all possible scattering paths, we calculate a fundamental error from photon scattering of 4.5×10^{-5} . This means we have a total fundamental error of 5.3×10^{-5} .

Combining errors

From the analysis above, it is clear we have several contributors to the technical error presently limiting our gate fidelity. To assess the contribution of each factor, we simulate the OTDF gate, introducing each type of error, starting with the smallest contributor adding each error source on top. We use the values inferred from the characterisation above which are summarised in Table 7.1 along with their estimated error contribution. While this combination of errors is sufficient to account for the

Error Source (Value)	Error Contribution (%)
Spin dephasing (50 ms T_2)	1.0
Detuning offset (200 Hz)	0.3
Additional motional dephasing (20 ms T_2)	0.1
Heating and motional dephasing (80 ms T_2)	0.2
Rabi freq offset ($1.02 \times \Delta_{\text{Ideal}}$)	0.1
Fundamental errors	0.005
Total	1.7
Measured error	1.0(7)

Table 7.1: Error sources in the OTDF gate along with the value which is used to simulate the error in the gate. The leading source of error is spin dephasing, which can also be viewed as a product of instability in these SDF pulses. This accounts for a 1% error. With improvements to mode stability and beam amplitude stability it should be possible to achieve errors around $\sim 2 \times 10^{-3}$ in our system.

measured error of 1.0(7)% (within error bars), it does not explain the trajectory taken by the qubit populations throughout the gate. In particular, $P_{\uparrow\downarrow\downarrow\uparrow}$ reaches a maximum population of ~ 0.25 in experiment, whereas in simulation it remains close to zero. At present, the only plausible explanation is a phase glitch in the pulse sequence, which would prevent the phases from closing loops in phase space. Such glitches have previously been observed in the Sinara HW phaser board [143], although none have yet been identified in our system. The only other possible explanation is a large detuning error; however, this would require a detuning on the order of several kilohertz, which we are confident is not present.

7.9 Summary

We have implemented the OTDF gate in an optical qubit with hyperfine structure for the first time. Such qubits can exhibit orders-of-magnitude longer coherence times without the need for magnetic-field noise cancellation or shielding. As such, the implementation of the OTDF gate, albeit with some technical errors has been successful in providing a path towards high-fidelity two-qubit gates in optical qubits with hyperfine structure.

In addition, for our experiment specifically, implementing a two-ion gate in the σ_Z basis enables inter-ion interactions in systems with mixed qubit types. This is because no phase coherence is required between the gate driving force and the qubit, leaving the technical difficulty of phase tracking to single-qubit operations only. Without this, we would need to implement two-qubit gates with the 1762 nm laser, or an MS gate with the 532 nm laser in a metastable qubit. Given the small Lamb–Dicke parameter of 0.012 for the 1762 nm laser, the former would not be feasible. As a result, the OTDF gate is an excellent fit for our optical-centric mixed-qubit-type system.

If we were to adopt a system centred around metastable qubits, then an MS gate implemented with the 532 nm beams would be a natural choice, as described in Ref. [47]. In this approach, phase coherence during mapping to an optical qubit is not required, allowing inter-ion gates to operate effectively in the σ_x basis (see Sec. 4.5).

In conclusion, the implementation of the OTDF gate – albeit with some technical errors – has been successful in providing a path towards high-fidelity two-qubit gates in optical qubits with hyperfine structure.

*A complimentary escape hatch will open in three...
two... one.*

— GLaDOS, Portal

8

Outlook

Creating a trapped ion system with mixed qubit types is a promising route to realising qubit specific and mid-circuit measurement in large scale quantum computers. This form of qubit measurement is critical to implementing quantum error correction which is in turn key to realising fault tolerant quantum computation. We have explored how optical and metastable qubits can be combined to implement systems with mixed qubit types.

By carefully selecting the states used to form both optical and metastable qubits, we are able to construct mixed-qubit-type schemes that combine long coherence times on the order of 100 ms–1 s with low-error interconversion between qubit types. When working in an optical centric system, pairing with a metastable qubit which can be mapped to coherently is vital for implementing qubit specific measurement. As a result, we have found this places a further restriction on our choice of optical and metastable qubit combination. However, the coherent conversion is feasible. Additionally, we are able to achieve 10^{-4} level errors in both qubit types, suitable for NISQ scale development and research purposes.

These qubits can be combined with our scalable addressing system to address ions in a chain. Our addressing system is compatible with the 532 nm light which can

be used to drive metastable qubit transitions, the OTDF gate in Chapter 7, and light shift optical qubit transitions in/out of resonance. The system demonstrates cross talk which is below 10^{-3} across all ion zones and as low as 10^{-5} in some cases, enabling high-fidelity operations on all qubits. This allows us to construct a universal gate set in a multi-qubit system, enabling small scale quantum computation in the near future.

We have developed a protocol to enable the high-fidelity state preparation and measurement (SPAM) of optical, ground and metastable qubits. This SPAM protocol can integrate directly with our existing architecture, and enables the mitigation of errors from qubit loss. Typically metastable decay or Raman scattering in trapped ions. Within the protocol we have explored exactly how the SPAM error is attributed, mitigated and ultimately limited. The only limiting factor after post-selection or repeat-until-success is the speed with which measurements can be performed and the result acted upon. The protocol allows us to eliminate some of the fundamental errors in our single-qubit operations and the two-qubit OTDF gate. However, the protocol may bias the observables we measure during a tomographic measurement. Fortunately, there are several ways to mitigate this, either by alternating the definition of $|\uparrow\rangle$ and $|\downarrow\rangle$, applying a correction to the statistics, or reducing the fidelity of our mapping pulses to match the worst case scenario. The protocol has demonstrated record low SPAM errors in all three qubit types, with the minimum error being $5(4)\times 10^{-6}$. When combined with erasure conversion [99, 100], the diagnostic measurements in our protocol can be used to perform in-sequence qubit loss correction, paving a way to scalable quantum computing, even with finite qubit lifetimes.

The OTDF gate presented in Chapter 7 enables us to directly integrate our optical qubit with the existing 532 nm laser system. While the gate is presently limited by technical error, performing this gate in $^{137}\text{Ba}^+$ will enable a high-fidelity gate in a system with long coherence times and metastable $D_{5/2}$ state lifetimes. Additionally, due to the lack of coherence required between driving field and qubit, this gate can be implemented in a variety of systems with minimal changes. For instance, it should be possible to extend this to qudit systems without much added complexity in phase

tracking, for example, this implementation of a 25-level qudit in $^{137}\text{Ba}^+$ [144]. Looking to the future, our experiment is currently shifting towards the implementation of optical based poly-qubits in $^{137}\text{Ba}^+$, and this gate mechanism is compatible, reducing the technical details of implementing two-ion gates in poly-qubits.

