Trapped Ion Quantum Information Processing Using Multiple Qubit Encodings



Ana S. Sotirova

Balliol College

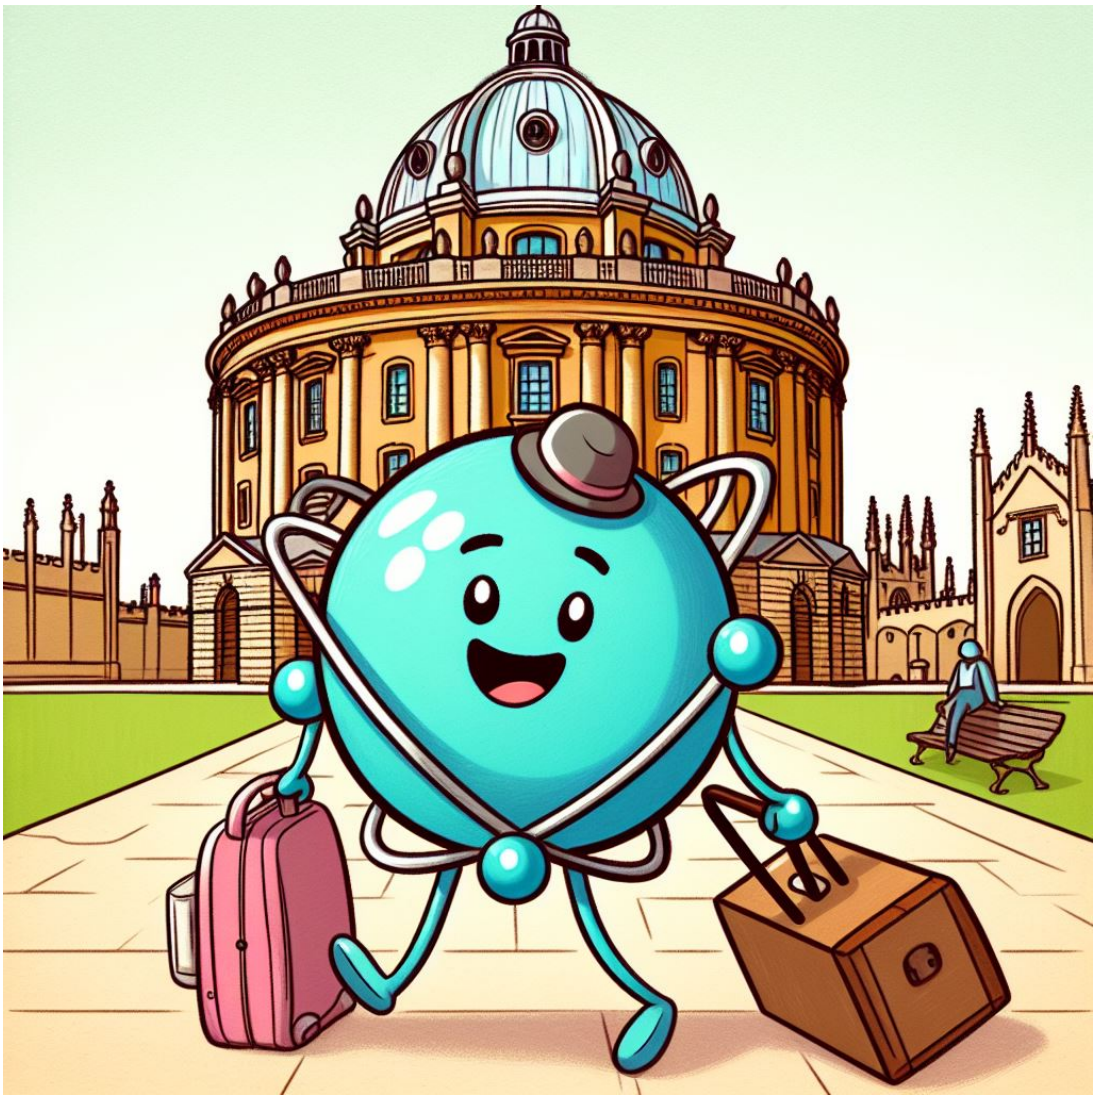
University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2024

This is the story of how Barry came to Oxford...



This thesis is dedicated to my sanity. May it rest in pieces.

Acknowledgements

The most valuable lesson that I learned over the past four years is that the true measure of any experience lies in the people you share it with. In this regard, I consider myself incredibly fortunate. Reflecting on these experiences fills me with deep gratitude and joy, and I want to take this opportunity to thank everyone who has been part of this journey.

First and foremost, I would like to express my heartfelt gratitude to my supervisor, Chris Ballance. Chris, your unwavering patience, invaluable help, and steadfast support over the past few years have been truly remarkable. Your passion for ion traps is not only evident but also contagious, and it always kept me excited about working in the lab. You have provided me with every opportunity I sought. I truly appreciate your willingness to offer help with the experiment and otherwise, even at the most ungodly hours. The amount I have learned from you is simply astounding, and I am deeply thankful for your guidance and mentorship. It will take me a lifetime to pay you back.

I would also like to extend my heartfelt thanks to my post-doc, Fabian Pokorny. Fabian, being your 'head minion' has been the best title I ever held. Your presence in the lab defined my experience, as you seamlessly balanced the roles of an exceptional teacher, a true friend, and an impromptu therapist. Thank you for your unwavering patience, for always listening to my woes, and for consistently offering your support and guidance. I miss working with you every day.

I had the immense pleasure of sharing the lab with three other DPhil students, Jamie, Andres, and Sophie. Jamie, thank you for upgrading our control system, even though it cost you a few white hairs. Andres, thank you for taming our trapping potentials - a lot of the work in this thesis would not have been possible otherwise. Sophie, it is unbelievable that you only joined the lab a few months ago. Thank you for your support in getting the two-qubit gates to work. It has been a true privilege working with you all. I am sure you will go on to achieve great things, and I am happy to have been a small part of your journeys.

My gratitude also extends to David Nadlinger, who made an excellent hostage postdoc during the last year of my DPhil. Without your help and support, much of this thesis would never have seen the light of day - thank you!

The whole ion trap quantum computing group at Oxford has been a great joy to work with. I have learned something from every one of you and I hope our

paths will cross again. I want to thank the other PI in the group, David Lucas, for introducing me into the group. Thank you to Raghu and Gabriel for always having my back, even when times were hard. Dougal, thank you for being my late-night buddy in the lab. Assembling a cryogenic system at 3 a.m. will always be one of my most cherished memories in the group. Thank you to Alan Francis for ensuring our deliveries made it to the lab. You have customs superpowers. I would also like to thank the mechanical workshop for their incredible flexibility with all the bits and pieces I needed made for the lab quickly.

During the last two years of my DPhil, I had the great pleasure of working as a part-time intern at Oxford Ionics, where some of the work presented in this thesis was born. I now work full-time at Oxford Ionics, and I thank the entire team for their immense support in ensuring that I could do my work while finishing this thesis. I would like to thank my manager, Maciej. Thank you for your great advice, understanding and flexibility in structuring my work around my DPhil. Neither this internship nor this thesis would have happened without your help. Timko, thank you for being so supportive and patient with me while I was settling in and for giving me photography lessons. Emily and Sophie, thank you very much for your great support and for keeping the place running as well as it does. Susanna, thank you for the many puppy pictures that got me through writing this thesis. Jacopo and Steven, thank you for introducing me to the ion trap systems. Working at the company has been an incredible experience, and I look forward to my future here.

Getting to do a DPhil was, of course, preceded by years of learning, and it is thanks to my great teachers that I am here today. I would first like to thank the physics fellows at Balliol College, Andre, Armin, David, and Chris, whose tutorials were a highlight of my undergraduate degree. During my undergraduate degree, I also undertook several research projects. I want to thank the people who supervised me during these projects: Chiara, Roland, Robin, and Jonathan at ETH; Chris, Tom, Clemens, and Jochen at the Oxford ion trap group; and Peter Spring and Peter Leek at the Oxford superconducting quantum devices group. Thank you all for teaching a toddler how to do experimental physics. None of this would have been possible without the foundation you helped me build. I would also like to thank my high school maths and physics teachers, Mr Kalinov and Mr Mihalkov. Thank you for believing in me and pushing me to do my best. I made it here, thanks to you.

Living at Balliol College over the past few years has been a great source of distraction and fun. First, I would like to thank our fantastic night porters, Mark and Steve. Thank you both for making Holywell Manor feel like home. Steve, thank you very much for all the friendly chats and impromptu mandolin concerts

that got me through many late-night study sessions. Mark, thank you for all the friendly chats and your help organising my travels on several occasions. Mark and Steve, you have been a constant source of joy and comfort over the past four years, and for that, I am truly grateful.

My friends at Balliol, Akari, Ana D., Zoe, Konrad, Keith, Calum, Elyse, Toby, Nick S., and Kumeren have provided a much-needed source of fun, distraction, and support over the past few years. Akari, thank you for being my private social secretary and for making sure I left my room from time to time over the past few months. Ana D., it was an honour being your flatmate for two years. Thank you for your friendship and support – I still miss your armchair.

Finally, I would like to thank my family and friends in Bulgaria. Mom, since I can remember, you have done everything you can to ensure I had whatever I needed. I know it wasn't always easy, and I cannot express how much I appreciate it. I love you. My aunt has taken endless care of me since I can remember and has always treated me like her own child. It is largely thanks to her support that I made it here today. I would also like to thank my grandparents for believing in me, teaching me never to give up, and showing me what family truly means. Finally, I would like to thank my best friend Maria, a.k.a. Shushy, who is the closest I will ever get to having a sister. Thank you for the friendship, the support, the craziness, and the experiences we shared. Thank you for listening to all of my nonsense and always having my back. I miss you, and I hope that one day we will live close to each other again.

Abstract

Trapped atomic ions are one of the most promising platforms for developing a quantum computer. On a small scale, they have shown the longest qubit coherence times [1] and the highest single- and two-qubit gate fidelities [2–5]. The challenge is finding ways to scale up the system while maintaining this level of performance to build usefully powerful quantum computers. In this thesis, we describe experimentally validated techniques to overcome some of the difficulties in scaling trapped-ion quantum computers.

We use the vast state space available within a single ion to encode qubits both in the ground level and the metastable level [6]. This architecture enables measurement, cooling, and reset of the ground-level qubits while protecting the state of the metastable level ones. We demonstrate the building blocks of this architecture using trapped $^{137}Ba^{+}$ ions. We implement a metastable level clock qubit in $^{137}Ba^{+}$ and measure a coherence time of 31(5) s. We show that the coherence time of this clock qubit in the presence of the laser light for cooling and measurement is 21.9(21) s. We further demonstrate a two-qubit entangling gate between two metastable level qubits necessary for building a universal gate set. Moreover, we perform an entangling gate between a ground-level qubit and a metastable level qubit, the equivalent of the mixed-species gate [7]. This gate facilitates the implementation of quantum non-demolition measurements (QNDs) in our architecture directly without having to redefine the qubit encoding.

Additionally, we present a new state preparation and measurement (SPAM) protocol that uses the ability to perform mid-circuit measurements to detect errors that occur during SPAM effectively. These errors can then be removed using post-selection or converted into erasures, one of the most straightforward errors to correct. We demonstrate the use of the protocol for three different qubit encodings in $^{137}Ba^{+}$: an optical qubit (O), a metastable level qubit (M), and a ground level qubit (G). For all three qubit types, we achieved the lowest reported SPAM infidelities of any qubit of O: $7(4) \times 10^{-6}$, M: $5(4) \times 10^{-6}$, and G: $8(4) \times 10^{-6}$.

The combination of these techniques can be used in building large-scale quantum computers based on trapped ions in the future.

Contents

1 Introduction **1**

1.1 Background 1

1.2 This work 5

1.3 Thesis outline 7

2 Architecture **8**

3 The barium ion **14**

3.1 Choice of ion 15

3.1.1 Choice of isotope 16

3.2 Photoionisation 18

3.3 Ion operations 19

3.3.1 Ion fluorescence 20

3.3.2 Doppler cooling 21

3.3.3 State preparation 22

3.3.4 Detection 23

3.3.5 Sideband cooling 25

3.3.6 Quadrupole transitions 25

3.4 Qubit encoding 26

3.4.1 Metastable level structure 26

3.4.2 Ground level structure 28

3.4.3 Qubit transitions 30

4 Apparatus **32**

4.1 The trap 33

4.1.1 Trap topology 33

4.1.2 Trap geometry 35

4.1.3 Trap signal delivery 37

4.1.4 Isotope-selective loading sequence 39

4.2 Imaging system 40

4.3 Laser systems 43

4.3.1 Photoionisation lasers 44

4.3.2 Ion operations lasers 45
4.3.3 Locking cavities 47
4.3.4 Raman lasers 48
4.3.5 Trap beam paths 52

5 Individual addressing system 55
5.1 Background 56
5.2 Individual addressing system design using a novel integrated photonic chip 58
5.2.1 Fibre network and photonic chip design 58
5.2.2 Coupling into the photonic chip 60
5.2.3 Integration with the ion trap system 62
5.3 Cross-talk characterisation using a single trapped ion 68
5.3.1 Robust phase estimation protocol 69
5.3.2 Knill Dynamical Decoupling Sequence 71
5.3.3 Measurement implementation 73
5.4 Summary 79

6 High-fidelity heralded state preparation and measurement 81
6.1 SPAM protocol overview 84
6.2 Protocol implementation in $^{137}Ba^{+}$ 87
6.3 High-fidelity population detection 89
6.3.1 Optical detection 90
6.3.2 Probability of decay of the $D_{5/2}$ level during detection . . . 96
6.3.3 Total expected SPAM error 99
6.4 Experimental results 100
6.4.1 SPAM results 100
6.4.2 SPAM data rates 107
6.5 Summary of SPAM results 109

7 Exploring a metastable level clock qubit 111
7.1 Choice of metastable level clock qubit 112
7.2 Experimental characterisation of the metastable level clock qubit . 114
7.2.1 Field-insensitive point and qubit frequency 114
7.2.2 Qubit coherence time 116
7.3 Isolation of the metastable level qubit from dissipative light . . . 122
7.4 Summary 125

8 Two-qubit gates **127**
8.1 Mølmer-Sørenson gate overview . . . 129
8.2 Raman beam setup . . . 131
8.2.1 Raman beam configuration . . . 131
8.2.2 Raman beam power . . . 132
8.2.3 Raman beam polarisation . . . 133
8.3 Motional mode structure . . . 138
8.3.1 Choice of gate mode . . . 139
8.3.2 Motional heating rates . . . 139
8.3.3 Lamb-Dicke parameter . . . 140
8.4 Gate implementation . . . 140
8.4.1 Qubit encoding . . . 140
8.4.2 General gate sequence . . . 141
8.5 Gate calibration . . . 141
8.5.1 Equal ion illumination . . . 141
8.5.2 Equal sideband strength . . . 142
8.5.3 Calibration of the gate detuning and duration . . . 142
8.6 G-G gates . . . 144
8.7 M-M gates . . . 144
8.8 G-M gates . . . 146
8.8.1 Experiment sequence . . . 148
8.8.2 Gate calibration . . . 148
8.8.3 G-M gate results . . . 150
8.9 Discussion . . . 150
8.10 Summary . . . 155

9 Outlook **156**

Appendices

A Barium clock qubits **169**
A.1 Metastable level clock qubits . . . 169
A.2 Ground level clock qubits . . . 171

List of Figures **172**

List of Tables **182**

List of Abbreviations **183**

In this sense the large-scale quantum machine, though it may be the computer scientist's dream, is the experimenter's nightmare.

— Serge Haroche and Jean-Michel Raimond

1 Introduction

Contents

1.1 Background 1
1.2 This work 5
1.3 Thesis outline 7

All the work in this thesis is driven by the desire to build a *quantum computer*. Quantum computing uses the full state space of nature, which is quantum, to perform computations, compared to the vastly reduced state space inhabited by classical computers. It is, hence, the most powerful form of computing allowed by the laws of physics.

1.1 Background

The fact that quantum systems have a computational advantage over classical systems was first observed by Feynman in the 1980s. He first discussed classical and quantum mechanical methods for simulating physical systems in his keynote speech at MIT in 1981 [8]. He postulated that a *useful* simulation of a physical system is one that satisfies the following rules:

- simulates the behaviour of the system *exactly*, as it would happen in nature;

- requires a number of computing elements proportional to the physical system's space-time volume; that is, there is no explosion in the required amount of computational resources.

As Feynman explained, the physical world is quantum mechanical; the classical behaviour of a given system is merely an approximation of its quantum-mechanical behaviour. Two critical phenomena in quantum mechanical systems set them apart from classical systems: quantum superposition and entanglement. There are a couple of ways to model quantum systems using a conventional (classical) computer. One could work in the classical approximation, describe the system using a set of differential equations, and use numerical methods to solve these equations. However, this method is not exact and breaks the first rule. Alternatively, one could write down the Hamiltonian describing the quantum system's behaviour and numerically calculate all probability amplitudes of the system as it evolves under the given Hamiltonian. However, to describe a quantum system of N particles, each of which has p states, would require p^N classical numbers. The exponential scaling results in resource explosion, hence breaking the second rule. This ruled out classical computers as a possible tool for universal exact simulation of quantum-mechanical systems. Feynman finished his talk with, *"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."* In other words, he suggested that the only way to exactly simulate nature, which is quantum, is by building a completely new type of computer, a *quantum computer*, whose basic computational unit, the *qubit*, itself exhibits quantum behaviour.

In his speech, Feynman only discussed quantum computing in the context of simulating physical systems. However, since then, quantum computing has found applications in mathematics [9], cryptography [10], machine learning [11], and direct applications of physical simulations, such as new material and drug discovery [12]. These theoretical developments prompted experimental efforts to

build this new type of machine. However, as Feynman said more than 40 years ago, this is a formidable task. To build a quantum computer, we need a scalable quantum mechanical system, which we can initialise in a well-known state, for example, $|000...\rangle$. We further need to be able to implement all necessary quantum logic gates on the system to form a 'universal' gate set. Additionally, we must have a way of measuring the states of the qubits at the end of the computation. Finally, we must ensure that the system is sufficiently well isolated from the environment so it can store quantum information over timescales much longer than the gate durations (known as the system's coherence time). These requirements are codified as DiVincenzo's criteria for building a quantum computer [13].

In summary, to build a quantum computer, we need a quantum system that we can isolate from the environment very well during operation but also interact with strongly for state preparation and measurement; these are two completely opposing requirements. As Serge Haroche and Jean-Michel Raimond framed it, *"... the large-scale quantum machine, though it may be the computer scientist's dream, is the experimenter's nightmare."* [14].

To date, many physical systems have been proposed as platforms for quantum computing that can, in principle, satisfy DiVincenzo's criteria, including trapped ions [15], neutral atoms [16], superconductors [17], NV-centres [18], and spins in silicon [19]. On a small scale, these systems have already been used in proof-of-principle demonstrations of several quantum computing algorithms, for example, a machine learning classification algorithm implemented on an 11-qubit trapped ion device [20] or a variational quantum eigensolver implemented on a 7-qubit superconducting device [21]. Therefore, the question we ask today is not whether we can build a quantum computer in principle but rather what tools we need to build a large-scale quantum computer.

How do we choose a size scale to aim for? In the long term, we want tens of thousands of very low-error qubits. However, this is a while away, and it's challenging to work out the optimum tools to get there right now. Another exciting threshold to reach is where we can start evaluating quantum algorithms

of large enough sizes that we would never be able to run them on a classical computer. This 'quantum discovery' threshold is set by the capabilities of current classical supercomputers. Current simulation tools can perform full simulation[1] of a perfect quantum system containing up to ~ 50 qubits [22, 23].

In addition to a large qubit count, we also need the error rates in the system to be sufficiently small. The error rates limit the circuit depth that can be executed before the output of the computation no longer has any statistical significance. Several quantum error correction (QEC) schemes have been developed to significantly reduce errors incurred during a quantum computation regardless of the error source, resulting in what is known as fault-tolerant quantum computing (FTQC) [24–26]. These error correction schemes work by treating several *physical qubits* as a single *logical qubit*. Entanglement between the physical qubits and measurements performed on some of the physical qubits during the circuit execution can be used to flag when an error occurs on the logical qubit and correct that error. However, this causes a significant overhead in the required size of the physical system: error correction schemes require anywhere between 3 and $10,000$ physical qubits per logical qubit, depending on the error sources, the physical error rates, and the target logical error rates. Ensuring that the error rates in the system are as low as possible minimises this overhead.

Of course, we cannot build a large-scale FTQC overnight. So, how do we build a system that provides near-term computational advantage while not hindering progress towards a large-scale fault-tolerant quantum computer? There is a regime in which quantum computers have enough qubits, say $50 - 500$, that are not error-corrected but have low enough error rates to facilitate the implementation of useful quantum algorithms. In 2018, John Preskill termed this regime NISQ to stand for noisy intermediate-scale quantum computers [27]. Such devices will allow us to explore new avenues of science that are not currently accessible to

[1] Here, we only consider a classical simulator of a quantum computer that calculates how each classical probability amplitude of the system evolves under the application of the input quantum circuit. That is, a classical simulator that represents the behaviour of a generic quantum computer. We are explicitly neglecting simulators optimised for very specific applications and, as a result, could handle more qubits.

classical computers, such as simulating molecular systems or enhancing some machine learning algorithms [28]. It should be noted that a NISQ device is not an FTQC; it is merely an intermediate stepping stone towards building bigger and more powerful devices while providing a computational advantage over classical systems as we go along.

1.2 This work

The main goal of the work presented in this thesis is to develop and experimentally demonstrate techniques that overcome some of the roadblocks to scaling trapped-ion systems. Trapped ions are one of the leading platforms for building a quantum computer. Small-scale trapped-ion devices have demonstrated the longest coherence times [1] and the highest single- and two-qubit gate fidelities of any quantum computing platform [2–5]. The challenge is finding ways to add more qubits to the system while maintaining this level of performance.

Both working with more qubits and running longer circuits compound errors in the system encountered on a small scale. For example, at the start of each sequence, the qubit register needs to be initialised in a well-known state, say $|000...\rangle$. Suppose the error with which we can initialise a single ion in the state $|0\rangle$ is ϵ. Then, the error with which we can initialise N ions in the state $|000...\rangle$ is $\approx N\epsilon$ for $\epsilon \ll 1$. Hence, if for a single qubit, we have a 0.1% state preparation error, then we would have a 1% error for a 10-qubit register and a 10% error for a 100-qubit register.

A second example is the heating of the ions. At the start of each sequence, the trapped ions are usually cooled down to their motional ground state. However, once cooling is complete and turned off for information processing, the ions start to thermalise with the environment. The temperature of the ions alters their interaction with the classical controls used to perform quantum logic gates, leading to errors. With no mitigation (e.g. sympathetic cooling during the sequence), the longer the sequence, the more the ions will heat up, resulting in larger gate errors.

The above two examples show that simply developing hardware to trap more ions is insufficient to build a NISQ device or any other form of large-scale FTQC. In fact, we have known how to trap thousands of ions for years, see e.g. [29, 30]. The challenge is in controlling large-scale systems with high precision and all-to-all connectivity to achieve true quantum computational advantage.

This thesis focuses on addressing precisely these problems. We have designed and implemented an architecture that combines existing methods for controlling trapped ions with novel tools and techniques to reduce the rate at which errors accumulate in large-scale systems. Our architecture replaces two atomic species (one for logic operations and one for mid-circuit cooling and measurement) with two different encodings in a single species, which significantly simplifies the complexity of the experimental apparatus. The ability to cool the ions during a sequence solves the issue of ion heating during long circuits we explained above. The ability to measure the state of some ions during a sequence can be used in error mitigation techniques to detect and reduce errors in the system further. We used this capability in a new protocol for qubit state preparation and measurement (SPAM). Implementing this protocol in our system resulted in the lowest SPAM errors recorded in any quantum computing platform. Additionally, we experimentally characterised the performance of the qubit encoding chosen in this work and implemented the types of two-qubit gates required to build a universal set and implement quantum non-demolition (QND) measurements in the system.

Therefore, we have designed and implemented a novel system architecture and demonstrated all necessary fundamental operations. As the first DPhil student on this project, I took charge of designing and setting up most of the experiment apparatus and implementing the techniques outlined above. In Chapter 9, I will discuss how the methods developed over the course of this work naturally fit together to enable scaling up the system.

1.3 Thesis outline

This thesis is structured as follows.

Chapter 2: Architecture gives an overview of the intended system architecture.

Chapter 3: The barium ion motivates the choice of barium for this work, provides a detailed overview of the barium level structure, and explains how all fundamental operations are implemented.

Chapter 4: Apparatus describes the design of the experimental apparatus used in this work.

Chapter 5: Individual addressing system presents the design and characterisation of a novel low cross-talk individual optical addressing system.

Chapter 6: High-fidelity heralded state preparation and measurement describes a novel SPAM protocol, which uses several mid-circuit measurements to herald successful SPAM. Using this protocol, we demonstrated the lowest-ever recorded SPAM error of any quantum computing platform.

Chapter 7: Exploring a metastable level clock qubit characterises the properties of one of the metastable level barium clock qubits and its isolation from mid-circuit measurements and cooling in the system.

Chapter 8: Two-qubit gates demonstrates the implementation of two-qubit gates between different qubit encodings within barium.

Chapter 9: Outlook summarises the results of the work presented in this thesis and discusses how they can be used in building large-scale devices.

The content of Chapters 5, 6, 7, and 8 of this thesis will result in three publications that I am lead author on.

Barbie is all these women. And all these women are Barbie.

— Barbie Movie Narrator

2 Architecture

In this chapter, we describe an architecture design that aims to address some of the challenges of scaling trapped-ion quantum computers. The content discussed here is separate from the remainder of this thesis, which focuses on experimentally demonstrating the fundamental building blocks of this architecture.

Any quantum computer must fulfil the fundamental requirements outlined by DiVincenzo in Ref. [13]. These requirements are as follows:

1. A scalable physical system with well-characterised qubits.
2. The ability to initialise the state of the qubits to a simple fiducial state, such as $|000...\rangle$.
3. Long relevant decoherence times, much longer than the gate operation time.
4. A 'universal' set of quantum gates.
5. A qubit-specific measurement capability.

On a small scale, atomic ions have already proven to satisfy DiVincenzo's criteria listed above [15]. The qubits are encoded in a pair of long-lived states within the ion's level structure, and the system can be scaled up by simply adding more ions to it, satisfying the first requirement. The ions can be prepared into

a well-known state by applying a set of lasers that pump the population out of all but one state within their level structure, satisfying the second criterion. Similarly, the ion's quantum state can be measured by applying a laser onto the ion that couples only one of the qubit states, say $|0\rangle$, to a short-lived excited state. The number of photons scattered as a result of the decay of this excited state can be used to infer whether the ion was in the state $|0\rangle$ (presence of photons) or $|1\rangle$ (lack of photons), hence satisfying the fifth criterion.

The lifetime of the atomic levels typically used to encode qubits is much longer than the gate operation time (usually seconds-years versus microseconds, respectively), enabling a long system coherence time generally limited by classical controls. Trapped ions have demonstrated the longest coherence times of any quantum computing platform. The record at the time of writing is > 1 hour [1]. Therefore, they satisfy the third criterion as well.

To build a universal set of gates, we need a set of single-qubit gates together with a pairwise entangling gate. Single-qubit gates are implemented using electromagnetic radiation to drive the transition between the two qubit states. Two-qubit gates are implemented by applying a spin-dependent force on the ions to excite their collective motion [31], which is coupled via their Coulomb repulsion, except for the case of Rydberg ions where the dipole-dipole interaction of ions in highly excited states is used instead [32]. Therefore, trapped ions satisfy the fourth DiVincenzo criterion as well. Trapped-ion qubits have demonstrated the lowest single- and two-qubit gate errors to date [2–5].

To build a functional large-scale quantum computer out of these primitives, we have to take further considerations into account. First, the system must be algorithm-agnostic to implement any generic quantum algorithm without limitations, for example, qubit connectivity. Second, as we discussed in Chapter 1, in large-scale systems, the errors from the different qubit operations add up, resulting in a higher net error of the computation. Furthermore, some errors, such as those due to ion heating, increase as a function of time during the circuit. We must, therefore, optimise the way in which errors add up in the system.

There are techniques we can build into quantum computers that aid in reducing the addition of errors with the number of qubits or the circuit depth. The first is to perform mid-circuit measurements on selected qubits (the ancilla qubits) while preserving the state of the remaining qubits (the computational qubits). The outcomes of these measurements can be used to detect errors that occur during the circuit and remove the corresponding shots by post-selection or convert them into erasures, one of the easiest kinds of errors to correct [33, 34]. Additionally, mid-circuit measurements are a core requirement for the implementation of QEC [24, 25]. They are also used in several types of quantum algorithms, for example, feed-forward circuits in quantum machine learning [35]. Hence, in addition to reducing system errors, this capability expands the range of practical applications of the quantum computer.

In the context of trapped-ion systems, another common source of error is the temperature of the ions. The ion temperature changes the interaction between the ion and the field used to perform quantum operations, leading to errors. This is why the ions are typically cooled down at the start of each sequence. However, during the sequence, the ions heat up as they start to thermalise with the environment. As we implement longer sequences, the ions heat up more, resulting in larger errors. This effect can be significantly reduced by cooling some of the ions during the sequence, which sympathetically cools down the entire register [36].

The challenge in performing mid-circuit measurements and cooling is that they are dissipative operations, meaning that they destroy any information stored in the qubit. We must, therefore, find a way to isolate the computational qubits, whose state we wish to preserve, from the light fields used to perform cooling and measurement on the ancilla qubits.

One of the most common approaches is to use two different ion species in the system, for example, both calcium (Ca) and strontium (Sr) ions [7, 37, 38]. Since they have very different transition wavelengths, they are insensitive to each other's light fields. For example, cooling or measurement of Sr in a Ca-Sr crystal does not cause population transfer within Ca and vice versa.

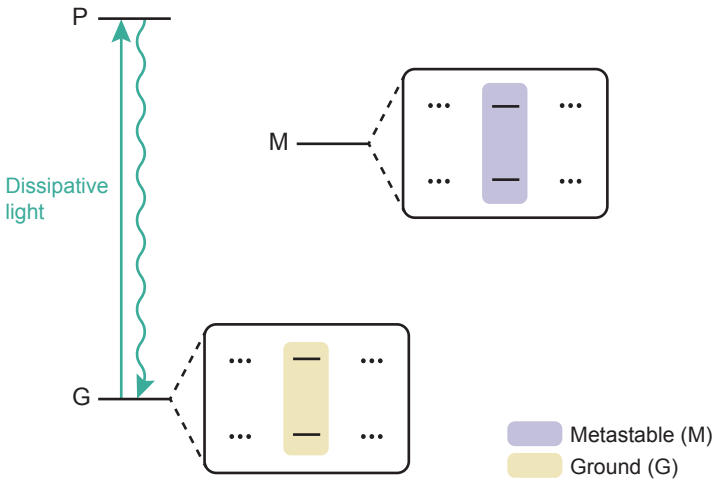


Figure 2.1: Fundamental concept of our architecture [6]. We encode qubits within two level manifolds of the ion: the ground level (G) and the metastable level (M). The laser light used to perform cooling, readout, and qubit reset (dissipative light) does not couple to the M level. Thus, we use the M level qubits as our computational qubits and the G level qubits as our ancilla qubits.

However, building systems using several ion species is challenging. They add significant experimental overhead in the classical control hardware required to run the experiment, including laser sources, laser light modulation components, and control electronics. They also make a lot of operations much harder to implement. Due to the different masses of different types of ions, they have unequal participation in the crystal's motional mode structure. This makes sympathetic cooling in such systems much less efficient, performing two-qubit gates with high fidelity more complex, and implementing the transport operations required in large-scale devices with low ion heating much more involved. These difficulties limit the scalability of mixed-species systems.

The architecture we are developing here takes a different approach. Rather than using extra elements, it uses the extra state space available within a single element. We make use of the fact that the ion has a long-lived level manifold, a metastable level (M), in addition to the ground level (G), within which we can store information. This concept is known as the *omg* architecture [6] and is shown in Fig. 2.1.

In trapped-ion systems, both cooling and measurement are performed by

continuously exciting the transition between the ion's G level and another short-lived level (P) using a laser, resulting in the scattering of photons. Crucially, this laser only couples to the G level and not to the M level. As a result, information encoded within the M level is protected during cooling and measurement.

Hence, we can build a system out of a single ion species that is equivalent to a mixed-species system by encoding the computational qubits within the M level, and the ancilla qubits within the G level. Conversion between the different qubit types is performed by simply driving transitions between the G and the M levels using laser pulses. The equivalent of the mixed-species gate in this architecture is the implementation of a direct entangling gate between a G qubit and an M qubit.

Using the same type of ion in the system removes the main disadvantages of working with multiple ion species and significantly simplifies the way in which operations are implemented. The equal ion mass makes sympathetic cooling, transport, and performing high-fidelity entangling gates much easier. Additionally, the qubit encoding in each ion is defined by a sequence of laser pulses. Thus, it can be changed in real-time without the need to completely reconfigure the ion register. For example, consider a register of 100 ions. We can easily switch from a configuration in which 50 are defined as ancilla qubits (in the G level) and 50 are defined as computational qubits (in the M level) to a configuration in which there are 20 ancilla qubits and 80 computational qubits by simply changing the laser pulses we apply to the ions, which can be done from software. This capability means that we can also change the qubit encoding in each ion on the fly during a circuit, allowing us to reuse qubits, which adds additional degrees of freedom in algorithm design.

Individual components of this architecture have previously been demonstrated, for example, the protection of a metastable level qubit from dissipative light [39] or the implementation of two-qubit entangling gates between a pair of metastable level qubits [40]. The work in this thesis is focused on developing and experimentally demonstrating all building blocks of this architecture in a way that can be used to build large-scale quantum computers using multiple

qubit encodings within the same ion species. All experimental demonstrations in this thesis are performed using $^{137}Ba^{+}$ ions. These ions are ideally suited to this architecture due to the long lifetime of the metastable level, making it an ideal place to encode qubits for both quantum information processing and storage.

We did a detailed study of the metastable level structure of $^{137}Ba^{+}$ and found several pairs of states to be excellent candidates for qubit encoding. We implemented one of these possibilities and evaluated the coherence time of the metastable level qubit and its isolation from dissipative light. We further demonstrated a two-qubit entangling gate between two M qubits (M-M gate) and, for the first time, between an M qubit and a G qubit (G-M gate). The M-M gates in our architecture are necessary to implement a complete universal set of unitary gates. The G-M gates are the equivalent of the mixed-species gates and, therefore, facilitate the implementation of quantum non-demolition (QND) measurements for QEC [41, 42] directly without the necessity for changing the ancilla or the computational qubit encoding for the QND. This makes the performance of our system dependent only on the qubit operations themselves and not on the quality of the laser pulses used to initialise the qubit in the correct level. We further showed how the ability to perform mid-circuit measurements can be used to reduce operation errors by developing a new SPAM protocol, which allowed us to achieve the lowest SPAM errors ever recorded.

Combining these techniques makes our system equivalent to a mixed-species system. The techniques we have developed and demonstrated can be used in building large-scale systems, and we discuss this further in Chapter 9.

A day without a friend is like a pot without a single drop of honey left inside.

— Winnie the Pooh

3 The barium ion

Contents

3.1 Choice of ion 15
3.1.1 Choice of isotope 16
3.2 Photoionisation 18
3.3 Ion operations 19
3.3.1 Ion fluorescence 20
3.3.2 Doppler cooling 21
3.3.3 State preparation 22
3.3.4 Detection 23
3.3.5 Sideband cooling 25
3.3.6 Quadrupole transitions 25
3.4 Qubit encoding 26
3.4.1 Metastable level structure 26
3.4.2 Ground level structure 28
3.4.3 Qubit transitions 30

Most trapped-ion qubits are implemented in atomic species with a single electron in the outermost shell upon ionisation, resulting in a hydrogen-like configuration. This leads to the simplest possible level structure, making fundamental operations much easier to carry out than in ions with more valence electrons. The atomic species meeting this criterion include group 2 alkaline earth metals, group 12 metals, and certain lanthanides.

With few exceptions, these ions share a similar gross level structure outlined

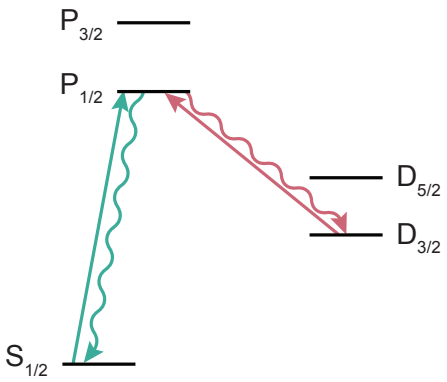


Figure 3.1: Gross level structure of most ions used for ion trap quantum computing. The ground level is $S_{1/2}$. There are two short-lived levels, $P_{1/2}$ and $P_{3/2}$, and two long-lived levels, $S_{1/2}$ and $D_{5/2}$ (lifetime of ~ 1 s). The *D*-levels do not exist in beryllium and magnesium. In ytterbium, they are very short-lived, and the $F_{7/2}$ level is used in cases where a long-lived level is necessary in addition to the ground level.

in Fig. 3.1. Ground level qubits, as defined in Chapter 2, are encoded within the $S_{1/2}$ level. Measurement and Doppler cooling are implemented by driving transitions between the ground $S_{1/2}$ level and the $P_{1/2}$ level. Some of the population excited to the $P_{1/2}$ level can decay to the long-lived $D_{3/2}$ level, and an additional laser is used to pump it back to the $P_{1/2}$ level. However, the $D_{5/2}$ (or $F_{7/2}$ in the case of ytterbium) level is decoupled from this cycle, making it a suitable place to store qubits in addition to the ground level, as outlined in Chapter 2.

While most species used as trapped-ion qubits share the gross level structure outlined above and all fundamental operations are implemented analogously, different ions have different masses, transition frequencies, and lifetimes of their metastable levels, and some ions exhibit hyperfine structure. These differences guide the choice of ion for a given architecture. In this chapter, we motivate our choice of ion for our apparatus and describe how we implement all fundamental operations. All atomic physics simulations in this chapter were performed using [43].

3.1 Choice of ion

As outlined in Chapter 2, the choice of ion for this work is primarily based on the capability to store the qubit states in the ion's metastable level, while utilising the ground level for carrying out mid-circuit dissipative operations. Hence, we

need an ion with an easily accessible metastable level and a long metastable level lifetime. Furthermore, to minimise the system complexity, we need an ion with transition frequencies that are relatively straightforward to generate and manipulate using optical components. Generally speaking, we prefer visible and/or infrared transitions over near-UV or UV transitions, since both high-power laser sources and optical components with high quality are much more difficult to manufacture for shorter wavelengths.

Table 3.1 presents the most commonly used atomic species for trapped-ion quantum information processing along with their respective properties. Of the ions presented in the table, barium and ytterbium have the longest metastable level lifetimes, of 30.14 s and days − years, respectively. However, the shortest wavelength needed for working with barium ions is 455 nm, close to the centre of the visible spectrum, compared to 369 nm for ytterbium. Thus, using barium would allow us to utilise off-the-shelf high-performance laser sources and components. Additionally, transferring population from the ion's ground $S_{1/2}$ level to the $F_{7/2}$ level in ytterbium requires multiple laser pulses driving narrow linewidth transitions [39], compared to the $D_{5/2}$ level in barium, which requires a single laser pulse. Finally, the metastable level lifetime of barium is already $\sim 10^7$ times longer than typical gate speeds in trapped ions. As we will see in Chapter 6, we can also detect and correct for its decay, thus enabling working with experiment sequences that are comparable in length to the metastable level lifetime. Hence, using ytterbium would not significantly improve the performance of the system as a quantum computer, but it would complicate the experimental setup. Therefore, we choose to work with barium ions. The gross level structure of barium is shown in Fig. 3.2.