The use of mixed qubit types in quantum computing has generated considerable excitement, not only in trapped-ion platforms but also in neutral-atom architectures [100]. The ability to coherently interconvert between different qubit types opens up new strategies for scalable control, mid-circuit measurement, and error correction that push the limitations of single-species systems. This thesis contributes an investigation of how *omg* schemes can be realised in $^{137}\text{Ba}^+$, representing a step toward harnessing these capabilities in future large-scale quantum processors.

Building on the foundations laid in this thesis, our aim is to integrate the key components – single-qubit gates, two-qubit gates, single-ion addressing, and high-fidelity state preparation and measurement – into a NISQ device capable of demonstrating novel functionalities in quantum computing. The OTDF interaction presented in Chapter 7 is directly compatible with qudits and poly-qubits, enabling the exploration of new gate schemes and investigations into resource trade-offs in poly-qubit-based error-correcting codes. Taken together, the techniques outlined in this thesis, along with our planned future developments, provide a pathway toward scaling trapped-ion systems into practical quantum computers.

Appendices

A

Motional coherence times

Here we provide data illustrating the motional coherence time of all four radial modes on two ions. Data is shown in Fig. A.1 illustrating the motional coherence time both with and without an echo pulse. We observe that the CoM modes have short ~ 5 ms coherence times that are not extended with an echo pulse. This is likely due to them being limited by their high heating rates of ~ 100 q/s. On the other hand, the rocking modes exhibit even shorter ~ 2 ms coherence times without an echo pulse, but this is greatly extended to ~ 70 ms with an echo. This large increase implies a slow frequency drift in the rocking modes that cause random shot-to-shot phase differences that are easily echoed out. Another important feature in the LF and HF rocking mode data is the 20 ms oscillations. This implies there may be a 50 Hz noise spike consistent with noise from the mains electricity. Given the sparsity of the data, further characterisation is required to draw a concrete conclusion on this.

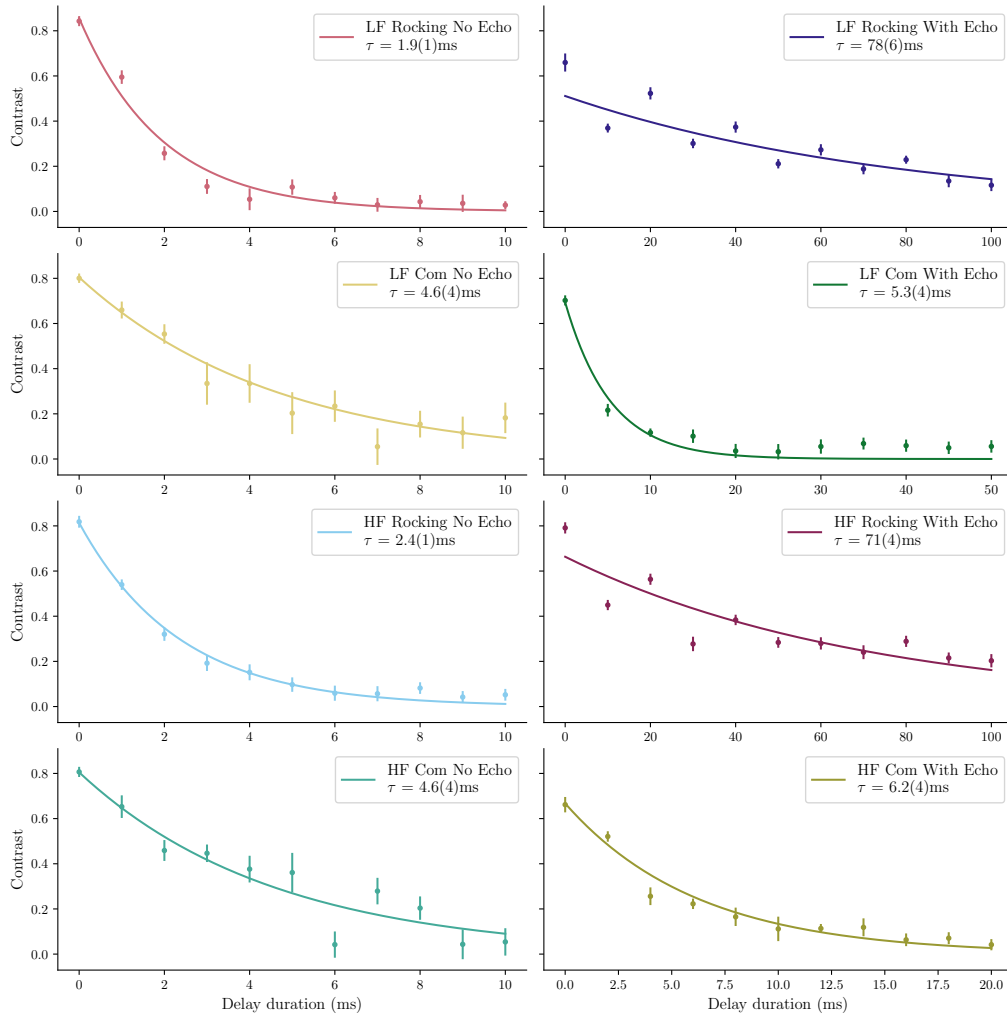


Figure A.1: Motional coherence time measurement for all four radial modes in a two ion chain. Measurements are taken both with and without an echo pulse to highlight the difference between the fast dephasing/heating on the CoM modes and a slower phase drift which is exhibited on the rocking modes.

B

Optical clock qubits

In Chapter 4, we provided a short summary of low field optical clock qubits in $^{137}\text{Ba}^+$. In Tab. B.1 we extend this to include all optical clock qubits we have found through simulation using Ref. [53] between 0.5 G and 125 G. We neglect clock qubits at a field of less than 0.5 G as these low magnetic fields will provide a weak projection that is difficult to work with experimentally. Furthermore, these very low field clock qubits have very high second order sensitivities, implying their clock like properties only hold over a small range of fields. All optical clock qubits must be drivable via a quadrupole transition, and so we limit the search to pairs of states with a difference in angular momentum projection $|\Delta m_F| \leq 2$.