3.1.1 Choice of isotope

The atomic properties of the most widely used barium isotopes are provided in Table 3.2. Several factors must be considered when selecting an isotope to work with. First, it is advantageous to choose an ion with a non-zero nuclear

Ion species	Metastable level lifetime (s)	Metastable level	Shortest wavelength (nm)
Be	N/A	N/A	313
Mg	N/A	N/A	279
Ca	1.2	$D_{5/2}$	393
Sr	0.39	$D_{5/2}$	408
Ba	30.14	$D_{5/2}$	455
Yb	days-years	$F_{7/2}$	369

Table 3.1: A table containing the most commonly used atom species for trapped-ion quantum information processing [6, 44–47]. The wavelengths listed correspond to the shortest wavelengths necessary for driving transitions between levels used for quantum information processing in the ions. They typically correspond to the $S_{1/2} \leftrightarrow P_{3/2}$ transitions. The two species with the longest lifetimes of the metastable level are barium and ytterbium. While the ytterbium metastable level lifetime is longer compared to barium, it requires work with UV lasers and several population transfer pulses to reach the metastable level from the ground level. Therefore, in this work we choose to work with barium ions.

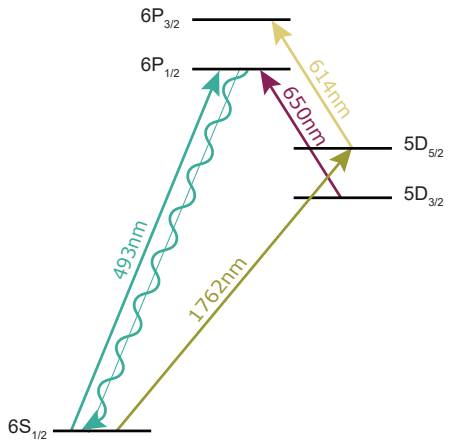


Figure 3.2: Gross level structure of a barium ion. All wavelengths we need to work with are in the visible range, with the exception of the 1762 nm laser, which is in the near-IR range.

spin, i.e. an ion that exhibits a hyperfine structure. The hyperfine structure results in pairs of states in the ion's level structure whose frequency splitting is first-order insensitive to external magnetic field variations, known as 'clock' qubits. Consequently, clock qubits are much less sensitive to magnetic field noise, which is the typical limit to the system coherence time, and can thus achieve longer coherence times in a given setup than non-clock qubits. Second, it is preferable to use ions with the smallest possible nuclear spin. The number

of states within a given level of the ion scales with its nuclear spin. A lower nuclear spin thus leads to a simpler level structure, making the implementation of essential operations, such as ion cooling or state preparation, easier. Lastly, we prefer isotopes with high natural abundance – this ensures a large amount of the desired isotope in the atomic source and hence eases the restrictions on the loading mechanisms that can be used.

Of the isotopes presented in Table 3.2, $^{138}Ba^{+}$ has the highest natural abundance of 71.7%. However, it has a nuclear spin of 0 and does not exhibit a hyperfine structure. Of the odd isotopes whose nuclear spin is non-zero, the one with the lowest nuclear spin of 1/2 is $^{133}Ba^{+}$. This isotope is not naturally abundant. It can, however, be created by neutron bombardment of ^{132}Ba. As such, it is weakly radioactive, with a half-life of 10.5 years [48]. This limits the size of the atomic source that can be used for loading to ~ 100 µg and hence necessitates loading via laser ablation of the sample. There have been several attempts to load barium ions via laser ablation[49, 50]. However, laser ablation is associated with issues such as a limited sample lifetime or trap charging after loading due to the charged particles produced during the ablation process. At the start of this project, we were aiming to use $^{133}Ba^{+}$. However, our ablation attempts ran into the exact problems listed above. After several months of ablation trials, we decided to turn to naturally abundant isotopes at the cost of a slightly higher nuclear spin.

The remaining two odd barium isotopes, $^{135}Ba^{+}$ and $^{137}Ba^{+}$, have analogous properties. We therefore choose to work with $^{137}Ba^{+}$ as it has a higher natural abundance compared to $^{135}Ba^{+}$.

3.2 Photoionisation

Before we can trap, we must ionise the neutral $^{137}Ba^{+}$ atoms. We do this using a two-step photoionisation process, as shown in Fig. 3.3. The two electrons in the outermost shell of barium are initially in the state $6s^2$. A 553 nm laser resonantly excites one of the electrons to the p-state. A subsequent 405 nm laser brings the p-electron to the continuum, leaving a singly-ionised barium ion. The 553 nm

Isotope	Natural abundance	Nuclear spin	Ground level hyperfine splitting
$^{138}Ba^{+}$	71.7%	0	N/A
$^{137}Ba^{+}$	11.2%	3/2	8.037 GHz
$^{136}Ba^{+}$	7.9%	0	N/A
$^{135}Ba^{+}$	6.6%	3/2	8 GHz
$^{134}Ba^{+}$	2.4%	0	N/A
$^{133}Ba^{+}$	0%	1/2	9.025 GHz
$^{132}Ba^{+}$	0.1%	0	N/A

Table 3.2: barium isotopes [48]. All of the isotopes listed in this table are naturally abundant except $^{133}Ba^{+}$, which is normally created by neutron bombardment of $^{132}Ba^{+}$.

laser has to be tuned on resonance with the $6s^2 \rightarrow 6s6p$ transition for the given isotope. This aids in isotope selectivity during ion loading. The wavelength of the 405 nm does not need to be set precisely – it needs to be short enough to promote the p-electron to the continuum but not too short to excite the s-electron.

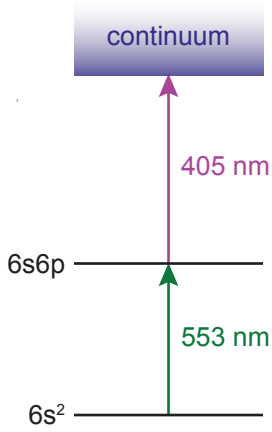


Figure 3.3: Two-step photoionisation of neutral barium atoms. The two valence electrons are initially in the state $6s^2$. A resonant 553 nm laser excites one of the electrons to the $6p$ state. This is followed by a 405 nm laser that brings the p-electron to the continuum.

3.3 Ion operations

A set of fundamental ion operations must be performed in order to enable quantum information processing. These include laser cooling, state preparation, and read-out. This section describes how we perform each of these operations in $^{137}Ba^{+}$.

3.3.1 Ion fluorescence

Three of the operations listed above (Doppler cooling, state preparation, and measurement) are implemented by continuously driving the cycling $S_{1/2} \leftrightarrow P_{1/2}$ transition using a 493 nm laser. Once excited, the electron rapidly decays back from the $P_{1/2}$ to the $S_{1/2}$ level with $\sim 70\%$ probability, resulting in 493 nm photons scattered from the ion, an effect known as fluorescence. The emitted photons are collected using a microscope objective and sent to an sCMOS camera (see Section 4.2 for details) used to detect the ions.

To address the transitions between states within all F-manifolds in $S_{1/2}$ and $P_{1/2}$, the 493 nm light is modulated using an electro-optic modulator (EOM). The laser carrier is tuned on resonance with the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. We then apply two microwave tones to the EOM: a tone at 8.037 GHz driving the $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions, and a tone at 6.56 GHz driving the $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 1\rangle$ transitions, as shown in Fig. 3.4. When applied simultaneously, these two tones also generate a sideband at 1.48 GHz driving the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 1\rangle$ transitions. The relative strengths of the microwave tones are fixed for all operations and set to generate approximately equal carrier and sideband strengths when all tones are on. Only the common-mode detuning and the total power of the 493 nm beam are tuned dynamically depending on the desired operation.

Approximately 30% of the population excited to the $P_{1/2}$ level decays to the $D_{3/2}$ level. Since the lifetime of the $D_{3/2}$ level is very long ($\approx$ 80 s) this can cause the ion to be stuck in the $D_{3/2}$ level. To repump the population that has decayed to the $D_{3/2}$ level, we apply a 650 nm laser that drives the $D_{3/2} \leftrightarrow P_{1/2}$ transitions. As with the 493 nm laser, we use an EOM to generate sidebands on the 650 nm light such that we are able to pump population out of all F-manifolds within $D_{3/2}$. However, contrary to the 493 nm laser, here we only apply a single tone to the EOM and use a high beam intensity ($\sim$ 1000 saturation intensities[1]) to power-broaden the transition and hence drive population out

[1]Saturation intensity I_{sat} definition as in Ref. [53].

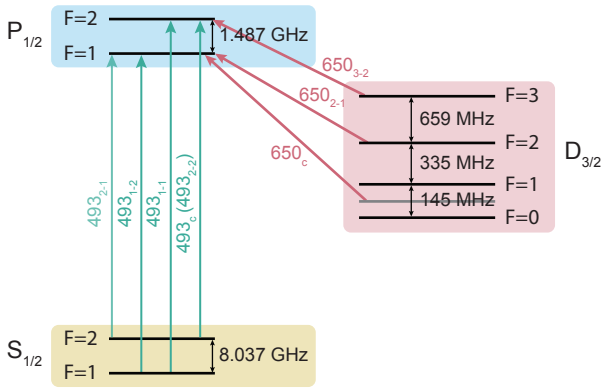


Figure 3.4: Sidebands used during operations using the ion's fluorescence cycle [51, 52]. The 493 nm laser carrier is tuned on resonance with the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. We apply two RF tones to an EOM, at 8.037 GHz and 6.56 GHz driving the $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ and the $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 1\rangle$ transitions, respectively. Due to sideband mixing in the EOM, these tones also generate a sideband at 1.48 GHz, which drives the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. The 650 nm laser carrier is set halfway in between transitions from $|D_{3/2}, F = 0\rangle$ and $|D_{3/2}, F = 1\rangle$ to $|P_{1/2}, F = 1\rangle$. We apply a single RF tone to an EOM at 430 MHz, which generates sidebands driving the $|D_{3/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 1\rangle$ and $|D_{3/2}, F = 3\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. We use ~ 1000 saturation intensities in the 650 nm beam such that we power-broaden the transition to address all states within $D_{3/2}$.

of all states in $D_{3/2}$. The laser carrier and the EOM frequencies used for the 650 nm beam are shown in Fig. 3.4.

3.3.2 Doppler cooling

During Doppler cooling [54–56] we apply a red-detuned 493 nm light to the ions. We turn on all EOM sidebands as described in Fig. 3.4 to pump population out of all states in $S_{1/2}$. Due to the Doppler shift, ions preferentially absorb photons whose direction is opposite to the ion's motion. Upon decay from $P_{1/2}$ to $S_{1/2}$, the ion scatters a photon whose frequency is closer to the resonance frequency of the transition. The difference between these two frequencies corresponds to the energy that is removed from the ion's motion. The cooling process continues until the energy removed per scattering event is equal to the energy gained due to the emitted photon's recoil. This steady-state temperature is referred to as the Doppler limit.

In our experiment, we use two sets of Doppler cooling parameters:

- 'Idle' or 'loading' Doppler cooling: high-intensity ($\approx 100 I_{\text{sat}}$[2]), high-detuning ($\approx -100$ MHz) cooling; these parameters are used during loading or idle ion storage in the trap. Due to the large detuning and power broadening of the lineshape, these parameters are able to span higher ion velocities encountered after photoionisation and before trapping as well as during collisions that significantly heat up the ions. However, the steady-state temperature of the ions as a result of this cooling is much higher than the Doppler limit. This is due to the much higher photon scattering rate at these laser parameters, resulting in a higher heating due to the scattered photons, and hence a higher steady-state temperature.

- 'Active' Doppler cooling: this is used to bring the ion's temperature from the steady-state temperature after idle cooling to approximately the Doppler limit. For this type of cooling we use $\approx 0.5 I_{\text{sat}}$ and a detuning of $\approx \Gamma/2$ where Γ is the FWHM of the 493 nm transition.

3.3.3 State preparation

At the start of an experiment sequence, the ions need to be prepared in a well-known state. To achieve this, we use the fact that at a low quantisation field (such that $\mu_B B \ll 8$ GHz where B is the quantisation field and μ_B is the Bohr magneton), transitions between $m_F = 0$ states with the same F quantum number are forbidden [57]. Therefore, we are able to prepare the $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ state by applying our purely π-polarised 493 nm beam and only turning on the sidebands driving the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ and $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions as shown in Fig. 3.5.

The fidelity of this process is limited by the polarisation purity of the 493 nm beam (typically limited by the angle between the beam and the external magnetic field), the suppression of transitions between $m_F = 0$ states at the given field, and off-resonant excitation of the $F = 2$ states in $\mathrm{S}_{1/2}$ through $F = 1$ states in $\mathrm{P}_{1/2}$. Off-resonant excitation can be minimised by using a low power of the

[2]Saturation intensity I_{sat} definition as in Ref. [53].

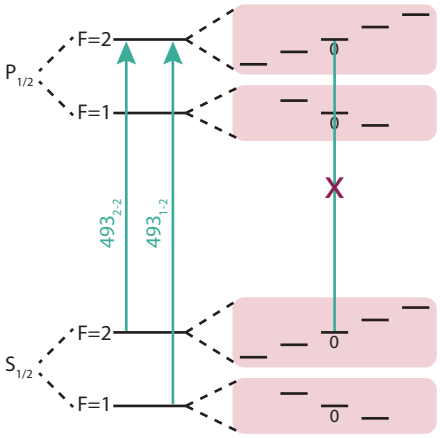


Figure 3.5: Preparation into a well-known state within $S_{1/2}$ using optical pumping. We turn on the 493 nm laser, turning on only the sidebands resonant with the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ and $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. Since the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ transition is forbidden, this process results in the majority of the population being pumped in the $|S_{1/2}, F = 2, m_F = 0\rangle$ state.

493 nm beam, and the polarisation purity of the beam can be optimised with good alignment. However, the suppression of transitions between $m_F = 0$ states for $\Delta F = 0$ at a given field poses a limit to the quality with which this process can be implemented. This is illustrated in Fig. 3.6.

In this work, we designed a new protocol that overcomes this limitation by harnessing the availability of a metastable level and performing a mid-circuit measurement to herald the success of the process. This protocol is described in detail in Chapter 6.

3.3.4 Detection

We can infer the presence of population within the $S_{1/2}$ level by detecting the 493 nm photons scattered by the ion. However, the 493 nm light does not couple to the metastable $D_{5/2}$ level. Before state detection, we ensure that one of the qubit states is encoded within $S_{1/2}$ and the other qubit state is encoded within $D_{5/2}$, a technique known as electron shelving[58, 59]. This is realised using population transfer between states within $S_{1/2}$ and states within $D_{5/2}$ using a 1762 nm laser resonant with the desired transition (see Section 3.3.6). The fidelity of the state detection process is then limited by the fidelity of the population transfer pulses,

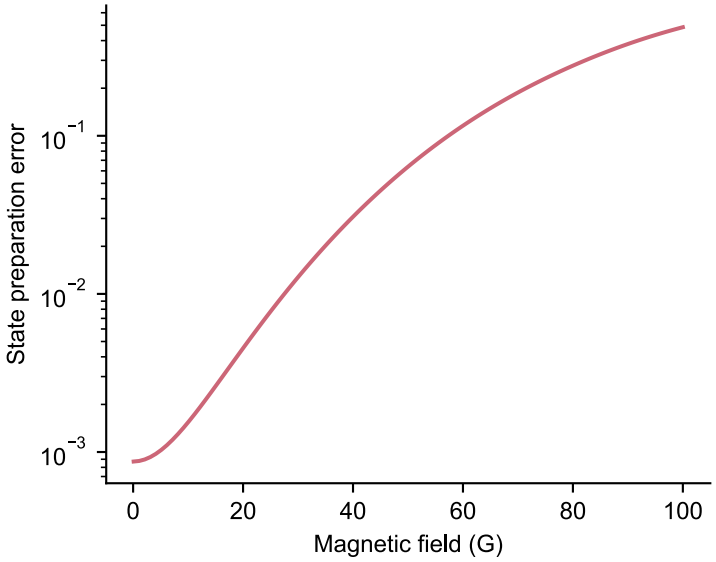


Figure 3.6: A rate equations simulation of the state preparation error in the $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ state as a function of the applied magnetic field. In this simulation, we use $1I_{\mathrm{sat}}$ of 493 nm light and $1000I_{\mathrm{sat}}$ of 650 nm light. We assume pure π-polarisation of the 493 nm beam. The 493 nm sidebands used are as shown in Fig. 3.5, and the 650 nm sidebands used are as shown in Fig. 3.4. We assume we have equal powers in each sideband (for both the 493 nm and the 650 nm beams). The state preparation duration used is 100 ms: this is much longer than necessary in practice; we used this duration for the simulation only to ensure that we have reached a steady state. As the magnetic field increases, the matrix element of the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ increases as well, and we enter the intermediate-field regime [57]. Hence, F is no longer a good quantum number, resulting in changes in the transition selection rules. The increase of this matrix element with magnetic field poses a fundamental limit on the fidelity of the state preparation scheme described in Section 3.3.3. The minimum state preparation error of $\sim 0.1\%$ is due to off-resonant excitation of the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 1, m_F = 0\rangle$ transition.

the discrimination between the presence and lack of fluorescence from the ion, and the decay of the metastable level during the fluorescence detection window.

The error due to the decay of the $\mathrm{D}_{5/2}$ level during the fluorescence detection window is approximately linear with the detection duration. To minimise the readout duration, we need to maximise the photon scattering rate. Therefore, we choose to work with $\sim 100I_{\mathrm{sat}}$ of 493 nm light, with only a small ≈ -5 MHz detuning relative to the transition resonance to ensure we do not heat the ion too much during this process, which could result in ion loss. The readout duration is then chosen as a trade-off between the error due to bright/dark discrimination and the error due to the decay of the metastable level and a detailed discussion

is provided in Section 6.3.1.

As with state preparation, the readout process can be further enhanced by performing a series of mid-circuit measurements to herald the successful transfer of the population to and from the metastable level. For further details on the state preparation and measurement (SPAM) procedure, see the SPAM protocol in Chapter 6.

3.3.5 Sideband cooling

We can reduce the ion's temperature below the Doppler limit where necessary (e.g. prior to two-qubit gates) by using sideband cooling. Sideband cooling is possible for transitions for which the FWHM Γ is much smaller than the ion's motional mode frequencies ω_i, i.e. $\Gamma \ll \omega_i$, such that the motional sidebands can be resolved. Then the technique consists of state preparation followed by driving the red sideband (RSB) transition for all motional modes that must be cooled. This technique is discussed in detail in Ref. [60]. The efficiency of the process (i.e. the number of phonons removed per cycle) is limited by any heating during state preparation between consecutive pulses as well as the trap heating rate.

In this experiment, we perform sideband cooling by driving the motional mode sidebands of the ground level qubit using a two-photon Raman process (see Section 3.4.3).

3.3.6 Quadrupole transitions

We must be able to transfer population between states in the $S_{1/2}$ level and states in the $D_{5/2}$ level for readout and for state preparation of the metastable level qubit. Preparation of the metastable level qubit is implemented by first preparing the state $|S_{1/2}, F = 2, m_F = 0\rangle$ as described in Section 3.3.3, followed by population transfer from this state to the desired qubit state in the $D_{5/2}$ level. In addition, the detection procedure described in Section 3.3.4 requires encoding one of the qubit states in $S_{1/2}$ and the other state in $D_{5/2}$ before readout. In

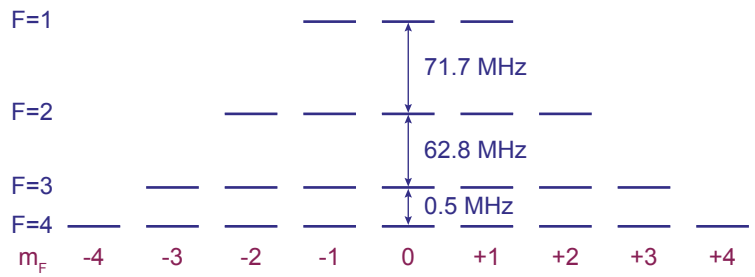


Figure 3.7: Hyperfine splitting of the metastable level of $^{137}Ba^{+}$ in the absence of an external magnetic field [52].

the case of a ground level qubit this requires population transfer from $S_{1/2}$ to $D_{5/2}$, and vice versa for a metastable level qubit.