$S_{1/2}$ state	$D_{5/2}$ state	Field insensitive point (G)	d^2f/dB^2 (kHz/G ²)	Matrix element Cm/h
$ 1, 0\rangle$	$ 3, -1\rangle$	0.7633	69.8	1.39
$ 2, 0\rangle$	$ 3, -1\rangle$	0.7742	64.7	0.574
$ 1, 0\rangle$	$ 2, 0\rangle$	1.3471	0.2	0.139
$ 2, 0\rangle$	$ 3, -1\rangle$	1.6354	26.4	0.794
$ 1, 0\rangle$	$ 3, -1\rangle$	1.6970	26.8	1.31
$ 2, 0\rangle$	$ 2, 0\rangle$	4.1795	2.1	1.09
$ 2, -2\rangle$	$ 4, 0\rangle$	4.6159	35.7	0.269
$ 2, -2\rangle$	$ 3, -2\rangle$	8.5406	35.2	2.52
$ 1, -1\rangle$	$ 1, 0\rangle$	8.8245	72.6	2.43
$ 2, 1\rangle$	$ 1, 0\rangle$	8.9014	71.7	0.454
$ 2, 1\rangle$	$ 3, 0\rangle$	10.7028	56.2	0.124
$ 1, -1\rangle$	$ 3, 0\rangle$	10.8269	55.6	0.239
$ 2, -1\rangle$	$ 3, -1\rangle$	14.7440	73.7	2.81
$ 1, 1\rangle$	$ 3, -1\rangle$	14.9046	73.5	1.10
$ 2, -2\rangle$	$ 1, -1\rangle$	17.0624	103.1	1.25
$ 2, 2\rangle$	$ 1, 0\rangle$	20.8627	44.8	0.319
$ 2, 0\rangle$	$ 3, 0\rangle$	22.8316	52.8	2.64
$ 2, -2\rangle$	$ 3, -1\rangle$	23.1050	90.3	0.817
$ 1, 0\rangle$	$ 3, 0\rangle$	23.2659	51.0	0.424
$ 2, 2\rangle$	$ 3, 1\rangle$	24.1870	27.8	0.624
$ 2, -2\rangle$	$ 2, -1\rangle$	24.6967	34.2	1.66
$ 2, -2\rangle$	$ 2, -2\rangle$	29.8295	37.9	2.56
$ 1, -1\rangle$	$ 2, 0\rangle$	35.5444	15.9	0.800
$ 2, 0\rangle$	$ 1, -1\rangle$	36.6492	40.3	0.116
$ 2, 1\rangle$	$ 2, 0\rangle$	37.3159	13.2	0.397
$ 2, -1\rangle$	$ 4, 1\rangle$	44.6498	4.6	0.156
$ 2, -1\rangle$	$ 3, 0\rangle$	46.8961	11.6	1.02

Table B.1: Transition properties of optical clock qubits in $^{137}\text{Ba}^+$ between 0.5 G and 125 G. Atomic states are written as $|F = i, m_F = j\rangle = |i, j\rangle$.

List of Figures

2.1	Level diagram of a typical trapped ion system implementing ground, optical, and metastable qubits.	11
3.1	(a) Diagram of beam geometry. (b) Neutral barium photoionisation levels. (c) $^{137}\text{Ba}^+$ atomic level diagram.	21
3.2	Neutral fluorescence resonance measurement.	22
3.3	Resonances with the 493 nm laser for (a) “idle” Doppler cooling and (b) “cold” Doppler cooling.	24
3.4	EIT cooling level diagram.	25
3.5	Fluorescence measurement while scanning the EIT probe detuning.	26
3.6	Temperature measurements vs EIT cooling time.	27
3.7	Diagram of the lasers used for state preparation.	28
3.8	State preparation success vs optical pumping duration.	28
4.1	Calculation of the scaling factor $\langle n' e^{-i\eta(a^\dagger+a)} n \rangle$ for the carrier and first sideband with (a) the 532 nm counter-propagating lasers (b) the 1762 nm laser.	38
4.2	An example three level system driven by a Raman transition	40
4.3	Diagram of the 1762 nm optical setup between laser amplifier and trap used for switching and frequency control.	43
4.4	Geometry polarisation factor for quadrupole transitions.	44
4.5	Level diagram with the 532 nm Raman lasers.	45
4.6	Simulation of the matrix element for three Raman transitions.	46
4.7	Diagram of the AOMs used to switch and frequency control the counter-propagating, co-propagating and metastable beams.	47
4.8	A visualisation of the RPE algorithm.	51
4.9	Simulation of the RPE protocol with added noise.	53
4.10	An example RPE frequency calibration with the metastable clock qubit.	54
4.11	An Example Rabi RPE calibration with the optical qubit.	55
4.12	Randomised benchmarking of the metastable clock qubit with optimal parameters.	58
4.13	The error from randomised benchmarking of the metastable qubit over a period of six hours.	58
4.14	Error per Clifford from RBM vs $\pi/2$ -pulse duration in the metastable qubit.	59
4.15	Error per Clifford from RBM vs detuning with the metastable qubit.	61
4.16	Metastable clock qubit coherence time measurement.	61
4.17	Optimal result of RBM with the optical qubit.	62
4.18	Error per Clifford from repeated RBM runs with the optical qubit over a 2 hour period.	63
4.19	Error per Clifford vs $\pi/2$ -pulse duration in the optical qubit.	64
4.20	Amplitude stability of the 1762 nm laser over 50 pulses.	64
4.21	Error per Clifford in the optical qubit vs detuning of the 1762 nm laser.	65

4.22	Coherence time of the Optical qubit and a non-clock metastable qubit.	66
4.23	Optical qubit frequency tracked over a 40+ hr period.	67
4.24	Simulation of the error contribution from ion temperature for the optical qubit.	68
4.25	Procedure to convert from metastable qubit to optical qubit in both (a) the Clock qubit configuration and (b) the Hybrid qubit configuration.	71
4.26	Results from mapping the Hybrid qubit and Clock qubit between their optical and metastable configurations.	73
4.27	State fidelity measurement after mapping the Hybrid and Clock qubit between their optical and metastable configurations.	74
5.1	(a) The ideal output of a laser based addressing system for trapped ions. (b) Diagram of the V-groove array (VGA) and photonic chip.	80
5.2	Telecentric imaging system which maps the SPIM-WG output onto the ions.	81
5.3	Ramsey and KDD sequences used to measure light shifts from the waveguide system.	83
5.4	Cross-talk results from scanning a single ion through the output of the SPIM waveguide array.	84
6.1	Diagram of the SPAM protocol.	90
6.2	Simulation of repeat iterations of the SPAM protocol under the Markov matrix treatment in Eq. 6.1.	93
6.3	Probability of a full successfully prepared register of qubits.	94
6.4	Simulation of observable bias in the SPAM protocol.	98
6.5	Histogram of the counts collected on camera for the second readout in three sequential population detection steps.	104
6.6	Probability of decay of the metastable level as a function of the delay after shelving from $S_{1/2}$ to $D_{5/2}$.	105
6.7	Histogram of counts from the camera when performing the SPAM protocol on the metastable qubit.	106
6.8	The SPAM error after each post-selection step for the metastable qubit.	107
6.9	Results of implementing the SPAM protocol in the optical qubit.	108
6.10	Results of implementing the SPAM protocol in the optical qubit.	109
6.11	Measurement of the observable bias as the error from mapping between $S_{1/2}$ and $D_{5/2}$ is artificially increased.	110
7.1	The divergence between the arithmetic mean and the geometric mean of two variables.	118
7.2	Simulation of the mean value of $\eta^2(2n + 1)$ considering a full thermal distribution as well as the 80 th and 95 th percentile of the distribution.	121
7.3	Step by step diagram of sideband cooling with the metastable qubit.	123
7.4	Sideband thermometry data for the lower frequency radial mode taken on a single ion.	124
7.5	OTDF gate sequence.	125
7.6	Data from a spin-dependent force applied to a single ion.	126
7.7	Simulation of the phase space trajectories for two applications of a spin-dependent force with a π -pulse on the qubit separating them.	127
7.8	Data illustrating the difference between the “discrete” and “continuous” applications of the SDF on a single ion.	128
7.9	Heating rate measurement of the four radial modes on two ions.	133
7.10	Coherence time measurement of the LF Rocking mode.	134
7.11	Calibration of the Radial LF Rocking mode frequency using the SDF.	135