In our experiment, we drive these quadrupole transitions directly using a resonant 1762 nm laser. We must be able to drive several different transitions, split by the ion's hyperfine structure, thus spanning several gigahertz in frequency. Therefore, the 1762 nm laser is frequency-modulated using an EOM to achieve the required frequency tuning range. The transitions we are able to drive then only depend on the beam geometry relative to the applied field. Details on this can be found in Chapter 4.

3.4 Qubit encoding

As described in Chapter 2, we will use qubits encoded in the metastable $D_{5/2}$ level in $^{137}Ba^{+}$ for quantum information processing and storage. The ground level qubit is only used for mid-circuit dissipative operations, such as cooling and measurement, during an algorithm, hence, its properties are not limiting to the system performance.

3.4.1 Metastable level structure

In the absence of an external magnetic field, the metastable $D_{5/2}$ level is split into four hyperfine manifolds with $F = 1, 2, 3, 4$ as shown in Fig. 3.7, and states with the same F are degenerate.

When an external magnetic field is applied to the ion, the degeneracy of the states within each F-manifold is lifted due to the Zeeman effect. Due to the

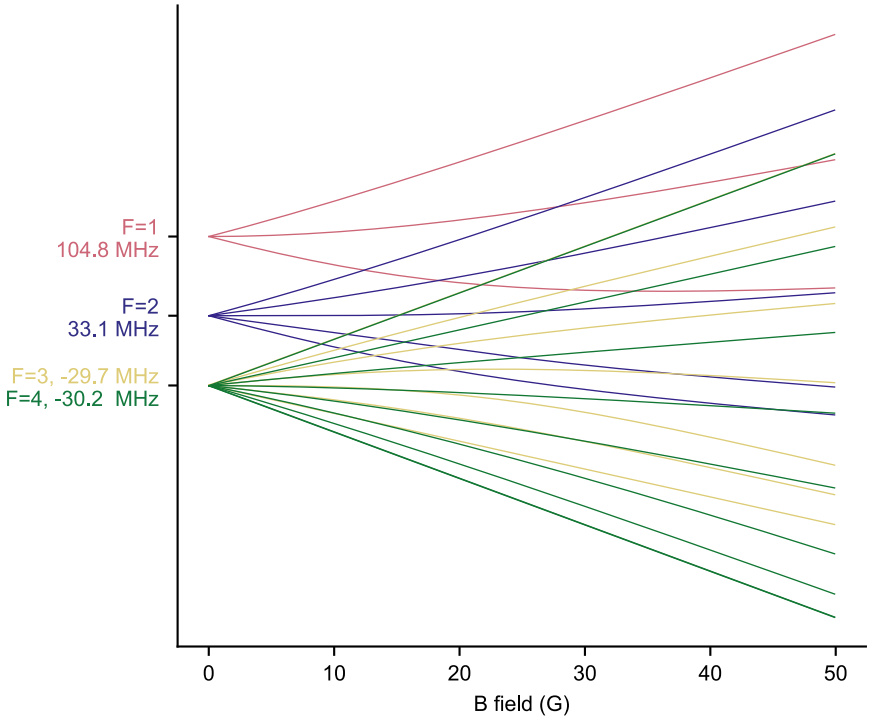


Figure 3.8: Evolution of the metastable level state energies as a function of the external magnetic field. The frequencies provided are relative to the $D_{5/2}$ state when only fine structure is considered.

very small hyperfine splitting in this level, even small magnetic fields ($\sim 1\,\mathrm{G}$) are enough to cause a Zeeman splitting comparable to, or larger than, the hyperfine splitting and hence F is no longer a good quantum number. For clarity, we will proceed to label the qubit states in terms of F and m_F anyways, denoting the quantum numbers they would have at zero field.

The frequencies of the states in the metastable $D_{5/2}$ level as a function of the external magnetic field are shown in Fig. 3.8. For some pairs of states, we can find a magnetic field at which their frequency splitting is first-order independent of small magnetic field changes, the so-called clock qubits mentioned earlier. An example of this behaviour is shown in Fig. 3.9. Since there are 24 states in the metastable level, there is a large number of clock qubits we can choose from, and a complete list is provided in Section A.1.

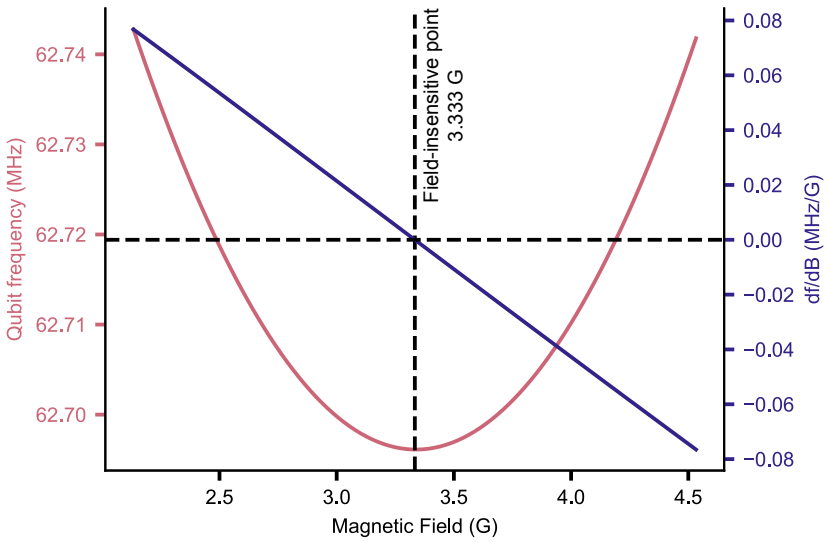


Figure 3.9: Evolution of the frequency splitting between the states $|F = 2, m_F = +2\rangle$ and $|F = 3, m_F = +2\rangle$ in the $D_{5/2}$ level of $^{137}Ba^+$. The frequency splitting reaches its minimum at $B \approx 3.333\,\mathrm{G}$. As can be seen from the blue curve, this point also corresponds to $df/dB = 0$, i.e. at this magnetic field, these states form a clock qubit.

3.4.2 Ground level structure

In the absence of an external quantisation field, the ground level of $^{137}Ba^+$ is split into two hyperfine manifolds, as shown in Fig. 3.10. The splitting between the two manifolds at zero quantisation field is $\Delta_{S_{1/2}} \approx 8.037$ GHz.



Figure 3.10: Hyperfine splitting of the ground $S_{1/2}$ level of $^{137}Ba^+$ in the absence of an external magnetic field [51].

At the low external magnetic fields required to work with a metastable clock qubit, the Zeeman splitting between states within the ground level is much smaller than the hyperfine splitting, and hence F remains a good quantum number. The evolution of the energies of the individual states within the ground $S_{1/2}$ level is shown in Fig. 3.11.

As with the metastable level, we can tune the quantisation field such that a given pair of states within the ground level becomes a clock qubit and a list of the available clocks is provided in Section A.2. However, since the required magnetic

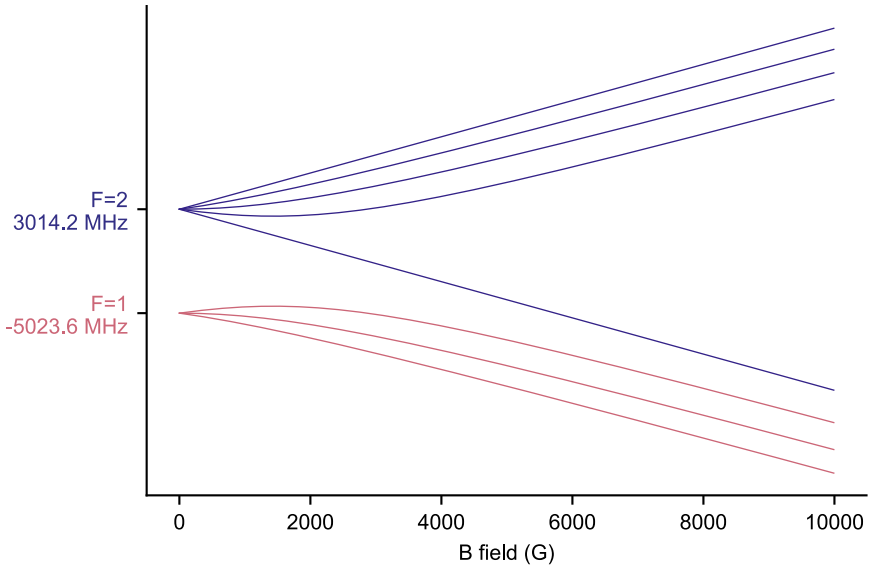


Figure 3.11: Evolution of the energies of the states within the ground $S_{1/2}$ level in $^{137}Ba^{+}$ as the applied quantisation field is increased. The frequencies provided are relative to the $S_{1/2}$ state when only fine structure is considered.

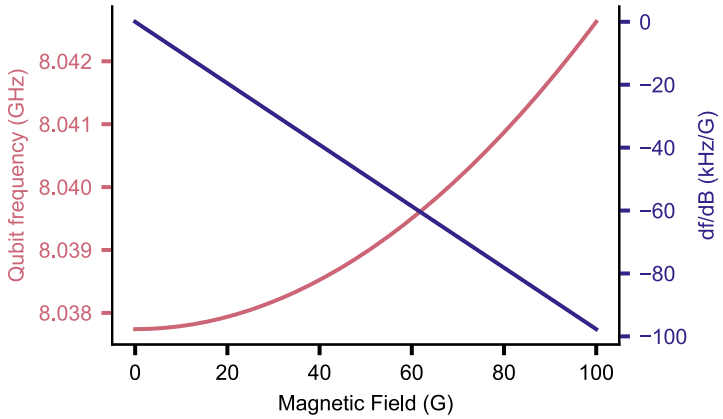


Figure 3.12: Evolution of the frequency splitting and the first-order magnetic field sensitivity of the $|F = 2, m_F = 0\rangle \leftrightarrow |F = 1, m_F = 0\rangle$ in the $S_{1/2}$ level. This pair of states forms a clock qubit at zero quantisation field. As the field is increased, this is no longer a perfect clock qubit and its first-order magnetic field sensitivity becomes non-zero. However, at the fields we want to work with (constrained by the metastable level), this is the ground level transition with the lowest field sensitivity.

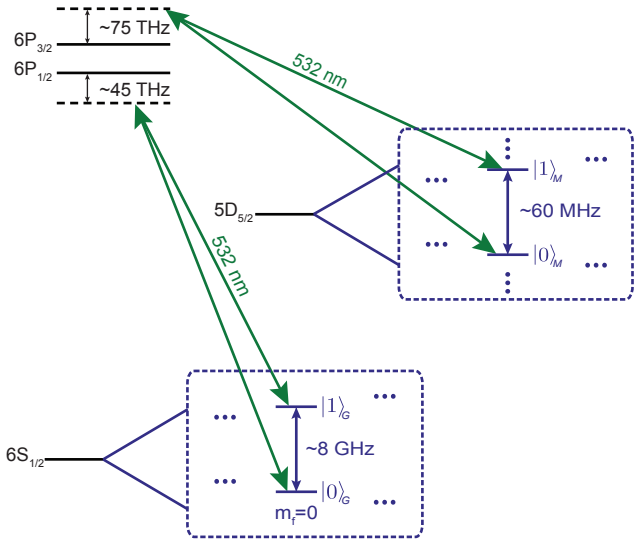


Figure 3.13: Raman transitions used to drive the two qubit types in $^{137}Ba^{+}$.

field is determined by the metastable level, here we choose to work with the qubit states $|F = 2, m_F = 0\rangle$ and $|F = 1, m_F = 0\rangle$. These two states form a clock qubit at zero magnetic field. As the field is increased, their frequency splitting and their first-order magnetic field sensitivity increase, as shown in Fig. 3.12. Still, at the low fields required for using clock qubits in the metastable level, this qubit has the lowest magnetic field sensitivity of all ground level qubits.

3.4.3 Qubit transitions

We drive both the ground and the metastable level qubit transitions using a two-photon Raman process. A unique feature of the level structure of barium is that the same laser wavelength at 532 nm drives Raman transitions for both the ground and metastable level qubits. That's because the detuning of the 532 nm beam relative to the $S_{1/2} \leftrightarrow P_{1/2}$ transitions is comparable to its detuning relative to the $D_{5/2} \leftrightarrow P_{3/2}$ transitions. This wavelength is at the centre of the visible range, and we can, therefore, use high-performance off-the-shelf laser and fibre components, and in particular, low-noise high-power lasers.

Details on the Raman laser system are provided in Section 4.3.4 and our novel individual addressing system is discussed in Chapter 5. The Raman beam

parameters (spot size, power, polarisation) vary throughout the thesis and are therefore discussed in detail in each chapter.

If you don't break things you aren't working fast enough.

— Dr Christopher James Ballance

4 Apparatus

Contents

4.1 The trap 33
4.1.1 Trap topology 33
4.1.2 Trap geometry 35
4.1.3 Trap signal delivery 37
4.1.4 Isotope-selective loading sequence 39
4.2 Imaging system 40
4.3 Laser systems 43
4.3.1 Photoionisation lasers 44
4.3.2 Ion operations lasers 45
4.3.3 Locking cavities 47
4.3.4 Raman lasers 48
4.3.5 Trap beam paths 52

In this chapter, we describe the design of the experiment apparatus. As we will see, many of the experimental tools we use are standard for confining and controlling trapped ions, for example, how we generate and manipulate the laser beams used for the fundamental ion operations described in Chapter 3. In such cases, we only briefly overview the exact apparatus used for the experiments in this thesis. This further shows that the main benefit of our architecture is not in the engineering of our classical controls but in how the qubits are encoded within the ion. In the rest of this chapter, we will only discuss in detail aspects

of the apparatus that are either non-trivial and specific to our experiment or are necessary for implementing the architecture primitives discussed in Chapter 2 and for understanding Chapters 5-8 of this thesis.

The main goal of the work described in this thesis is to design and implement techniques that help overcome some of the roadblocks in scaling trapped-ion quantum computers, as outlined in Chapter 2. This means that on the one hand, we want an experiment apparatus that would allow us to quickly get to a point where we can test these techniques. On the other hand, we also want to ensure that this apparatus can be used with a larger number of qubits in the near term. Therefore, the apparatus used for this thesis was designed to minimise the experimental complexity while also enabling near-term flexibility.

4.1 The trap

In this section, we discuss the choice of trap for our apparatus and describe the trap vacuum system and the signal delivery systems. The trap is chosen in a way that minimises the complexity of the setup while also enabling near-term scalability.

4.1.1 Trap topology

There are three main trap topologies that can be used to build a trapped-ion quantum computer:

- Confining many ions within the same trapping potential [59, 61].
- Using devices with multiple zones where each zone hosts only a few ions [59, 62]. Communication between zones is facilitated by shuttling ions across the chip. This is known as a quantum charge-coupled device, QCCD.
- Networking separate trapped-ion nodes using photonic interconnects [63–65].

The first approach, adding more ions to the same linear chain, has been proven to work with tens of ions using standard ion trap devices widely used in the field over the past 20 years. However, there is a limit to the number of ions that

can be stably confined and manipulated with high fidelity in a single chain [59]. The second approach of QCCDs overcomes this limit by only confining a small number of ions in any trap zone and facilitating connectivity between zones by shuttling ions across the device. There is no fundamental limit to the number of qubits that can be achieved using this approach. However, manufacturing devices capable of confining ions in several zones in parallel and transporting these ions across the chip requires complex manufacturing procedures and much more complicated classical control. Finally, networking several ion trap nodes changes the constraints in qubit count that can be achieved in either a linear trap or a QCCD architecture. However, each ion trap node in a network requires its own apparatus and set of classical controls (laser and electronic controls). Therefore, this approach is more suitable for overcoming limits in already large-scale systems rather than increasing the qubit count from, e.g., ten to fifty qubits, which would require five nodes.[1]

Here, we chose to work with ions confined in a single linear chain. This approach enables near-term scalability using standard techniques with minimal experimental complexity. In the long term, this approach can be combined with the QCCD approach by creating more trapping zones and controlling the qubits in each zone as in a linear trap.

Typically, ion traps come in one of two flavours. The first is three-dimensional mechanically assembled traps. These traps usually have a large ion-to-electrode distance and thus offer deep confining potentials while maintaining low heating rates. However, they have limited degrees of freedom for modifying the potentials experienced by the ions. The second is surface electrode traps. These traps have multiple electrodes and hence offer more degrees of freedom for generating the trapping potentials. However, the typical ion-electrode is $40 - 80$ µm [68, 69], leading to shallower confining potentials and higher heating rates.

In our apparatus, we work with a trap that takes the best from both worlds – we use a three-dimensional microfabricated segmented ion trap [70–73]. The

[1]Note that while quantum networking is not a suitable scaling approach for small increases in qubit count, it is fundamental to a lot of security protocols, see e.g. Ref. [66, 67].

three-dimensional layout and the large ion-to-electrode distance facilitate the generation of deep confining potentials. The segmented electrode structure provides additional degrees of freedom for modifying the confining potential. For example, it enables the generation of anharmonic potentials, which can be used to keep the ion spacing uniform across the chain. It also allows for ion shuttling operations, which can improve our loading procedures, as discussed in Section 4.1.4, or be used to build a small QCCD.

The trap vacuum chamber housing is based on the one presented in Ref. [74]. We made small modifications to the vacuum chamber to allow beam access at $> 45°$ and using an objective with a smaller working distance, therefore enabling fluorescence collection from a larger solid angle (see Section 4.2 for details). To create vacuum, we use the pumping assembly designed in Ref. [63]. The assembled vacuum system is shown in Fig. 4.1.

4.1.2 Trap geometry

The trap chip used in this experiment is designed and manufactured by the National Physical Laboratory in the UK. Details on the design and fabrication of the chip can be found in Ref. [70–73]. This section provides an overview of the trap geometry relevant to understanding the rest of this chapter.

The layout of the ion trap is shown in Fig. 4.2. The ion-to-electrode distance is 247.5 µm. The trap can be logically split up into four types of zones:

- Trap operation zones where ions are stored during experiments.
- A loading zone. This is the only place atomic flux from the oven can get to due to shielding fitted within the trap chamber. This zone is hence used to load ions before merging them into one of the operation zones.
- A transfer zone. As the name suggests, this zone is used to transport ions between the loading and the operation zones.

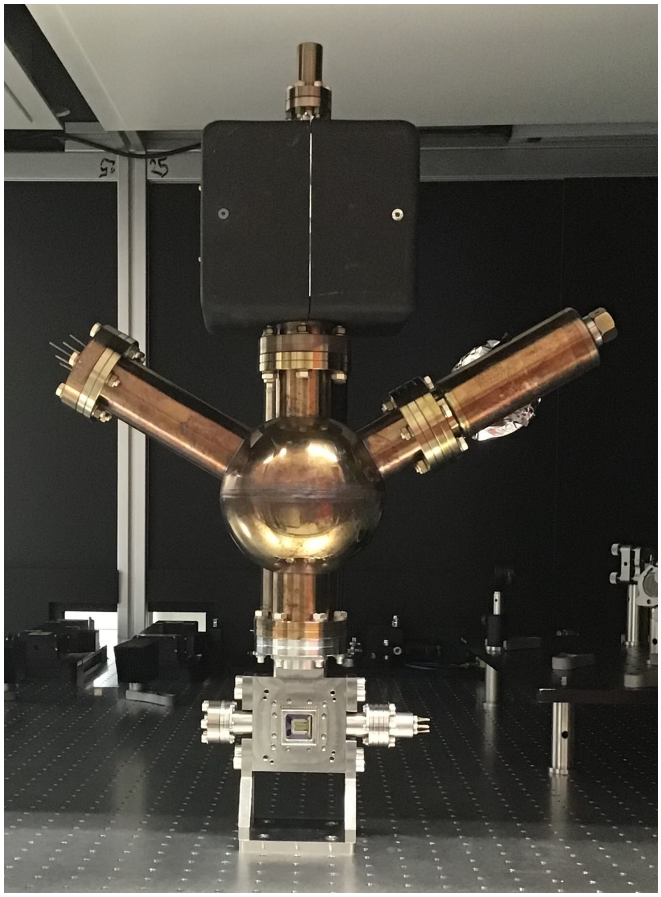


Figure 4.1: The trap vacuum chamber and pumping assembly mounted on the optics table.