7.12	Average light shift vs optical power on two optical qubits when applying either the counter-propagating beam or co-propagating beam.	137
7.13	OTDF gate results.	139
7.14	Qubit populations vs power under the OTDF gate implementation.	140
7.15	Simulation of a perfectly calibrated gate in our system with no added noise from the 532 nm lasers.	141
7.16	Error of the OTDF gate vs increasing detuning error.	142
7.17	Lower frequency radial mode stability over a 13 hour period.	143
7.18	Gate error vs a constant offset error in Rabi frequency.	145
7.19	Measurement of the counter- and co-propagating beam amplitudes across 50 shots measured on a photodiode.	146
7.20	Gate error vs Rabi frequency asymmetry in the two SDF's which make up the OTDF gate.	147
7.21	Ramsey phase measurement with the OTDF gate force far detuned. We measure a phase offset of 0.22(2) rad from the gate sequence implying there is a consistent amplitude offset between the two pulses. We compensate this single-qubit rotation in the final gate implementation in Fig. 7.13.	148
7.22	Gate error vs coherence time of the LF Rocking mode.	148
7.23	Coherence time measurement when implementing the OTDF sequence.	149
7.24	Gate error vs spin coherence time.	150
A.1	Motional coherence time measurement for all four radial modes in a two ion chain.	159

References

- [1] D. T. C. Allcock et al. “*omg* blueprint for trapped ion quantum computing with metastable states”. In: *Applied Physics Letters* 119.21 (Nov. 2021), p. 214002.
- [2] A. S. Sotirova et al. “Low cross-talk optical addressing of trapped-ion qubits using a novel integrated photonic chip”. In: *Light: Science & Applications* 13.1 (Aug. 2024), p. 199.
- [3] A. S. Sotirova & J.D. Leppard et al. “High-fidelity heralded quantum state preparation and measurement”. In: *arXiv preprint arXiv:2409.05805* (2024).
- [4] B. C. Sawyer and K. R. Brown. “Wavelength-insensitive, multispecies entangling gate for group-2 atomic ions”. In: *Phys. Rev. A* 103 (2 Feb. 2021), p. 022427. URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.022427>.
- [5] J. Jumper et al. “Highly accurate protein structure prediction with AlphaFold”. In: *Nature* 596.7873 (Aug. 2021), pp. 583–589. URL: <https://doi.org/10.1038/s41586-021-03819-2>.
- [6] J. Abramson et al. “Accurate structure prediction of biomolecular interactions with AlphaFold 3”. In: *Nature* 630.8016 (June 2024), pp. 493–500. URL: <https://doi.org/10.1038/s41586-024-07487-w>.
- [7] T. Jones et al. “QuEST and High Performance Simulation of Quantum Computers”. In: *Scientific Reports* 9.1 (July 2019), p. 10736. URL: <https://doi.org/10.1038/s41598-019-47174-9>.
- [8] R. P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6 (June 1982), pp. 467–488. URL: <https://doi.org/10.1007/BF02650179>.
- [9] R. P. Feynman. “Quantum mechanical computers”. In: *Foundations of Physics* 16.6 (June 1986), pp. 507–531. URL: <https://doi.org/10.1007/BF01886518>.
- [10] P. Benioff. “Quantum mechanical Hamiltonian models of Turing machines”. In: *Journal of Statistical Physics* 29.3 (1982), pp. 515–546.
- [11] D. Deutsch. “Quantum theory, the Church–Turing principle and the universal quantum computer”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 400.1818 (1985), pp. 97–117.
- [12] D. P. DiVincenzo. “The Physical Implementation of Quantum Computation”. In: *Fortschritte der Physik* 48.9-11 (2000), pp. 771–783.
- [13] H. J. Kimble. “The quantum internet”. In: *Nature* 453.7198 (June 2008), pp. 1023–1030. URL: <https://doi.org/10.1038/nature07127>.