- End electrodes. These electrodes provide additional degrees of freedom that can be used, for example, for electric field compensation. They are not used in the rest of this thesis.

During an experiment sequence, we store the ions in zone 5 of the trap (as labelled in Fig. 4.2). Hence, all laser beams, except for the photoionisation beams, are aligned to this zone (see Section 4.3.5). When we need to load an ion, we first load an ion in zone 10 before merging it with the remaining ions in the chain in zone 5. No Doppler cooling is performed in zone 10 of the trap before transporting the ion. Further details on the loading sequence can be found in Section 4.1.4.

In our apparatus, we modified the device such that the silicon substrate of the trap is also connected to a DC voltage source. Biasing of the trap substrate enables efficient rotation of the ion's radial motional modes. In the future, this

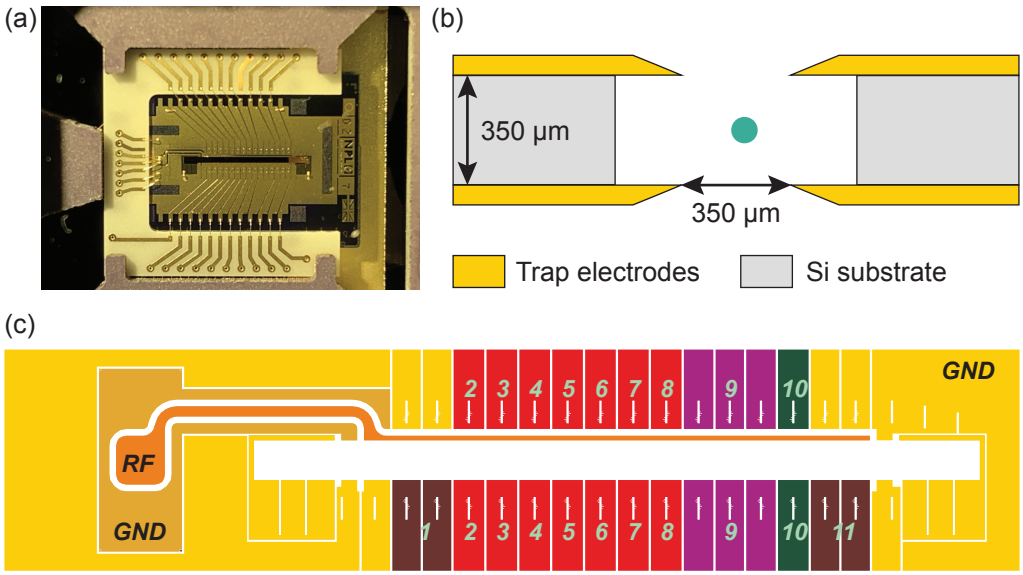


Figure 4.2: An overview of the trap layout. (a) A photo of the trap chip. (b) A cross-section of the trap layout in a plane perpendicular to the ion chain. The distance between the trap electrodes is 350 µm resulting in an ion-electrode distance of $\approx$ 247.5 µm. (c) Trap electrode map [same view as in (a)]. The ion chain is at the centre of the slot, parallel to the RF electrode. Zones 1 and 11 (brown) are end electrodes that could be used as additional degrees of freedom in generating electric field compensation voltagesets. Zone 10 (green) is the loading zone. Zones $2-8$ (red) are the trap operation zones. Zone 9 (purple) is the 'transfer' zone used for ion shuttling between the loading zones and the trap operation zones. Unless otherwise specified, during experiments the ions are confined in Zone 5. Zone 10 is only used during loading, and zones 3, 4, and 6 are used as overflow zones for ion merging, as described in Section 4.1.4.

could be useful for ensuring that the Raman lasers used for entangling gates only couple to one set of radial motional modes. This reduces the number of modes that need to be sideband cooled before a gate operation and makes decoupling from all motional modes at the end of a two-qubit gate easier.

4.1.3 Trap signal delivery

The trap is indium-sealed to the vacuum chamber. This allows for the pads of the chip carrier, which the trap is wire-bonded to, to remain outside of vacuum. This significantly simplifies both the design and the exchange of the PCB used for signal delivery to the trap, as shown in Fig. 4.3.

The operation of the ion trap requires the delivery of individual DC signals to each of the DC electrodes and an RF signal to the RF electrode in Fig. 4.2. The DC

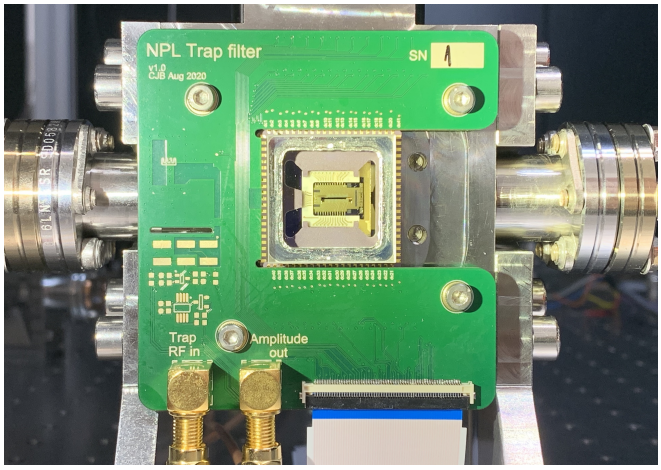


Figure 4.3: A photo of the PCB used to deliver signals to the trap chip.

voltages are generated using two Sinara Fastino boards, each containing 32 16-bit DAC channels. The signals generated by the Fastino boards are first connected to a small breakout PCB. The output is a single flat flex cable that is then connected to a PCB (the 'trap PCB') mounted on the vacuum system and connected to the trap chip carrier. The input DC signal from each Fastino channel passes through an RC filter[2] on the trap PCB before being delivered to the trap to attenuate high-frequency noise and thus minimise the trap heating rate, as shown in Fig. 4.4.

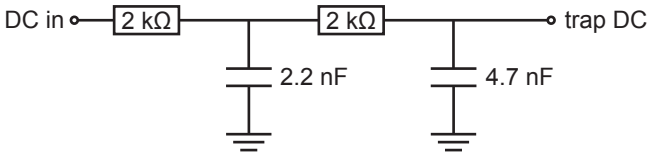


Figure 4.4: RC filter used to filter the input DC voltage to each electrode before delivery to the trap. The component values were chosen to minimise high-frequency noise, thus minimising the trap heating rate while facilitating shuttling operations.

The trap RF voltage source is a single DDS[3] channel of a Sinara Urukul board. The output of the DDS is amplified[4] to $\sim 30\,\mathrm{dBm}$ and connected to the trap PCB. On the trap PCB, it passes through an LC resonator circuit, which converts the input low-voltage high-current signal into a high-voltage low-current signal, which is then connected to the RF pin on the trap chip carrier,

[2]This filter was designed by Chris Ballance.
[3]AD9910
[4]Minicircuits ZHL-03-5WF+

as shown in Fig. 4.5. The resonator was designed to have a resonance frequency of 20 MHz. This was chosen as a compromise between having a high radial motional mode frequency while keeping the RF frequency much higher than the mode frequency [59]. The measured resonance frequency of the resonator, when connected to the trap, is 17.685 MHz. Part of the output signal from the LC resonator circuit is sent to an oscilloscope to monitor the voltage applied to the trap. The ratio between the pickoff voltage and the trap voltage is 1 : 540. During experiment operation, the inferred RF voltage amplitude at the trap is 200 V. This is a conservative maximum voltage that can be applied to the chip without causing dielectric breakdown.

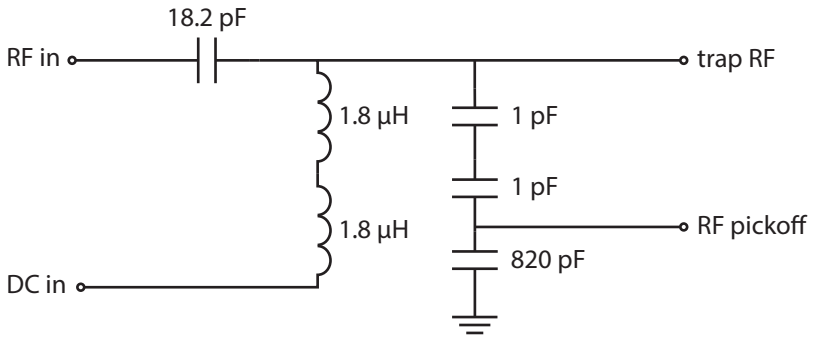


Figure 4.5: LC resonator circuit used to convert the input low-voltage high-input current RF signal into a high-voltage low-current signal before delivery to the trap. The resonance frequency of this circuit, when connected to the trap chip, is 17.685 MHz. A small amount of the signal delivered to the trap is picked off using the 820 pF capacitor and is used to monitor the applied voltage. The measured ratio between the RF amplitude at the trap and the pickoff signal is ≈ 540. The discrepancy between the capacitor values shown in the circuit above and the measured pickoff fraction is most likely due to the capacitance of the pickoff cable to ground.

4.1.4 Isotope-selective loading sequence

The trap electrode layout shown in Fig. 4.2 (b) allows ions to be loaded into a dedicated loading zone and then transported to the trap operation zone. This enables using a shield inside the trap chamber to constrain the atomic flux to the loading zone only. This minimises the effect of the particle flux generated during loading on the operation zones. The flux generated during loading can have two detrimental effects. First, if loading is performed directly in the operation

zone, the particle flux can cause collisions with the ions that are already trapped, leading to ion loss. Second, in the case of charged particle flux (as is often the case if ablation loading is used), it can cause charging of the trap itself, leading to changes in the stray fields at the trapping region.

In addition to shielding the operation zone from unwanted particle flux, using a dedicated zone for loading followed by shuttling to the operation zone enables the ions to be loaded one by one. That way, we are able to load the exact number of ions required for a given experiment. Furthermore, before merging each ion we can check whether it is of the correct isotope. We do this by measuring the ion fluorescence using a camera (see Section 4.2) and ensuring that it is above a certain threshold. By using a camera for imaging, we are also able to determine the number of ions in the chain at any time and terminate the sequence once the required number has been loaded. This capability allows us to build isotopically pure chains of the exact number of ions required. The detailed loading sequence is shown in Fig. 4.6. This sequence has been successfully tested for loading up to 10 ions in the trap operation zone. When attempting to load longer crystals, we found that we were predominantly losing ions during the merging step. This could be improved in the future by optimising the transport waveforms.

4.2 Imaging system

We make the barium ions fluoresce by continuously driving the $S_{1/2} \leftrightarrow P_{1/2}$ transition using a 493 nm laser as shown in Fig. 3.2. The short-lived $P_{1/2}$ level decays predominantly back to the $S_{1/2}$ level, resulting in the emission of 493 nm photons. To detect the ions, we need to collect and image the light they scatter in the process. The ion imaging system is shown in Fig. 4.7(a).

The 493 nm light is close to the centre of the visible range, making it compatible with off-the-shelf microscope objectives we can use to image the ions. In our experiment, we collect light from the ions using a commercial microscope objective[5] with $NA = 0.5$. This objective is glass-thickness compensated for 3.5 mm N-BK7

[5]Mitotoyo G Plan Apo 50

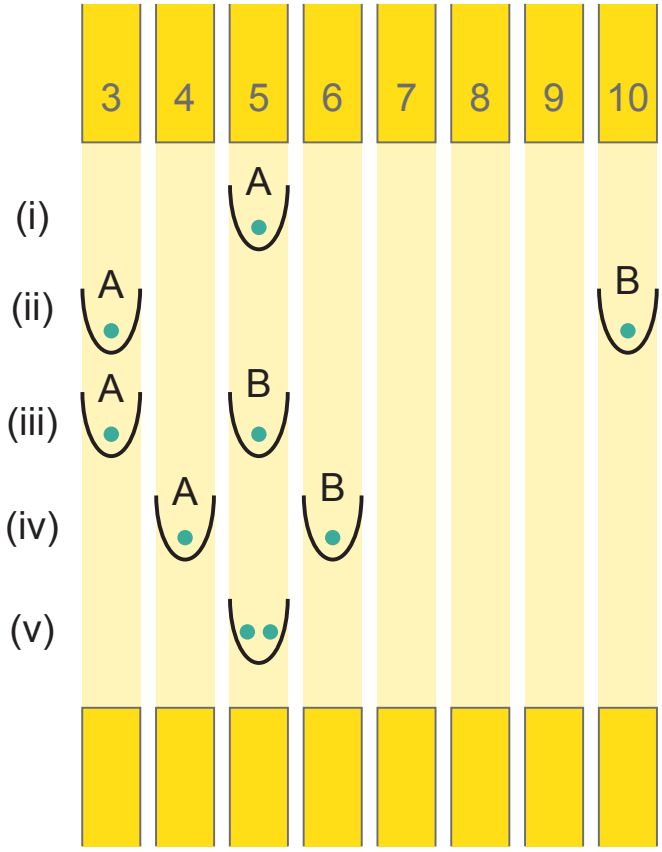


Figure 4.6: Isotope-selective loading sequence. (i) We start with a trapping potential at Zone 5 of the trap that contains a certain number of ions (one in the example above). (ii) When more ions need to be loaded, the ions that are currently trapped are transported to Zone 3 for storage, and a trapping potential is created in Zone 10 (the loading zone). (iii) After a chosen loading duration ($\sim$ 10 ms in our experiment), the trap in Zone 10 is transported to Zone 5, where it is cooled using the 493 nm and the 650 nm beams at their idle cooling parameters, as defined in Section 3.3.2. The ion is then imaged using a camera, and the resulting fluorescence counts are recorded. If the fluorescence counts are below a certain threshold (corresponding to either the absence of an ion or an ion of the wrong isotope), the ion is transported back to Zone 10 and then kicked out of the trap by applying DC voltages on the electrodes in Zones 9 and 10, resulting in an anticonfining potential in these zones. (iv) If the fluorescence counts are above this threshold, indicating that the ion is a $^{137}Ba^{+}$ ion, the ions stored in Zone 3 are transported to Zone 4, and the newly trapped ion is transported from Zone 5 to Zone 6. (v) Finally, a merging waveform is used to merge the ions in Zones 4 and 6 into Zone 5. This sequence is repeated as many times as necessary until the required number of ions is reached.

glass. Thus, we also use 3.5 mm N-BK7 vacuum chamber viewports (see Fig. 4.1). The objective is placed on a stack of linear[6] and tip/tilt[7] positioning stages. The linear positioning stages are used to ensure we image the correct region of the trap and to focus the objective on the ions. The tip/tilt stages are optimised to minimise aberrations. After the objective, we use a narrow line filter[8] to filter out any light except for 493 nm and then we re-image the ion on a low-noise scientific CMOS camera (sCMOS)[9]. The fact that the $^{137}Ba^{+}$ fluorescence wavelength is 493 nm helps us once again: like most scientific cameras, our sCMOS camera is optimised for visible wavelengths, and the estimated quantum efficiency at 493 nm is $\approx$ 90%.

An image of an ion taken with this camera is shown in Fig. 4.7(b). To measure the amount of fluorescence from the ion, we sum the counts from the pixels in an $n \times n$ region around the ion's centre. Using a larger value of n increases the magnitude of the signal but also increases the amount of noise we add to our measurement since the noise per pixel is broadly uniform over all pixels. Therefore, to maximise the signal-to-noise ratio, we must ensure that the image of the ion is spread over the smallest possible number of pixels [75]. In this setup, we use an $f = 35$ mm achromatic lens[10]. This is the shortest focal length lens compatible with the camera housing. The measured magnification of the imaging system is ≈ 6.5[11]. If a lower magnification is needed in the future, a multi-lens system could be used instead. A detailed discussion of the exact parameters used during ion readout is provided in Section 6.3.1.

[6]Newport M-561D-XYZ
[7]Newport M-561-TILT
[8]Semrock FF01-494/20-25
[9]Hamamatsu ORCA Fusion BT
[10]Thorlabs AC254-035-A
[11]This measurement was performed by measuring the separation of two $^{137}Ba^{+}$ ions on the camera as a function of their axial frequency and comparing it with the analytically calculated separation.

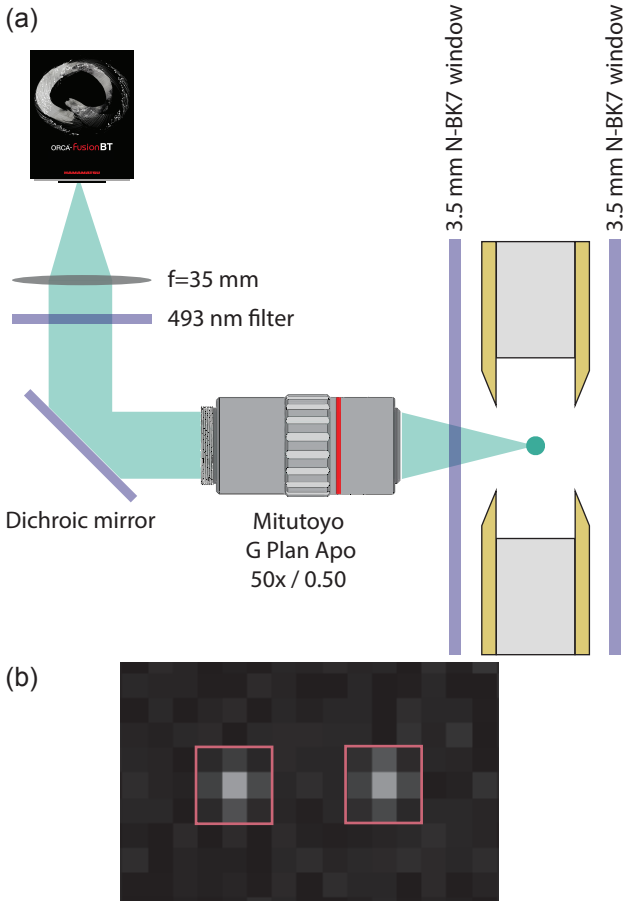


Figure 4.7: (a) Ion imaging system. Light from the ions is collected using an $NA = 0.5$ commercial microscope objective. The light then passes through a narrow line filter and is focused on an sCMOS camera using an $f = 35$ mm lens. The measured magnification of the imaging system is ≈ 6.5. (b) An image of two trapped $^{137}Ba^{+}$ ions. To measure the ion fluorescence, we sum the counts from a 3×3 region around the centre of the ion image.

4.3 Laser systems

Similar to Chapter 3, we logically split the lasers used in this experiment into three groups: the photoionisation lasers, the ion operations lasers (used for all fundamental operations such as cooling, state preparation, or readout), and the Raman lasers (used for implementing quantum logic operations on the qubits).

We further split the laser systems into two separate beam paths: the *laser* beam path and the *trap* beam path. The laser beam paths contain the laser source and all necessary frequency and power modulation components, typically located on a separate optical table from the trap. The light from the laser beam

paths is delivered to the trap using optical fibres, thus decoupling the laser beam path alignment from the trap alignment. The trap beam paths contain only passive optical components used for polarisation setting and generating the required spot size at the ion position.

4.3.1 Photoionisation lasers

The scheme used to ionise neutral barium atoms is shown in Fig. 3.3. The 553 nm laser is used to resonantly excite the $6s^2 \rightarrow 6s6p$ transition. Hence, its frequency must be set precisely within the 20 MHz linewidth of the transition. On the other hand, the frequency of the 405 nm laser must be high enough to excite the electron in the $6p$ state to the continuum but low enough so that the electron in the $6s$ state is not affected by this process. Hence, its frequency does not need to be set precisely.

There are no direct 553 nm laser diodes. To generate 553 nm light, we instead use a diode laser at 1106 nm[12] that is frequency doubled to 553 nm using a fibre-coupled waveguide doubler[13]. A subsequent AOM in double-pass configuration is used to fine-tune the frequency of this laser on resonance with the required isotope and to switch the light to the trap in real-time. During loading, this allows us to precisely define when and for how long we ionise the neutral barium atoms. Since the frequency of this laser needs to be set precisely, it is locked to a medium-finesse cavity (see Section 4.3.3).

The 405 nm laser source is a single-mode diode laser[14]. Since the tolerance in the wavelength of this laser is a few nanometres, it is not locked to a cavity.

The 553 nm laser and the 405 nm laser are combined using a dichroic mirror and then coupled into a fibre that goes to the trap table. A TTL-driven servo motor is used as a shutter to block the 405 nm light and the 553 nm light when no ions are being loaded.

[12]Toptica DL Pro
[13]HC Photonics Corp.
[14]Toptica iBeam smart

4.3.2 Ion operations lasers

493 nm

The 493 nm laser is used to drive transitions between the $S_{1/2}$ and $P_{1/2}$ levels, necessary for state preparation, measurement, and readout, as described in Chapter 3. As such, it must be able to drive transitions out of all states within the $S_{1/2}$ manifold of $^{137}Ba^{+}$. The 493 nm laser is a diode laser[15]. Similarly to the 553 nm laser, it is locked to a medium-finesse cavity (see Section 4.3.3).

The carrier laser frequency is set on resonance with the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ transition. To generate the frequencies necessary for driving the remaining transitions, we use a fibre EOM[16]. We use an RF synthesizer[17] to generate tones at 8.037 GHz and 6.56 GHz that are amplified and combined using an RF splitter and are then connected to the EOM. Due to sideband mixing within the EOM, the output light has sidebands at 8.037 GHz, 6.56 GHz, and 1.48 GHz. The strengths of the two RF tones are chosen to generate roughly equal amounts of power in all frequency components. Thus, even with purely π-polarised light (necessary for our state preparation scheme, see Chapter 3), we can drive all possible transitions between the $S_{1/2}$ and $P_{1/2}$ levels, as shown in Fig. 3.4.

During state preparation, we must turn off the tone at 6.56 GHz in real-time, as shown in Fig. 3.5. This is done using a TTL-triggered[18] RF switch[19] connected to the output of the corresponding RF synthesiser channel. We use an AOM in a double-pass configuration for adjusting the common frequency offset of all sidebands to tune the laser between being on resonance (for readout) and being red-detuned (for Doppler cooling).

A final AOM in a single-pass configuration switches the light on and off. This AOM is not strictly necessary and was only installed to provide further degrees

[15]Toptica DL PRO HP
[16]ADVR WPM-K0493-P48P48AL0
[17]Windfreak SynthHD PRO
[18]Sinara DIO-BNC
[19]Minicircuits ZASWA2-50DR-FT+

of freedom in case an additional 493 nm trap beam path was necessary. The light is then coupled into a fibre and delivered to the trap.

650 nm

The 650 nm laser used to repump population that decays to the $D_{3/2}$ level after excitation to the $P_{1/2}$ level. The 650 nm laser is a diode laser[20] and is locked to a medium-finesse cavity. Similar to the 493 nm laser, the 650 nm light is coupled into a fibre EOM[21]. However, contrary to the $S_{1/2}$ level, the hyperfine splitting in the $D_{3/2}$ level is much smaller (see Fig. 3.4). Therefore, rather than generating all possible sidebands, we only apply a single tone at 430 MHz[22] and use $\sim 1000 I_{\mathrm{sat}}$ to power-broaden the transition linewidth, as shown in Fig. 3.4. The strength of the sidebands is chosen to maximise the fluorescence of the ions. As with the 493 nm laser, an AOM in double-pass configuration is used for fine-tuning the offset of all sidebands, power stabilisation, and switching the beam on and off.

614 nm

The 614 nm light is used to deshelve any population in the $D_{5/2}$ level, as shown in Fig. 3.2. There are no diodes available at 614 nm, and hence we generate this wavelength by using an amplified diode laser[23] at 1228 nm that is frequency-doubled to 614 nm[24]. As before, the output of the 1228 nm laser is locked to a medium-finesse cavity. Due to the very small hyperfine splitting of the $D_{5/2}$ level, we do not use an EOM here to generate sidebands but rather power-broaden the transition. As with the 650 nm laser, an AOM in double-pass configuration is used for fine-tuning the laser frequency, power stabilisation, and switching the beam on and off. The 614 nm and the 650 nm laser beam paths are combined using a 50 : 50 beam splitter and coupled into the same fibre for delivery to the trap.

[20] Toptica DL PRO
[21] ADVR WPM-K0650-P63P63AL0
[22] Sinara Mirny
[23] Toptica TA Pro
[24] HC Photonics Corp.

1762 nm

The 1762 nm laser is used for population transfer between states in the $S_{1/2}$ and the $D_{5/2}$ level. These are quadrupole transitions with much lower matrix elements compared to dipole transitions. Hence, a high intensity is required to drive these transitions reasonably fast ($\sim$ 10 µs).

The 1762 nm laser is a diode laser[25]. We amplify the laser with a fibre amplifier[26] whose maximum output power is 2 W. A small part of the output light is sent to a high-finesse locking cavity (see Section 4.3.3). We use a fibre EOM to span the transition frequency range between the $S_{1/2}$ and $D_{5/2}$ manifolds. A subsequent AOM in single-pass configuration is then used to turn the light on and off and to stabilise the total 1762 nm power at the trap.

4.3.3 Locking cavities

In our experiment we use two types of locking cavities[27] for laser frequency stabilisation: a four-bore medium finesse cavity used for the lasers that drive broad linewidth transitions (1106 nm, 493 nm, 650 nm, 1228 nm); and a single-bore high-finesse cavity that is used for the 1762 nm laser that drives the narrow linewidth $S_{1/2} \leftrightarrow D_{5/2}$ transitions.

Four-bore cavity

The four-bore cavity comprises four mirror pairs mounted on a single cavity spacer. The estimated finesse of these cavities by the manufacturer is between 100 and 1000. Each mirror pair consists of a flat mirror and a spherical mirror with a radius of curvature of 50 cm, spaced by 100 mm, resulting in a cavity free spectral range (FSR) of $\approx$ 1.5 GHz.

The lengths of the cavity bores are fixed. Hence, we must ensure that the output light of the laser, when tuned on resonance with the desired atomic

[25]Toptica DL PRO
[26]Precilasers
[27]Stable Laser Systems

transition, matches one of the cavity resonances. To achieve this, we use an AOM in a double-pass configuration before delivering the light to the cavity.

We lock each laser to the cavity using a standard Pound-Drever-Hall lock[28].

High-finesse cavity

Contrary to the lasers locked to the four-bore cavity, the 1762 nm laser drives very narrow-linewidth transitions (linewidth of $\approx 30\,\text{mHz}$). We, therefore, need a significantly narrower linewidth of this laser to ensure that we can drive the quadrupole transitions efficiently. Thus, we lock this laser to a cavity with a finesse of ~ 300000. In addition to the higher cavity finesse, we use a fast laser locking controller[29] together with the standard PDH lock. Here, we use an EOM[30] in the optical beam path between the laser and the cavity to ensure that the target laser carrier frequency matches a cavity resonance.

4.3.4 Raman lasers

The two types of qubits we aim to work with have transition frequencies of ~ 8 GHz and ~ 50 MHz. As described in Section 3.4.3, the barium level structure enables driving the qubit transition using a two-photon Raman process with 532 nm light for both the ground and metastable level qubits. This greatly simplifies the qubit control laser systems, since we only need to generate two different Raman beatnotes, rather than work with completely separate laser sources and beam paths for each qubit type. Due to the large detuning of the 532 nm laser from the 493 nm and the 614 nm transitions, we need a substantial amount of intensity to drive the qubit transitions reasonably quickly, in ~ 1 µs. Using [43], we estimate that we require ~ 1 mW in a beam with 1 µm waist radius to drive the ground-level $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ transition in ~ 1 µs provided the beam polarisations are optimised to maximise its Rabi

[28]Topica PDH/DLC Pro
[29]Toptica FALC Pro
[30]We decided to use an EOM in this case, rather than an AOM as in the four-bore cavity, due to the generally lower tuning range and efficiency of AOMs at 1762 nm.

frequency. Therefore, we use a high-power laser source to give us flexibility when working with a large number of ions.

The ~ 50 MHz beatnote necessary for driving the metastable level qubit can be easily generated from a single laser source using a set of AOMs. However, driving the ground level qubit requires an ~ 8 GHz beatnote. At these frequencies, optical frequency modulators are generally inefficient, so we designed a new scheme to generate the required frequency offset before a final light amplification stage.

The laser system we designed to generate the 532 nm light is shown in Fig. 4.8. We use a 1064 nm laser source[31] that is amplified and split into two arms, one with an output power of 50 mW (the 'carrier') and one with an output power of 450 mW (the 'offset'). The carrier is connected to a second amplification stage and is then frequency-doubled to produce ≈ 10 W of 532 nm light.

The offset is connected to a fibre EOM[32] driven at 4 GHz using an RF synthesizer[33]. The output of the EOM is connected to a home-built cavity. This cavity consists of two concave mirrors[34] with a radius of curvature of 500 mm, spaced by approximately 6 cm. The 1064 nm laser is locked to the $+1^{\text{st}}$ EOM sideband[35]. The cavity output is fibre-coupled, connected to a second amplification stage, and frequency doubled to produce 10 W of power at 532 nm, offset by 8 GHz relative to the carrier.

We then use a series of AOMs[36] shown in Fig. 4.9 to shift the frequencies of the individual beams to match them to the qubit transition frequencies and modulate the phase and amplitude of the light. We will call the output of the AOM network after the offset arm the 'ground' beam. The light from the carrier is split up into two beam paths, a 'global' beam path and a 'metastable' beam path. The AOM frequencies in the 'global' beam path are such that:

[31] Toptica DL Pro
[32] iXblue
[33] HP 4536B
[34] Layertec 108864
[35] Note that the absolute laser frequency here is irrelevant. For driving Raman transitions we only care about the frequency splitting between the two beams. Therefore, when locking the laser to the cavity, we only need to ensure that we lock to the correct sideband to generate the desired frequency offset relative to the carrier laser.
[36] AA Optoelectronic

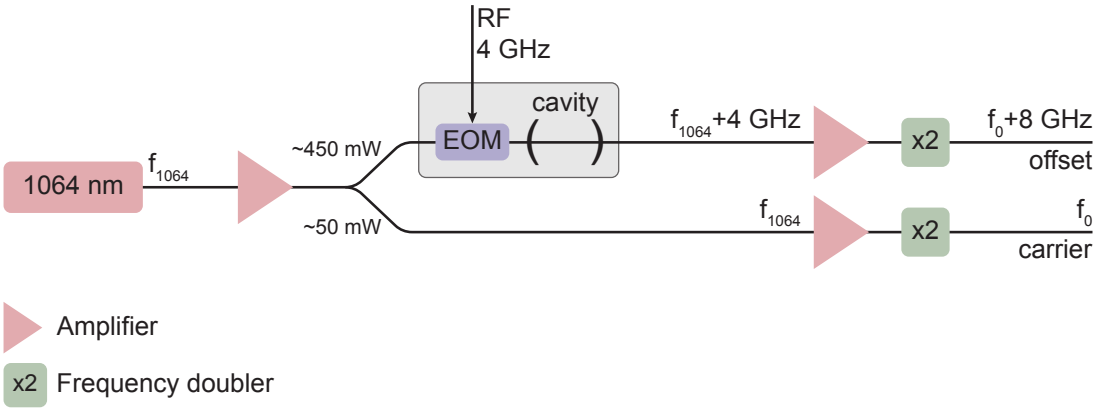


Figure 4.8: Raman laser system. The source 1064 nm laser is amplified and split up into two arms. One arm (the carrier) is amplified and frequency-doubled to produce $\sim$ 10 W of 532 nm light. The second arm (the offset) is frequency-modulated using an EOM driven at 4 GHz. It then passes through a cavity and the laser is frequency-locked to the $+1^{\text{st}}$ sideband of the cavity. The output of the cavity is then amplified again and frequency-doubled to generate $\sim$ 10 W of 532 nm light detuned by 8 GHz relative to the carrier.

- the frequency difference between the global beam and the metastable beam matches the metastable level qubit frequency;

- the frequency difference between the global beam and the ground beam matches the ground level qubit frequency.

The global beam path is further split into two separate paths with identical net frequency shifts. At the ion position, the output of one of them, the 'co-propagating' global beam, is incident parallel to the ground and the metastable beams. This configuration results in a net zero $\vec{k}$-vector difference between the Raman beams used to drive the qubit transition, making the transitions insensitive to the ion's motion. This is the configuration used for single-qubit gates. The output of the second path, the 'counter-propagating' global beam, is anti-parallel to the ground and the metastable beams. The resulting $\vec{k}$-vector difference in this configuration is $\vec{k} = 2\vec{k}_{532}$, where $\vec{k}_{532}$ is the $\vec{k}$-vector of 532 nm light. Thus, this configuration enables very strong coupling to the ion's motion. This is the configuration used for both sideband cooling (which requires driving the motional mode sidebands) and two-qubit entangling gates (which rely on

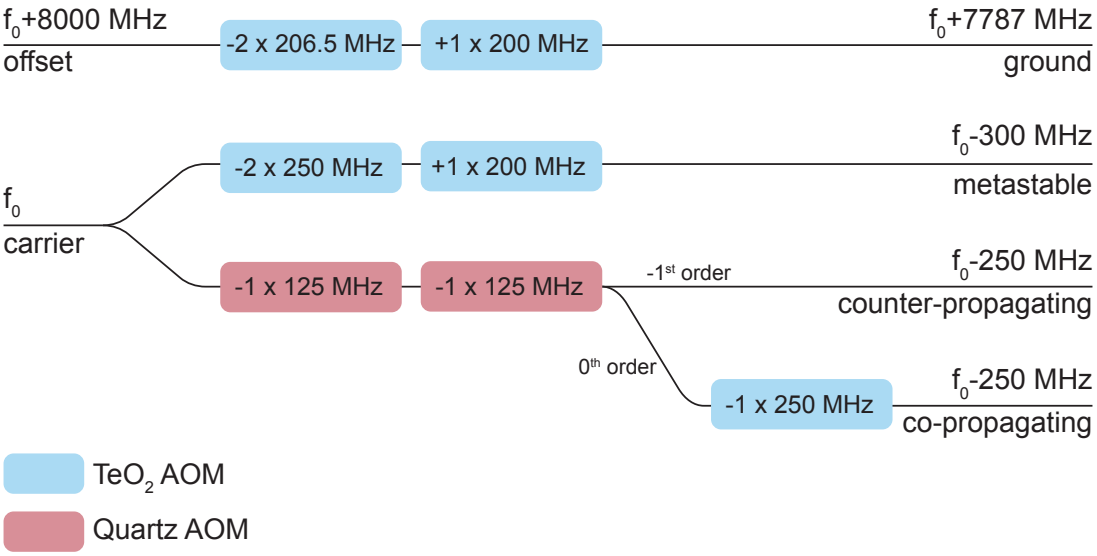


Figure 4.9: AOM network used to tune the frequencies of the Raman lasers on resonance with the ground and the metastable level qubits. The carrier and the offset lasers are denoted in Fig. 4.8. The frequencies shown in the diagram above are the centre AOM frequencies generating roughly the right frequency splitting. The double-pass AOMs in the ground and the metastable level beam paths are used for fine-tuning the beam frequencies on resonance with the desired transition. The ground beam, together with either the co- or counter-propagating beam, is used to drive ground-level transitions. The metastable beam, together with either the co- or counter-propagating beam, is used to drive metastable-level transitions. The ground, metastable, and co-propagating beams are tightly focused to ~ 1 µm spot size onto the ions using our individual addressing system (see Chapter 5) and therefore require little power. Hence, we use TeO_2 AOMs for these beam paths. The counter-propagating beam is global for the qubit register and, therefore, has a much larger spot size (~ 10 µm). Hence, it requires much higher power, as we will see in Chapter 8. Therefore, we use a set of quartz AOMs to modulate its frequency.

applying a spin-dependent force on the ion to excite the ion's motion). The directions of these beams relative to the trap are shown in Fig. 4.11.

Implementing a generic algorithm in a quantum computer requires the ability to perform an arbitrary gate on a subset of qubits, without affecting the remaining qubits in the register. To achieve this, we developed a novel integrated photonic chip capable of producing many closely-spaced tightly-focused beams with very low cross-talk. Due to the complexity of the individual addressing system, it is described separately in Chapter 5.

With this design of the Raman beam paths, we are able to tune which ion, what type of qubit, and what type of gate we drive in real-time by simply

modulating the amplitudes and frequencies of the tones applied to the AOMs in Fig. 4.9, which can be done from software.

4.3.5 Trap beam paths

At the end of the laser beam paths, all beams are coupled into optical fibres and delivered to the trap. All trap beam paths follow a similar general pattern shown in Fig. 4.10. The beam first goes through a PBS, which cleans up the beam polarisation and turns any polarisation fluctuations into power fluctuations. The following beam sampler sends a small fraction of the light to a photodetector. The signal on the photodetector is used to provide feedback on the AOM amplitude to stabilise the power delivered to the trap[37]. This is followed by a set of waveplates used to set the desired polarisation of the light. A final set of lenses is then used to set the size of each beam at the ion position.

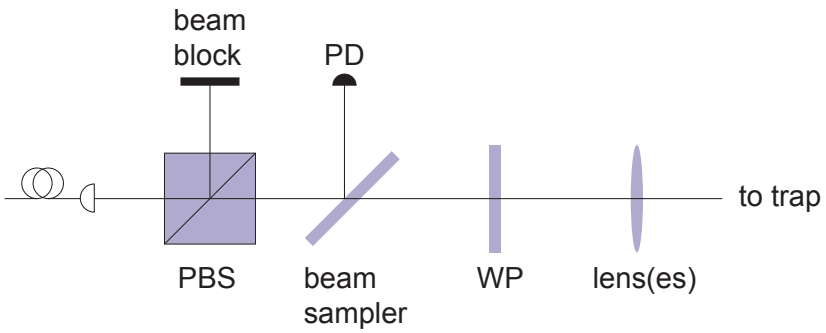


Figure 4.10: Generic trap beam path. The polarisation of the beam is cleaned up using a PBS, which also turns any polarisation fluctuations into power fluctuations. A small amount of the beam power is sent to a photodetector (PD) used in a feedback loop for power stabilisation. A subsequent set of waveplates and lenses is used to set the beam polarisation and the spot size at the ion position.

The directions of the individual beams relative to the trap chamber are shown in Fig. 4.11. The scheme we use for qubit state preparation described in Section 3.3.3 requires a 493 nm beam with pure π polarisation. Therefore, the 493 nm beam must be perpendicular to the applied magnetic field. The trap design does not allow access for a beam parallel to the ion chain, and we must avoid a 493 nm beam parallel to the path we use to collect ion fluorescence to minimise

[37] Artiq SuServo

background light during readout. These considerations constrain the directions of the magnetic field and the 493 nm beam to the ones shown in Fig. 4.11.