- [14] P. Drmota et al. “Verifiable Blind Quantum Computing with Trapped Ions and Single Photons”. In: *Phys. Rev. Lett.* 132 (15 Apr. 2024), p. 150604. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.132.150604>.
- [15] D. Main et al. “Distributed quantum computing across an optical network link”. In: *Nature* 638.8050 (Feb. 2025), pp. 383–388. URL: <https://doi.org/10.1038/s41586-024-08404-x>.
- [16] D. P. Nadlinger et al. “Experimental quantum key distribution certified by Bell’s theorem”. In: *Nature* 607.7920 (July 2022), pp. 682–686. URL: <https://doi.org/10.1038/s41586-022-04941-5>.
- [17] M. Kjaergaard et al. “Superconducting qubits: Current state of play”. In: *Annual Review of Condensed Matter Physics* 11.1 (2020), pp. 369–395.
- [18] R. Acharya et al. “Quantum error correction below the surface code threshold”. In: *Nature* 638.8052 (Feb. 2025), pp. 920–926. URL: <https://doi.org/10.1038/s41586-024-08449-y>.
- [19] S. Pezzagna and J. Meijer. “Quantum computer based on color centers in diamond”. In: *Applied Physics Reviews* 8.1 (2021).
- [20] N.-C. Chiu et al. “Continuous operation of a coherent 3,000-qubit system”. In: *arXiv preprint arXiv:2506.20660* (2025).
- [21] D. Bluvstein et al. “Architectural mechanisms of a universal fault-tolerant quantum computer”. In: *arXiv preprint arXiv:2506.20661* (2025).
- [22] C. D. Bruzewicz et al. “Trapped-ion quantum computing: Progress and challenges”. In: *Applied physics reviews* 6.2 (2019).
- [23] S. Kotochigova, S. Gupta, and B. Blinov. “Neutral atoms in optical tweezers as messenger qubits for scaling up a trapped ion quantum computer”. In: *arXiv preprint arXiv:2501.04223* (2025).
- [24] M. C. Smith et al. “Single-Qubit Gates with Errors at the 10^{-7} Level”. In: *Phys. Rev. Lett.* 134 (23 June 2025), p. 230601. URL: <https://link.aps.org/doi/10.1103/42w2-6ccy>.
- [25] M. M. et al. “Scalable electronic control of trapped-ion qubits”. In: *Quantum 2.0 Conference and Exhibition*. Optica Publishing Group, 2024, QM3A.4. URL: <https://opg.optica.org/abstract.cfm?URI=QUANTUM-2024-QM3A.4>.
- [26] P. Wang, C.-Y. Luan, M. Qiao, et al. “Single ion qubit with estimated coherence time exceeding one hour”. In: *Nature Communications* 12.1 (Jan. 2021), p. 233. URL: <https://doi.org/10.1038/s41467-020-20330-w>.
- [27] P. W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In: *SIAM Journal on Computing* 26.5 (1997), pp. 1484–1509. eprint: <https://doi.org/10.1137/S0097539795293172>. URL: <https://doi.org/10.1137/S0097539795293172>.
- [28] D. Kielpinski, C. Monroe, and D. J. Wineland. “Architecture for a large-scale ion-trap quantum computer”. In: *Nature* 417.6890 (2002), pp. 709–711.

- [29] J. M. Pino et al. “Demonstration of the trapped-ion quantum CCD computer architecture”. In: *Nature* 592.7853 (Apr. 2021), pp. 209–213. URL: <https://doi.org/10.1038/s41586-021-03318-4>.
- [30] L.-M. Duan and C. Monroe. “Colloquium: Quantum networks with trapped ions”. In: *Rev. Mod. Phys.* 82 (2 Apr. 2010), pp. 1209–1224. URL: <https://link.aps.org/doi/10.1103/RevModPhys.82.1209>.
- [31] I. V. Inlek et al. “Multispecies Trapped-Ion Node for Quantum Networking”. In: *Phys. Rev. Lett.* 118 (25 June 2017), p. 250502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.250502>.
- [32] Y.-Y. Huang et al. “Realization of a functioning dual-type trapped-ion quantum network node”. In: *arXiv preprint arXiv:2506.23562* (2025).
- [33] A. Reiserer. “Colloquium: Cavity-enhanced quantum network nodes”. In: *Rev. Mod. Phys.* 94 (4 Dec. 2022), p. 041003. URL: <https://link.aps.org/doi/10.1103/RevModPhys.94.041003>.
- [34] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52 (4 Oct. 1995), R2493–R2496. URL: <https://link.aps.org/doi/10.1103/PhysRevA.52.R2493>.
- [35] A. M Steane. “Error correcting codes in quantum theory”. In: *Physical Review Letters* 77.5 (1996), p. 793.
- [36] J. Preskill. “Beyond NISQ: The Megaquop Machine”. In: *ACM Transactions on Quantum Computing* 6.3 (Apr. 2025). URL: <https://doi.org/10.1145/3723153>.
- [37] M. J. McDonnell et al. “High-Efficiency Detection of a Single Quantum of Angular Momentum by Suppression of Optical Pumping”. In: *Phys. Rev. Lett.* 93 (15 Oct. 2004), p. 153601. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.93.153601>.
- [38] S. A. Moses et al. “A Race-Track Trapped-Ion Quantum Processor”. In: *Phys. Rev. X* 13 (4 Dec. 2023), p. 041052. URL: <https://link.aps.org/doi/10.1103/PhysRevX.13.041052>.
- [39] A. Paul and C. Noel. “Ion-chain sympathetic cooling and gate dynamics”. In: *Phys. Rev. Appl.* 22 (4 Oct. 2024), p. 044033. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.22.044033>.
- [40] V. M. Schäfer. “Fast gates and mixed-species entanglement with trapped ions”. PhD thesis. University of Oxford, UK, 2018.
- [41] A. C. Hughes. “Benchmarking Memory and Logic Gates for Trapped-Ion Quantum Computing”. PhD thesis. University of Oxford, UK, 2021.
- [42] M. Palmero et al. “Fast transport of mixed-species ion chains within a Paul trap”. In: *Phys. Rev. A* 90 (5 Nov. 2014), p. 053408. URL: <https://link.aps.org/doi/10.1103/PhysRevA.90.053408>.
- [43] C. D. Bruzewicz et al. “Dual-species, multi-qubit logic primitives for Ca⁺/Sr⁺ trapped-ion crystals”. In: *npj Quantum Information* 5.1 (Nov. 2019), p. 102. URL: <https://doi.org/10.1038/s41534-019-0218-z>.

- [44] A. C. Hughes et al. “Benchmarking a High-Fidelity Mixed-Species Entangling Gate”. In: *Phys. Rev. Lett.* 125 (8 Aug. 2020), p. 080504. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.080504>.
- [45] H.-X. Yang et al. “Realizing coherently convertible dual-type qubits with the same ion species”. In: *Nature Physics* 18.9 (Sept. 2022), pp. 1058–1061.
- [46] Z.-Y. Chen et al. “Non-invasive mid-circuit measurement and reset on atomic qubits”. In: *arXiv preprint arXiv:2504.12538* (2025).
- [47] A. S. Sotirova. “Trapped Ion Quantum Information Processing Using Multiple Qubit Encodings”. PhD thesis. University of Oxford, UK, 2024.
- [48] S A King et al. “Absolute frequency measurement of the $2S_{1/2}$ – $2F_{7/2}$ electric octupole transition in a single ion of $^{171}\text{Yb}^+$ with 1015 fractional uncertainty”. In: *New Journal of Physics* 14.1 (Jan. 2012), p. 013045. URL: <https://doi.org/10.1088/1367-2630/14/1/013045>.
- [49] D. Hucul et al. “Spectroscopy of a synthetic trapped ion qubit”. In: *Physical review letters* 119.10 (2017), p. 100501.
- [50] J. E. Christensen et al. “High-fidelity manipulation of a qubit enabled by a manufactured nucleus”. In: *npj Quantum Information* 6.1 (Apr. 2020), p. 35.
- [51] Brendan M. White et al. “Isotope-selective laser ablation ion-trap loading of $^{137}\text{Ba}^+$ using a BaCl_2 target”. In: *Phys. Rev. A* 105 (3 Mar. 2022), p. 033102. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.033102>.
- [52] X. Shi et al. “Ablation loading of barium ions into a surface-electrode trap”. In: *Applied Physics Letters* 122.26 (June 2023), p. 264002. URL: <https://doi.org/10.1063/5.0149778>.
- [53] T. P. Harty. *Atomic Physics*. https://github.com/OxfordIonTrapGroup/atomic_physics/. Version 1.0.2. ORCID: <https://orcid.org/0000-0003-4077-226X>. 2017.
- [54] T. L. Nicholson et al. “Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty”. In: *Nature communications* 6.1 (2015), p. 6896.
- [55] D. G. Matei et al. “ $1.5 \mu\text{m}$ Lasers with Sub-10 mHz Linewidth”. In: *Phys. Rev. Lett.* 118 (26 June 2017), p. 263202. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.263202>.
- [56] S. Shivam & F. Pokorny et al. *Utility of virtual qubits in trapped-ion quantum computers*. 2024. arXiv: 2406.19332 [quant-ph]. URL: <https://arxiv.org/abs/2406.19332>.
- [57] Martin Ringbauer et al. “A universal qudit quantum processor with trapped ions”. In: *Nature Physics* 18.9 (Sept. 2022), pp. 1053–1057. URL: <https://doi.org/10.1038/s41567-022-01658-0>.
- [58] A. S. Nikolaeva, Evgeniy O. Kiktenko, and Aleksey K. Fedorov. “Efficient realization of quantum algorithms with qudits”. In: *EPJ Quantum Technology* 11.1 (June 2024), p. 43. URL: <https://doi.org/10.1140/epjqt/s40507-024-00250-0>.

- [59] P. Hrmo et al. “Native qudit entanglement in a trapped ion quantum processor”. In: *Nature Communications* 14.1 (Apr. 2023), p. 2242. URL: <https://doi.org/10.1038/s41467-023-37375-2>.
- [60] M. A. Aksenov et al. “Realizing quantum gates with optically addressable $^{171}\text{Yb}^+$ ion qudits”. In: *Phys. Rev. A* 107 (5 May 2023), p. 052612. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.052612>.
- [61] W. C. Campbell and E. R. Hudson. “Polyqubit quantum processing”. In: *arXiv preprint arXiv:2210.15484* (2022).
- [62] N. Greenberg et al. “Trapping \backslash
 Ba^+
with seven – fold enhanced efficiency utilizing an autoionizing resonance”. In: *Quantum Science and Technology* 9.3 (2024), p. 035023.
- [63] B. Wang et al. “Highly efficient and isotope selective photo-ionization of barium atoms using diode laser and LED light”. In: *Opt. Express* 19.17 (Aug. 2011), pp. 16438–16447. URL: <https://opg.optica.org/oe/abstract.cfm?URI=oe-19-17-16438>.
- [64] A. V. Steele et al. “Photoionization and photoelectric loading of barium ion traps”. In: *Phys. Rev. A* 75 (5 May 2007), p. 053404. URL: <https://link.aps.org/doi/10.1103/PhysRevA.75.053404>.
- [65] K. Yamada et al. “Separation of Barium Isotopes by Selective Two-Step Photoionization Process”. In: *Journal of Nuclear Science and Technology* 25.8 (1988), pp. 641–648.
- [66] C. Huang et al. “Electromagnetically Induced Transparency Cooling of High-Nuclear-Spin Ions”. In: *Phys. Rev. Lett.* 133 (11 Sept. 2024), p. 113204. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.133.113204>.
- [67] C. Monroe et al. “Resolved-Sideband Raman Cooling of a Bound Atom to the 3D Zero-Point Energy”. In: *Phys. Rev. Lett.* 75 (22 Nov. 1995), pp. 4011–4014. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.75.4011>.
- [68] A. J. Rasmusson et al. “Optimized pulsed sideband cooling and enhanced thermometry of trapped ions”. In: *Phys. Rev. A* 104 (4 Oct. 2021), p. 043108. URL: <https://link.aps.org/doi/10.1103/PhysRevA.104.043108>.
- [69] S. R. Vizvary et al. “Eliminating Qubit-Type Cross-Talk in the *omg* Protocol”. In: *Phys. Rev. Lett.* 132 (26 June 2024), p. 263201. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.132.263201>.
- [70] K. Rudinger et al. “Experimental Demonstration of a Cheap and Accurate Phase Estimation”. In: *Phys. Rev. Lett.* 118 (19 May 2017), p. 190502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.190502>.
- [71] E. Magesan, J. M. Gambetta, and Joseph Emerson. “Scalable and Robust Randomized Benchmarking of Quantum Processes”. In: *Phys. Rev. Lett.* 106 (18 May 2011), p. 180504. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.106.180504>.

- [72] D. A. Steck. *Quantum and Atom Optics*. Available online at <https://steck.us/teaching> (revision 0.16.3, 16 April 2025). June 2007.
- [73] C. J. Foot. *Atomic Physics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, June 2005.
- [74] R. Loudon. *The Quantum Theory of Light*. 3rd ed. Oxford University Press, 2000. Chap. 4.9.
- [75] Y. Yu et al. “In-situ mid-circuit qubit measurement and reset in a single-species trapped-ion quantum computing system”. In: *arXiv preprint arXiv:2504.12544* (2025).
- [76] C. J. Ballance et al. “High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits”. In: *Physical Review Letters* 117.6 (Aug. 2016), p. 060504.
- [77] D. J. Wineland et al. “Experimental issues in coherent quantum-state manipulation of trapped atomic ions”. In: *Journal of research of the National Institute of Standards and Technology* 103.3 (1998), p. 259.
- [78] D. F. James and J. Jerke. “Effective Hamiltonian theory and its applications in quantum information”. In: *Canadian Journal of Physics* 85.6 (2007), pp. 625–632. eprint: <https://doi.org/10.1139/p07-060>. URL: <https://doi.org/10.1139/p07-060>.
- [79] D. F. V. James. “Quantum dynamics of cold trapped ions with application to quantum computation”. In: *Applied Physics B* 66.2 (Feb. 1998), pp. 181–190. URL: <https://doi.org/10.1007/s003400050373>.
- [80] S. Kimmel, Guang H. Low, and T. J. Yoder. “Robust calibration of a universal single-qubit gate set via robust phase estimation”. In: *Phys. Rev. A* 92 (6 Dec. 2015), p. 062315. URL: <https://link.aps.org/doi/10.1103/PhysRevA.92.062315>.
- [81] B. C. Nichol. “An elementary quantum network of entangled optical atomic clocks”. PhD thesis. University of Oxford, UK, 2012.
- [82] W. M. Itano et al. “Quantum projection noise: Population fluctuations in two-level systems”. In: *Physical Review A* 47.5 (May 1993), pp. 3554–3570.
- [83] T. Hurant et al. “Few-Shot, Robust Calibration of Single Qubit Gates Using Bayesian Robust Phase Estimation”. In: *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Vol. 01. 2024, pp. 1244–1253.
- [84] D. Main. “Distributed Quantum Computing across a Two-Node Trapped-Ion Quantum Network”. PhD thesis. University of Oxford, UK, 2024.
- [85] C. Dankert et al. “Exact and approximate unitary 2-designs and their application to fidelity estimation”. In: *Phys. Rev. A* 80 (1 July 2009), p. 012304. URL: <https://link.aps.org/doi/10.1103/PhysRevA.80.012304>.
- [86] A. D. Leu et al. “Fast, High-Fidelity Addressed Single-Qubit Gates Using Efficient Composite Pulse Sequences”. In: *Phys. Rev. Lett.* 131 (12 Sept. 2023), p. 120601. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.120601>.
- [87] C. M. Löschnauer et al. “Scalable, high-fidelity all-electronic control of trapped-ion qubits”. In: *arXiv preprint arXiv:2407.07694* (2024).