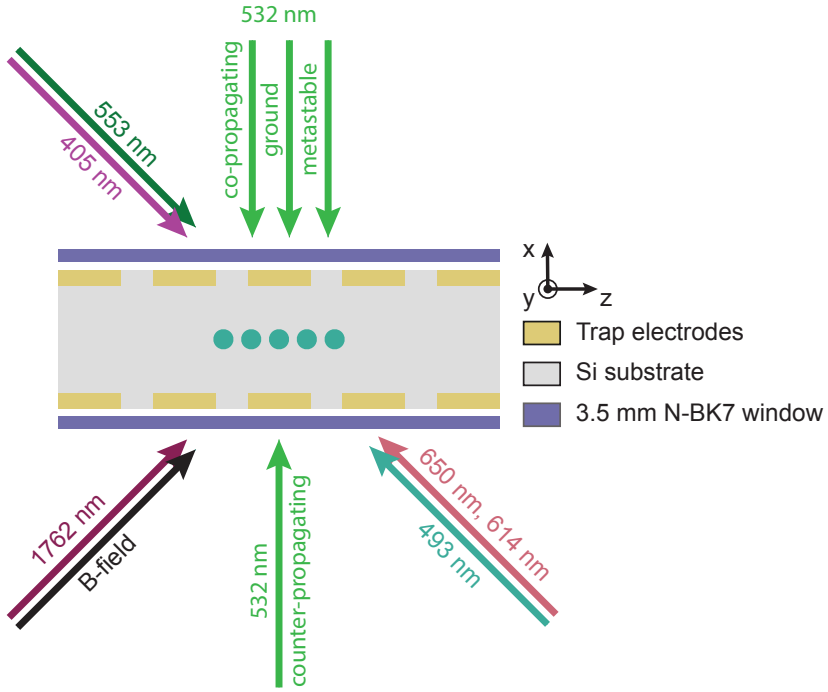


Figure 4.11: Directions of incidence of the individual laser beams onto the ions. The ground, metastable, and co-propagating Raman beams overlap in practice but are drawn separately here for clarity.

The 650 nm and the 614 nm beams are combined in a single fibre at the end of the laser beam paths (see Section 4.3.2) and therefore share a single trap beam path. Their direction and polarisation do not matter much, as long as we can pump population out of all states within $D_{3/2}$ and $D_{5/2}$. We combine them with the 493 nm beam path and we set their polarisation so that we have roughly equal powers in each of the three polarisation components.

The 1762 nm beam direction and polarisation constrain the quadrupole transitions between $S_{1/2}$ and $D_{5/2}$ we can drive [76]. When the experiment was set up, the direction of the 1762 nm was chosen for convenience. It is parallel to the quantisation field and has horizontal polarisation. As a result, we can only drive transitions with $\Delta m = \pm 1$.

The 532 nm Raman beams are incident perpendicular to the ion chain. First, this ensures that we only couple to the ions' radial modes during two-qubit

entangling gates and not to the axial modes. That way, we are not sensitive to the temperature of the axial modes, which are lower-frequency, heat up more, and are harder to cool. This allows us to use the same high-precision objective for both imaging the ions (see Section 4.2) and focusing down the small beams from our addressing system with minimal aberrations (see Chapter 5). As we discussed in Section 4.3.4, we support two possible configurations of the Raman beams: a configuration in which the two Raman beams used to drive the qubit transition are co-propagating and therefore do not couple to the ion's motion (used for single-qubit gates); and a configuration in which the two beams are counter-propagating and therefore couple to the ion's motion (used for ground-state cooling and two-qubit gates). The polarisations used for the Raman beams vary throughout the thesis and are discussed within each chapter, where relevant.

Finally, the direction and polarisation of the photoionisation beams are not critical. Contrary to the remaining beams, which are aligned to the trap operation zone [Zone 5 in Fig. 4.2(b)], the PI beams need to be aligned to the trap loading zone [Zone 10 in Fig. 4.2(b)]. Hence, they have their dedicated beam path, as shown in Fig. 4.11.

With the exception of the Raman beams, all beam paths are designed to generate a spot with a $1/e^2$ radius of ≈ 50 µm at the ion position. The spot sizes used for the Raman beams vary throughout the thesis and are discussed within each chapter, where relevant.

Baby, it's you
You're the one I love
You're the one I need
You're the only thing I see

— Beyonce, *Love on Top*

5 Individual addressing system

Contents

5.1 Background 56
5.2 Individual addressing system design using a novel integrated photonic chip 58
5.2.1 Fibre network and photonic chip design 58
5.2.2 Coupling into the photonic chip 60
5.2.3 Integration with the ion trap system 62
5.3 Cross-talk characterisation using a single trapped ion 68
5.3.1 Robust phase estimation protocol 69
5.3.2 Knill Dynamical Decoupling Sequence 71
5.3.3 Measurement implementation 73
5.4 Summary 79

In this chapter, we describe the design of a novel scalable optical system used to focus individual 532 nm beams on the ions and, therefore, achieve individual qubit addressing with very low cross-talk. As described in Chapter 3, $^{137}Ba^{+}$ allows us to use a 2-photon Raman process to drive the transitions of both the ground and the metastable level qubits and hence to implement both single-qubit and two-qubit entangling gates.

For a many-qubit ion trap device to be useful as a quantum computer, it must be algorithm-agnostic. In other words, there must be no experimental limit to the algorithms that can be implemented, such as limited qubit connectivity or qubit

selectivity. One of the crucial capabilities required to achieve this is individual qubit addressing: the ability to execute an arbitrary unitary on a subset of qubits while leaving the remaining qubits in the register intact.

5.1 Background

Individual addressing in ion trap quantum computers requires locally modifying the interaction between the ions and the radiation used to perform operations with very low cross-talk, where the natural inter-ion spacing is a few micrometres [59, 77]. This can be done by altering the properties of the magnetic and/or electric fields experienced by each of the ions to change their response to a globally applied field [78–83], by focusing down the radiation used to perform quantum gates on each ion [61, 84–89], as shown in Fig. 5.1, or by using a combination of the two approaches. The former approach works with both laser and microwave radiation but requires intricate local control of electric and magnetic fields. Hence, it is only compatible with microfabricated traps with numerous electrodes. The latter approach is only viable when optical wavelengths are used for operations, either when the qubit transition is driven with a two-photon Raman process or the qubit transition itself corresponds to an optical wavelength. It requires focusing down laser beams to micrometre scales, therefore working close to the diffraction limit. However, this method is applicable to all ion trap architectures.

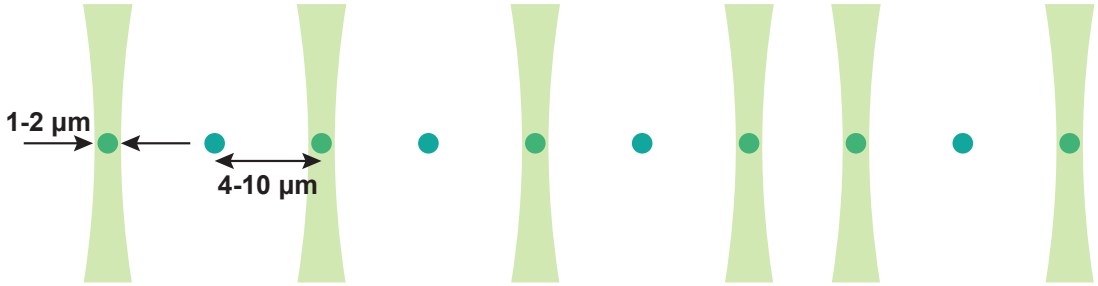


Figure 5.1: Requirements for individual optical addressing in long chains of trapped ions. The spacing between the individual laser beams needs to match the ion separaion, typically $4 - 10$ µm. The spot size needs to be small enough to minimise the intensity at the neighbouring ions, but larger than the Fourier limit for the laser wavelength in use to avoid cross-talk due to aberrations in the optical system.

Several methods have been developed for achieving individual optical addressing in ion trap systems, including micro-mirror beam steering [84–86], acousto-optic deflectors [87], multi-channel acousto-optic modulators [61] (AOMs), and micro-lens arrays [88]. These methods vary in cross-talk performance and scalability. Using micro-mirror arrays to steer the lasers onto the target ions offers individual addressing with very low cross-talk. However, the time required to reconfigure the beam positions or to tune the amplitude of the beams is comparable to the timescale of the gate operations, therefore accounting for a significant amount of the sequence run time. Acousto-optic deflectors offer similarly low cross-talk but also enable fast intensity control of the individual beams. However, they lack individual beam frequency control, therefore limiting the set of unitaries that can be implemented in parallel. Furthermore, only a few ions can be addressed at any one time due to power losses in higher diffraction orders. Multi-channel AOMs solve the problem of fast parallel control. However, they exhibit an order of magnitude larger cross-talk compared to the previous two approaches due to electronic cross-talk between channels in existing devices. Additionally, they are only manufactured with a fixed and equal inter-channel spacing, and readily available with only up to 32 channels, limiting both scalability and ability to address non-evenly spaced chains. The micro-lens array approach solves the majority of the problems listed above by enabling parallel operations with very low cross-talk. However, the system presented in Ref. [88] has a very high insertion loss limiting the obtainable light intensity at the ion, and a fixed, uniformly spaced output pattern that cannot be easily adapted to non-uniform ion crystals.

In this work, we demonstrate an approach to individual optical addressing in trapped-ion chains with minimal cross-talk using a network of fibre-coupled modulators connected to a high-performance photonic chip. We employ spherical phase-induced multiscan waveguides [90] (SPIM-WGs) that provide precise control over the optical mode and enable much higher refractive index (RI) contrast modifications in optical glass compared to conventional ultrafast laser-written

waveguides. A similar approach using laser-written waveguide devices has been reported [89], however matching the output of the waveguides presented in Ref. [89] to an ion chain while maintaining a good spot quality and low cross-talk is yet to be demonstrated. The use of a fibre network for light delivery to the photonic chip further simplifies the process of exchanging devices by minimising the necessary realignment, thus allowing for rapid iteration of the system design.

5.2 Individual addressing system design using a novel integrated photonic chip

5.2.1 Fibre network and photonic chip design

Performing laser-driven targeted operations in long chains of trapped ions imposes several competing requirements on the individual addressing setup. First, to minimise errors on the idle qubits, it is crucial to minimise the cross-coupling between neighbouring ion sites. Second, the spacing between neighbouring channels must match the ion-ion spacing, typically on the order of a few micrometres [59, 77]. This requires generating a series of closely spaced beams, each focused to a ~ 1 µm waist radius as shown in Fig. 5.1. As a result, this approach necessitates focusing the beams near the diffraction limit. Working with tightly focused beams can also increase errors in the quantum operations due to intensity, phase, and/or polarisation modulation of the light at the ion position. This modulation can be caused by mechanical drifts in the optical system or by the secular motion of the ions [91, 92]. Hence, it is desirable to make the beam waist radius as large as possible. The requirements on the beam spacing and beam waist radius constrain the waist-to-spacing ratio of the device output. A larger ratio makes integration into the trapped-ion system easier and more robust to errors but also increases the cross-talk within the device. Furthermore, to ensure that the required input laser power scales favourably with the qubit register size, it is important to maintain a high throughput efficiency of the system. Finally, the ability to control the intensity, phase, and frequency of individual channels in

parallel enables the simultaneous application of different unitary operations on different target qubits, therefore reducing the algorithm run time.

To address these challenges, we use a fibre network coupled to a novel integrated photonic chip, as shown in Fig. 5.2. Light from the free-space 532 nm laser source is split up into the required number of channels using a series of fibre splitters. The output fibres pass through polarisation controllers[1] to enable compensation of differential polarisation changes introduced to each of the channels. Each channel is then connected to a fibre AOM that provides individual amplitude, phase, and frequency control, thus enabling parallel operations. The outputs of the fibre AOMs are connected to a fibre V-groove array (VGA) whose output is an array of fibres with mode-field diameter (MFD) of 3.6 µm spaced by 127 µm. To map between the output of the VGA and the ions we use a novel photonic chip whose input matches the channel pattern of the VGA and whose output matches the distribution of the ions in the chain up to an overall scaling factor, as shown in Fig. 5.2.

As outlined above, there must be minimal light leakage from an active channel into the neighbouring ones to minimise any errors introduced onto the states of the idle qubits in the system. We employ spherical phase-induced multiscan waveguides (SPIM-WGs) designed and manufactured in collaboration with the Engineering Science department at the University of Oxford. Details on the fabrication of these devices can be found in [90]. This technique enables much higher refractive index contrast modifications in optical glass compared to conventional ultrafast laser-written waveguides and provides precise control over the shape of the optical mode. The photonic chip adopts individual adiabatic mode converters as light guiding channels (see Fig. 5.2) that exhibit excellent optical mode confinement and remarkably low cross-talk even at micrometre-scale channel separation, as we will see in Section 5.3.3. The improved optical performance ensures that the errors due to nearest-neighbour cross-talk do not limit the performance of the trapped ion device. Furthermore, the manufacturing

[1]Thorlabs CBC900

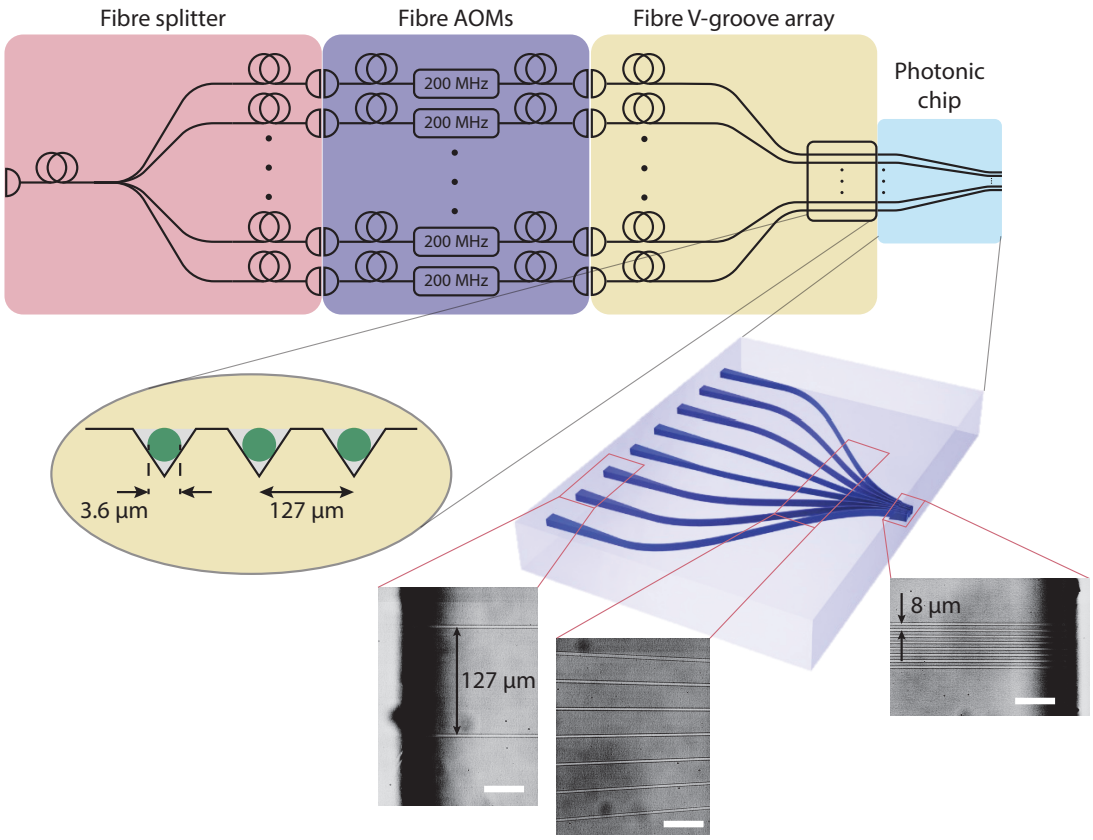


Figure 5.2: An outline of the fibre network used for individual addressing. The free-space 532 nm laser source is coupled to a series of fibre splitters. Each of the resulting channels is then connected to a fibre AOM that enables individual channel switching as well as individual phase and amplitude control. The outputs of the fibre AOMs are connected to a fibre VGA whose output is an array of fibre cores with an MFD of 3.6 µm and a pitch of 127 µm. To map between the output of the VGA and the ion pattern, we use an integrated photonic chip whose input matches the VGA output pattern and whose output matches the ion pattern (up to an overall scaling factor).

process of SPIM-WGs offers high flexibility, facilitating easy reconfiguration of the channel positions and mode shapes, making the device suitable for most trapped-ion setups.

5.2.2 Coupling into the photonic chip

To minimise the amount of light required per qubit in the system, we need to maximise the coupling efficiency between the VGA and the photonic chip. This coupling efficiency is extremely sensitive to relative position changes between the two components. In addition, changes in the coupling efficiency affect not only

the amount of light coupled into each channel but also the polarisation and the phase that the light accumulates as it travels through the channel. Therefore, it is crucial to minimise the variations in the coupling efficiency during operations. In this section, we describe the procedure used for coupling light from the VGA into the photonic chip and for glueing the photonic chip to the VGA to keep the coupling efficiency constant.

The system we use to optimise the coupling between the VGA and the photonic chip is shown in Fig. 5.3(a). Both components are placed on custom-made individual vacuum chucks. The vacuum chuck holding the waveguide is fixed, while the vacuum chuck holding the VGA is mounted onto a 6-axis positioning stage[2]. To evaluate the coupling efficiency, the output of the photonic chip is imaged using a commercial microscope objective[3] with an effective focal length (EFL) of 20 mm and is then refocused using an achromatic $f = 50$ mm lens onto a CMOS camera[4]. We then use all degrees of freedom of the positioning stages to maximise the light coupled into each waveguide channel. The excellent mode matching between the VGA and the photonic chip enables high single-channel coupling efficiencies (observed $> 80\%$). However, tolerances in the VGA and photonic chip fabrication reduce the average coupling efficiency that can be achieved for all channels simultaneously. In the 8-channel device presented here, we achieve an average coupling efficiency of 45(3)%.

Once the coupling between the VGA and the photonic chip is optimised, they are glued together using a UV-curing adhesive[5] as shown in Fig. 5.3(b). We cure the adhesive very slowly to avoid changes in the coupling efficiency during the curing process. We start with a 10-minute cure at 10% of the full power of a UV lamp[6]. We then gradually increase the UV power in steps of 10% and cure for ten minutes at each power level. Finally, we leave the UV lamp at full power for

[2]Thorlabs MAX603D/M
[3]Mitutoyo MY10X-803
[4]FLIR Blackfly S BFS-U3-200S6M-C
[5]Norland NOA61
[6]Thorlabs CS2010

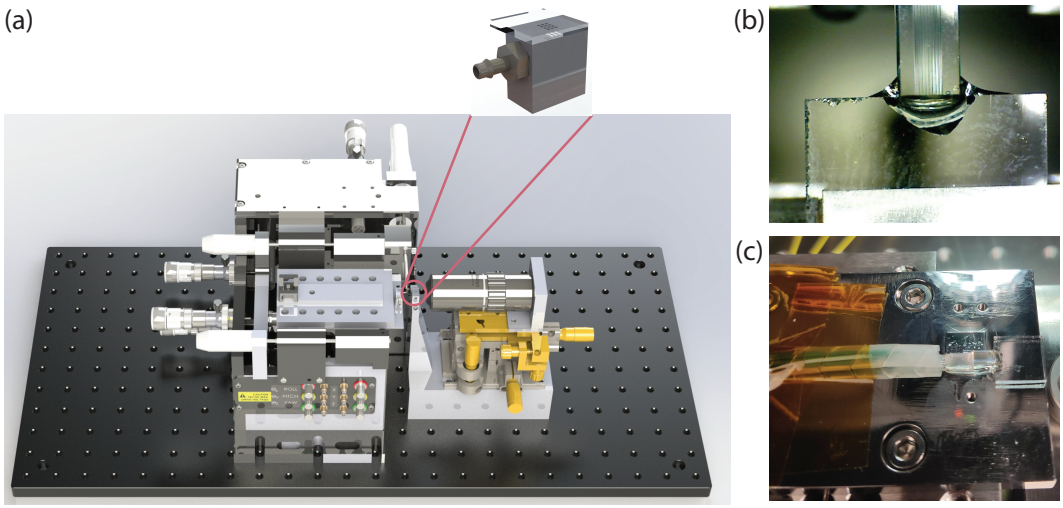


Figure 5.3: (a) Setup used to couple light from the VGA into the photonic chip. The VGA is placed on a 6-axis positioning stage (left) while the photonic chip is held stationary (platform on the right in front of the objective). The objective on the right is used to image the photonic chip output onto a camera (not shown) to evaluate the coupling quality. (b) A photo of the VGA glued to the photonic chip after the coupling is optimised. (c) A photo of the VGA and photonic chip assembly glued onto a stainless steel plate before integrating with the optical system in the ion trap setup.