- [88] I. Pogorelov et al. “Compact Ion-Trap Quantum Computing Demonstrator”. In: *PRX Quantum* 2 (2 June 2021), p. 020343. URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.020343>.
- [89] Y.-L. Chen et al. “Low-crosstalk optical addressing system for atomic qubits based on multiple objectives and acousto-optic deflectors”. In: *Phys. Rev. Appl.* 22 (5 Nov. 2024), p. 054003. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.22.054003>.
- [90] Y.-H. Hou et al. “Individually addressed entangling gates in a two-dimensional ion crystal”. In: *Nature Communications* 15.1 (Nov. 2024), p. 9710. URL: <https://doi.org/10.1038/s41467-024-53405-z>.
- [91] S. Crain et al. “Individual addressing of trapped 171Yb^+ ion qubits using a microelectromechanical systems-based beam steering system”. In: *Applied Physics Letters* 105.18 (Nov. 2014), p. 181115. eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.4900754/14304206/181115_1_online.pdf. URL: <https://doi.org/10.1063/1.4900754>.
- [92] Y. Wang et al. “High-Fidelity Two-Qubit Gates Using a Microelectromechanical-System-Based Beam Steering System for Individual Qubit Addressing”. In: *Phys. Rev. Lett.* 125 (15 Oct. 2020), p. 150505. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.150505>.
- [93] C.-Y. Shih et al. “Reprogrammable and high-precision holographic optical addressing of trapped ions for scalable quantum control”. In: *npj Quantum Information* 7.1 (Apr. 2021), p. 57. URL: <https://doi.org/10.1038/s41534-021-00396-0>.
- [94] A. Binai-Motlagh et al. “A guided light system for agile individual addressing of Ba^+ qubits with 10- 4 level intensity crosstalk”. In: *Quantum Science and Technology* 8.4 (2023), p. 045012.
- [95] C. L. Craft et al. “Low-Crosstalk, Silicon-Fabricated Optical Waveguides for Laser Delivery to Matter Qubits”. In: *arXiv preprint arXiv:2406.17607* (2024).
- [96] A. M. Souza, G. A. Álvarez, and D. Suter. “Robust Dynamical Decoupling for Quantum Computing and Quantum Memory”. In: *Phys. Rev. Lett.* 106 (24 June 2011), p. 240501. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.106.240501>.
- [97] B. T. Torosov and N. V. Vitanov. “Narrowband composite two-qubit gates for crosstalk suppression”. In: *Phys. Rev. A* 107 (3 Mar. 2023), p. 032618. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.032618>.
- [98] J. Flannery et al. “Physical coherent cancellation of optical addressing crosstalk in a trapped-ion experiment”. In: *Quantum Science and Technology* 10.1 (Oct. 2024), p. 015012. URL: <https://dx.doi.org/10.1088/2058-9565/ad8371>.
- [99] A. Kubica et al. *Erasure Qubits: Overcoming the T_1 Limit in Superconducting Circuits*. Nov. 2023. URL: <https://link.aps.org/doi/10.1103/PhysRevX.13.041022>.

- [100] Y. Wu et al. “Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays”. In: *Nature Communications* 13.1 (Aug. 2022), p. 4657.
- [101] A. H. Myerson et al. “High-Fidelity Readout of Trapped-Ion Qubits”. In: *Physical Review Letters* 100.20 (May 2008), p. 200502.
- [102] A. H. Burrell et al. “Scalable simultaneous multiqubit readout with 99.99% single-shot fidelity”. In: *Physical Review A* 81.4 (Apr. 2010), p. 040302.
- [103] A. Burrell. “High fidelity readout of trapped ion qubits”. PhD thesis. University of Oxford, UK, 2010.
- [104] S. Olmschenk et al. “Manipulation and detection of a trapped Yb + hyperfine qubit”. In: *Physical Review A* 76.5 (Nov. 2007), p. 052314.
- [105] F. A. An et al. “High Fidelity State Preparation and Measurement of Ion Hyperfine Qubits with $I > 1/2$ ”. In: *Physical Review Letters* 129.13 (Sept. 2022), p. 130501.
- [106] T. P. Harty et al. “High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit”. In: *Physical Review Letters* 113.22 (Nov. 2014), p. 220501.
- [107] A. D. Leu et al. *Polarisation-insensitive state preparation for trapped-ion hyperfine qubits*. arXiv:2406.14448. June 2024. URL: <http://arxiv.org/abs/2406.14448>.
- [108] M. Acton et al. *Near-Perfect Simultaneous Measurement of a Qubit Register*. arXiv:quant-ph/0511257. Feb. 2006. URL: <http://arxiv.org/abs/quant-ph/0511257>.
- [109] W. Nagourney, J. Sandberg, and H. Dehmelt. “Shelved optical electron amplifier: Observation of quantum jumps”. In: *Physical Review Letters* 56.26 (June 1986), pp. 2797–2799.
- [110] Th. Sauter et al. “Observation of Quantum Jumps”. In: *Physical Review Letters* 57.14 (Oct. 1986), pp. 1696–1698.
- [111] D. Hayes et al. “Eliminating Leakage Errors in Hyperfine Qubits”. In: *Physical Review Letters* 124.17 (Apr. 2020), p. 170501.
- [112] N. C. Brown and K. R. Brown. “Leakage mitigation for quantum error correction using a mixed qubit scheme”. In: *Physical Review A* 100.3 (Sept. 2019), p. 032325.
- [113] B. M. Varbanov et al. “Leakage detection for a transmon-based surface code”. In: *npj Quantum Information* 6.1 (Dec. 2020), pp. 1–13.
- [114] M. Kang, W. C. Campbell, and K. R. Brown. “Quantum Error Correction with Metastable States of Trapped Ions Using Erasure Conversion”. In: *PRX Quantum* 4.2 (June 2023), p. 020358.
- [115] S. Aaronson. *Quantum Computing, Postselection, and Probabilistic Polynomial-Time*. arXiv:quant-ph/0412187. Dec. 2004. URL: <http://arxiv.org/abs/quant-ph/0412187>.
- [116] Y. L. Lim, Almut Beige, and Leong Chuan Kwek. “Repeat-Until-Success Linear Optics Distributed Quantum Computing”. In: *Physical Review Letters* 95.3 (July 2005), p. 030505.

- [117] C. W. Hogle et al. “High-fidelity trapped-ion qubit operations with scalable photonic modulators”. In: *npj Quantum Information* 9.1 (2023), p. 74.
- [118] S. Saner et al. “Breaking the Entangling Gate Speed Limit for Trapped-Ion Qubits Using a Phase-Stable Standing Wave”. In: *Phys. Rev. Lett.* 131 (22 Dec. 2023), p. 220601. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.220601>.
- [119] M. J. Boguslawski et al. “Raman Scattering Errors in Stimulated-Raman-Induced Logic Gates in $^{133}\text{Ba}^+$ ”. In: *Phys. Rev. Lett.* 131 (6 Aug. 2023), p. 063001. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.131.063001>.
- [120] I. D. Moore et al. “Photon scattering errors during stimulated Raman transitions in trapped-ion qubits”. In: *Physical Review A* 107.3 (Mar. 2023), p. 032413.
- [121] Z. Zhang et al. “Branching fractions for $P_{3/2}$ decays in Ba^+ ”. In: *Physical Review A* 101.6 (June 2020), p. 062515.
- [122] P. Grangier, Juan Ariel Levenson, and Jean-Philippe Poizat. “Quantum non-demolition measurements in optics”. In: *Nature* 396.6711 (Dec. 1998), pp. 537–542.
- [123] V. Negnevitsky et al. “Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register”. In: *Nature* 563.7732 (Nov. 2018), pp. 527–531.
- [124] M. Cerezo et al. “Variational quantum algorithms”. In: *Nature Reviews Physics* 3.9 (Sept. 2021), pp. 625–644.
- [125] R. R. Ferguson et al. “Measurement-Based Variational Quantum Eigensolver”. In: *Physical Review Letters* 126.22 (June 2021), p. 220501.
- [126] J. Řeháček et al. “Diluted maximum-likelihood algorithm for quantum tomography”. In: *Phys. Rev. A* 75 (4 Apr. 2007), p. 042108. URL: <https://link.aps.org/doi/10.1103/PhysRevA.75.042108>.
- [127] W. Neuhauser et al. “Optical-Sideband Cooling of Visible Atom Cloud Confined in Parabolic Well”. In: *Physical Review Letters* 41.4 (July 1978), pp. 233–236.
- [128] D. J. Wineland, R. E. Drullinger, and F. L. Walls. “Radiation-Pressure Cooling of Bound Resonant Absorbers”. In: *Physical Review Letters* 40.25 (June 1978), pp. 1639–1642.
- [129] D. J. Wineland and W. M. Itano. “Laser cooling of atoms”. In: *Physical Review A* 20.4 (Oct. 1979), pp. 1521–1540.
- [130] D. Gottesman. *Stabilizer Codes and Quantum Error Correction*. arXiv:quant-ph/9705052. May 1997. URL: <http://arxiv.org/abs/quant-ph/9705052>.
- [131] J. Roffe. “Quantum error correction: an introductory guide”. In: *Contemporary Physics* 60.3 (July 2019), pp. 226–245. URL: <http://dx.doi.org/10.1080/00107514.2019.1667078>.
- [132] K. Mølmer and A. Sørensen. “Multiparticle Entanglement of Hot Trapped Ions”. In: *Phys. Rev. Lett.* 82 (9 Mar. 1999), pp. 1835–1838. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.82.1835>.

- [133] A. Sørensen and K. Mølmer. “Quantum Computation with Ions in Thermal Motion”. In: *Phys. Rev. Lett.* 82 (9 Mar. 1999), pp. 1971–1974. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.82.1971>.
- [134] G Milburn. “Simulating nonlinear spin models in an ion trap (1999)”. In: *arXiv preprint quant-ph/9908037* ().
- [135] D. Leibfried et al. “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate”. In: *Nature* 422.6930 (Mar. 2003), pp. 412–415. URL: <https://doi.org/10.1038/nature01492>.
- [136] C. H. Baldwin et al. “High-fidelity light-shift gate for clock-state qubits”. In: *Phys. Rev. A* 103 (1 Jan. 2021), p. 012603. URL: <https://link.aps.org/doi/10.1103/PhysRevA.103.012603>.
- [137] O. Băzăvan et al. “Synthesizing a $\hat{\sigma}_z$ spin-dependent force for optical, metastable, and ground-state trapped-ion qubits”. In: *Phys. Rev. A* 107 (2 Feb. 2023), p. 022617. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.022617>.
- [138] C. R. Clark et al. “High-Fidelity Bell-State Preparation with $^{40}\text{Ca}^+$ Optical Qubits”. In: *Phys. Rev. Lett.* 127 (13 Sept. 2021), p. 130505. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.130505>.
- [139] P. C. Haljan et al. “Spin-Dependent Forces on Trapped Ions for Phase-Stable Quantum Gates and Entangled States of Spin and Motion”. In: *Phys. Rev. Lett.* 94 (15 Apr. 2005), p. 153602. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.94.153602>.
- [140] W. Magnus. “On the exponential solution of differential equations for a linear operator”. In: *Communications on Pure and Applied Mathematics* 7.4 (1954), pp. 649–673. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.3160070404>.
- [141] O. Băzăvan. “Easy as XYZ: Exploring non-commuting coherent interactions in trapped ions”. PhD thesis. University of Oxford, UK, 2024.
- [142] S. A. Schulz et al. “Sideband cooling and coherent dynamics in a microchip multi-segmented ion trap”. In: *New Journal of Physics* 10.4 (Apr. 2008), p. 045007. URL: <https://dx.doi.org/10.1088/1367-2630/10/4/045007>.
- [143] M. F. Gely et al. “In situ characterization of qubit-drive phase distortions”. In: *Phys. Rev. Appl.* 22 (2 Aug. 2024), p. 024001. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.22.024001>.
- [144] P. J. Low et al. “Quantum logic operations and algorithms in a single 25-level atomic qudit”. In: *arXiv preprint arXiv:2507.15799* (2025).