60 minutes to ensure the glue is fully cured before transporting the assembly. We observe no measurable change in coupling efficiency during the curing process.

We analyse the device output after glueing using a CMOS camera, as shown in Fig. 5.4. We measure a beam waist radius of 1.25(1) µm along the axial direction (x axis in Fig. 4.11) and 1.76(7) µm along the radial direction (y axis in Fig. 4.11), and a channel separation of 7.84(3) µm, which are consistent with the design specifications. This measurement shows no observable nearest neighbour cross-talk compared to the noise floor of 2×10^{-3}.

Finally, the VGA and photonic chip assembly is glued to a stainless steel plate using the same UV curing adhesive and is transported to the ion trap system. This stainless steel plate is then bolted onto a positioning stage that is part of the setup, as shown in Fig. 5.3(c).

5.2.3 Integration with the ion trap system

After the VGA and the photonic chip are glued together, they must be integrated into an optical system that images the chip output onto the ion chain.

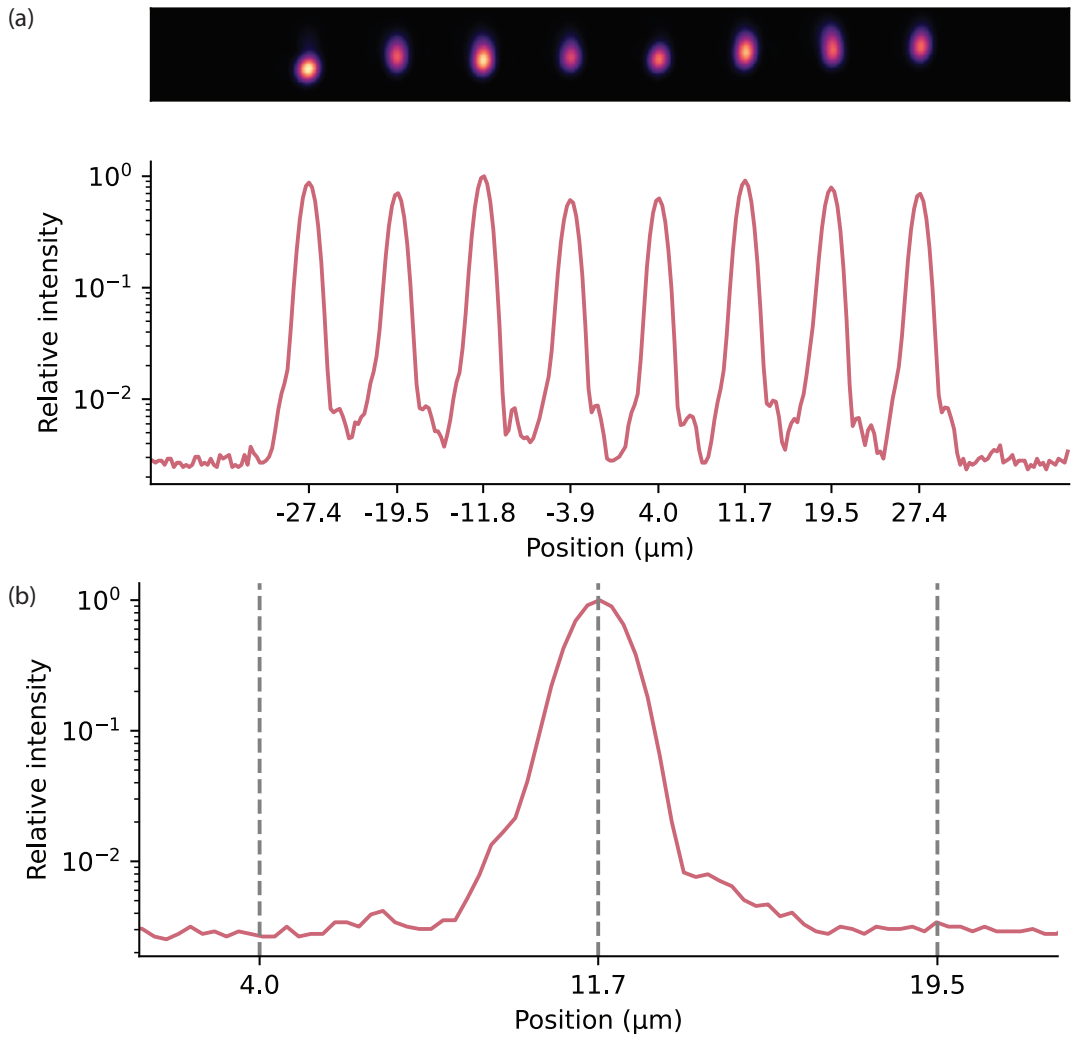


Figure 5.4: (a) A photo of the photonic chip output after glueing to the VGA taken using a CMOS camera, and a profile taken by summing the camera counts along the vertical axis. The slight difference in the vertical position of the channels comes from non-uniform coupling into each of the channels, resulting in multi-mode propagation. (b) A zoomed-in view of one of the channels (channel 6 left-to-right). We observe no measurable cross-talk compared to the baseline in this measurement.

Since the ratio of the waveguide output spot size to spot spacing is designed to match the requirements at the ion position, the waveguide output can be imaged onto the ions using a global linear optical system of fixed magnification. In this experiment, we aim to work with 4 µm ion spacing. This provides a good compromise between having a large enough ion spacing to allow for low cross-talk addressing and readout while maintaining a high level of confinement in the axial direction. Therefore, given the 8 µm channel spacing of the chip,

we use an optical system with a scaling factor of 1/2.

In addition to achieving the correct magnification, the imaging system must be designed such that the $\vec{k}$-vectors of all beams are parallel to each other, as shown in Fig. 5.5. This ensures that the laser beams used for quantum gates couple to the same modes for all ions.

We can then choose the beam directions so that we only couple to the radial motion of the ions. This minimises the sensitivity of the gate fidelity to the axial modes, which are lower in frequency and, as a result, have higher heating rates and higher temperatures after cooling. We can further choose the beam directions to only couple to one set of radial modes, which makes decoupling of the spin from all motional modes at the end of an entangling gate much easier.

Since the $\vec{k}$-vectors of the beams at the photonic chip output are parallel to each other, this requirement translates into the requirement for a bi-telecentric optical system.

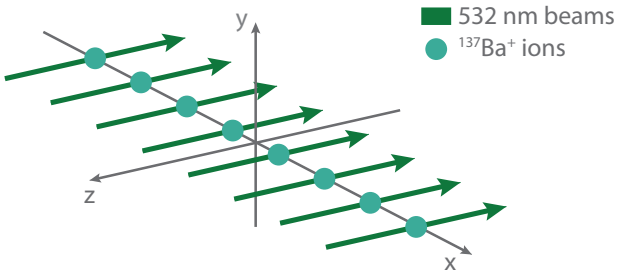


Figure 5.5: Telecentricity requirement for the optical system used to image the photonic chip output onto the ion chain. The ion chain is parallel to the x axis. The motion of the ions in the chain can be decomposed into three sets of normal modes along the x, y, and z axes. To ensure all 532 nm beams couple to only one set of motional modes, their $\vec{k}$-vectors should be parallel to each other and to the motional mode of interest (here along the z axis).

Optical system design

At the trap, the 532 nm beams from the waveguides are overlaid with the 493 nm imaging beam path shown in Fig. 4.7 using a dichroic mirror and are then focused on the ions using the $EFL = 4$ mm objective as shown in Fig. 5.6. In principle, utilising another objective with $EFL = 8$ mm would be the simplest solution to

achieving a 0.5 scaling factor. However, this solution cannot satisfy the system's bi-telecentricity requirement. This is only satisfied when the distance between the principal planes of the objectives is equal to the sum of their effective focal lengths, i.e. 12 mm. This cannot be achieved due to the mechanical design of the objectives. Hence, we introduce two additional lenses to form a system that satisfies both the requirements for magnification and telecentricity. Each of the two pairs of an objective and a lens is set up in a way that forms a bi-telecentric imaging system. Therefore, we use two bi-telecentric systems in series.

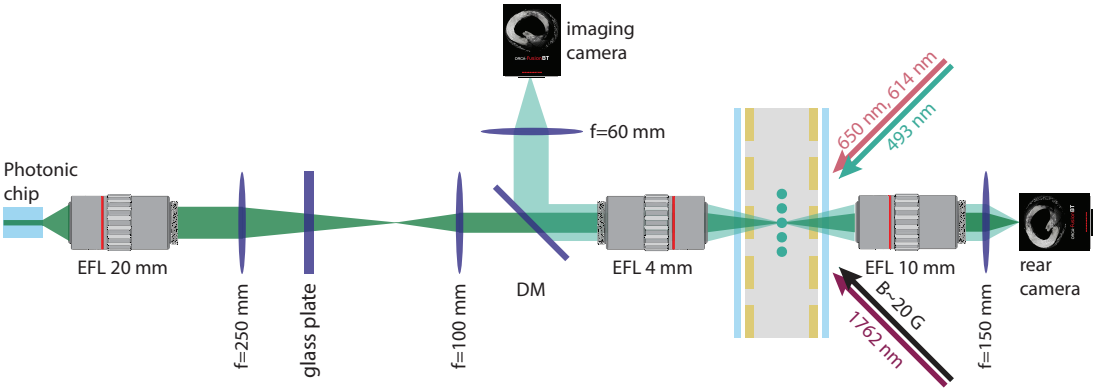


Figure 5.6: Trapped-ion setup with the individual optical addressing system. Light from the ions at 493 nm is collimated using a Mitutoyo NA0.5 objective with an effective focal length (EFL) of 4 mm and refocused onto an sCMOS camera for fast spatially-resolved readout. The 532 nm light from the photonic chip is first collimated using a Mitutoyo objective with an EFL of 20 mm to minimise aberrations. The focal lengths and positions of the subsequent $f = 250$ mm and $f = 100$ mm lenses are chosen to achieve the required 2:1 scaling and satisfy the telecentricity requirement of the optical system. A dichroic mirror (DM) overlaps the 532 nm beam path with the 493 nm fluorescence beam path. This allows us to focus the 532 nm light onto the ions using the same NA0.5 objective we use for imaging them. A 3.5 mm N-BK7 glass plate is used to fine-tune the beam positions at the trap with nanometre precision. A pair of mirrors (not shown) placed after the glass plate provide the necessary degrees of freedom for larger adjustments.

In the setup shown in Fig. 5.6, the focal lengths of the lenses used are chosen to accommodate the mechanical restrictions of the setup, while allowing for achieving the required magnification and telecentricity. We use a microscope objective with $EFL = 20$ mm and $NA = 0.2$ to collimate the photonic chip output beams with minimal aberrations. The combination of the $EFL = 20$ mm objective and the $f = 250$ mm lens magnifies the waveguide output by a factor of

12.5. The distance between the principal planes of the objective and the lens is set to 270 mm to ensure parallel $\vec{k}$-vectors of the beams at the focal plane of the $f = 250$ mm lens. The subsequent $f = 100$ mm lens and $EFL = 4$ mm objective achieve a demagnification by a factor of 25 resulting in a total scaling factor of 0.5. As before the distance between the principal planes of the objective and the lens is set to 104 mm to satisfy telecentricity (see Fig. 5.5).

Alignment procedure

To optimise the alignment of the optical system, we need a sufficient number of degrees of freedom. The position of the trap objective is optimised by looking at the ion image on the imaging camera. The $f = 100$ mm lens is rigidly mounted to the trap objective at the required distance to satisfy the telecentricity requirement. We now describe the procedure we use to optimise the alignment by using the degrees of freedom of the remaining components.

First, the VGA and waveguide assembly is mounted onto a tip/tilt stage[7]. This ensures that the output plane of the waveguide is parallel to the waveguide objective, and hence all spots are in focus. It also ensures that the output $\vec{k}$-vectors are perpendicular to the principal plane of the waveguide objective. The waveguide objective itself is placed on an xyz positioning stage to enable centering of the waveguide output onto the aperture of the objective as well as fine-tuning of the focus of the objective. This part of the alignment is performed by redirecting the beams after the $f = 250$ mm lens in Fig. 5.6 to a separate camera and minimising both the spot size and the aberrations of the system for all 8 spots simultaneously.

Next, for initial rough alignment of the addressing system onto the ions, we use a rear imaging system that allows us to image both the ion in the trap and the waveguide output as shown in Fig. 5.6. Now that the positions of the tip/tilt positioners and the waveguide objective positioners have been set in the previous step, we use a pair of mirrors to overlap one of the waveguide channels with

[7]Newport 9063-PY-M

the ion. We then adjust the position of the $f = 250$ mm lens along the optic axis to ensure the beams are focused at the ion position. While this alignment is good enough to generate a signal at the ion, due to imperfections in the rear imaging system a perfect overlap of the waveguide spot and the ion on the rear camera is not sufficient to optimise the alignment. That's why, once the alignment on the rear camera is good enough for us to see the effect of the beam on the ion, we switch to using the ion signal instead as a measure of the quality of the alignment. In particular, we use the AC Stark shift introduced by a single 532 nm beam onto the ion to quantify the intensity of the beam at the ion position. Details on exactly how this works and the experiment sequence used are provided in Section 5.3. We start by adjusting the positions of the pair of mirrors to maximise this frequency shift. Once we have reached the optimal positions, we scan the focus of the waveguide objective to maximise the frequency shift. The position of this objective is in 2:1 correspondence with the beam position at the trap. Hence moving the focus will inevitably also cause displacement of the objective perpendicular to the optic axis. Therefore this procedure has to be done iteratively, i.e. for each focus setting of the objective we also reoptimise the positions of the two mirrors.

Finally, to achieve alignment with nanometre precision (required due to the $\sim$ 1 µm diameter beams), we use an N-BK7 glass plate of thickness 3.5 mm mounted on a precision motorised mirror mount[8] as shown in Fig. 5.6. By adjusting the angle of this plate, we introduce a parallel displacement of all beams, where based on simulations we expect a displacement of 0.72 µm/°. The resolution of this mount is 0.001° corresponding to a step size at the trap of 0.7 nm. A measurement of the axial beam centre position at the trap as a function of the angle of this plate is shown in Fig. 5.7 and confirms our estimate. We measure a gradient of 0.69(7) µm/° which is in very good agreement with simulations within the errorbar.

[8] CONEX-AG-M100D

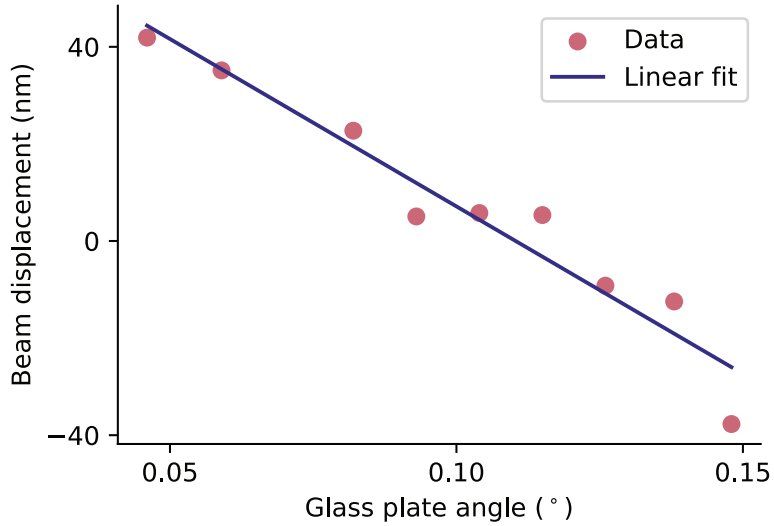


Figure 5.7: Displacement of the beam centre at the trap as a function of the angular displacement of the glass plate in Fig. 5.6. The distribution of datapoints along the vertical axis is dominated by positional drifts in the optical system between consecutive measurements.

5.3 Cross-talk characterisation using a single trapped ion

While the measurement of the device cross-talk using a camera presented in Fig. 5.4 shows excellent performance, aberrations in the optical system used to reimage the waveguide output onto the ions, as well as light leakage into neighbouring channels, can cause coherent cross-talk between channels. It is, therefore, important to characterise the performance of the entire addressing system at the ion position.

In this section, we use a single trapped $^{137}Ba^{+}$ ion to characterise the performance of the addressing system. We use the fact that when a single 532 nm beam is applied onto the ion, it introduces a differential AC Stark shift on the quadrupole transitions that is proportional to the intensity of the beam. By shuttling a single ion across the trap and measuring this AC Stark shift as a function of position, we can measure the beam waist radius and spacing as well as the cross-talk of the device. Accurately measuring the addressing system's cross-talk requires a very high measurement dynamic range spanning several orders of

magnitude in beam intensity. This requires careful design of the measurement procedure and calibration routines, as we will see below.

5.3.1 Robust phase estimation protocol

Let's consider a single trapped $^{137}Ba^{+}$ ion and for the purpose of this chapter only, we define the qubit states $|0\rangle = |S_{1/2}, F = 2, m_F = 0\rangle$[9] and $|1\rangle = |D_{5/2}, F = 4, m_F = +1\rangle$[10]. When a single 532 nm beam is applied onto the ion, it introduces an AC Stark shift on the frequency of the transition between $|0\rangle$ and $|1\rangle$. This effect can be represented as a Z_ϕ rotation with an angle of rotation $\phi \propto \delta_{\mathrm{ac}}\tau_{\mathrm{d}}$ where δ_{ac} is the amount of frequency shift and τ_{d} is the duration for which the beam is applied. In order to calculate the frequency shift, and hence the intensity of the 532 nm beam, it is therefore necessary to calculate the phase of this Z-rotation.

To measure this phase we employ a robust phase estimation (RPE) protocol as described in [93] and demonstrated in [94]. This protocol enables accurate and efficient determination of the amount of phase acquired in the presence of additive errors such as errors accumulated during state preparation and measurement. The RPE protocol relies on the system having a long coherence time, thus enabling application of the unitary that needs to be characterised multiple times in series. That way small variations (in this case small phases introduced to the qubit state) accumulate coherently into large observables whose expectation values can be easily measured. In this section we only look into how RPE works for estimating the phase while ignoring decoherence effects in the system for the time being. We leave the discussion of decoherence for Section 5.3.2.

For the purposes of this measurement we only need to estimate the angle of a Z rotation. We therefore only require good state preparation into the state $|{\rightarrow}\rangle = (|0\rangle + i\,|1\rangle)/\sqrt{2}$ and measurement in both the x and y bases on the Bloch sphere as outlined in [93]. The RPE sequence is shown in Fig. 5.8. We prepare the ion into the state $|0\rangle$ using the state preparation sequence outlined in Section

[9] This is the state we end up in after state preparation.

[10] We choose this state as $|1\rangle$ here as this corresponds to the least magnetically sensitive quadrupole transition that we are able to drive with the beam geometry shown in Fig. 4.11.

3.3.3. We then apply a $\pi_x/2$ rotation about the x axis using the 1762 nm laser to prepare the ion into the state $|{\rightarrow}\rangle$. The 532 nm beam is then turned on for a duration $\tau_{532} = 2^k\tau_0$ where $k = 0, 1, 2, ..., N$ and τ_0 is the *base* duration. The base duration τ_0 sets the maximum light shift that can be unambiguously resolved, $\delta_{\text{ac, max}} = 1/\tau_0$. In the absence of any errors the maximum exponential multiplier N sets the root-mean-square error (RMSE) of the phase estimate, $\epsilon_\phi \propto 1/2^N$ as shown in Ref. [94]. After the end of the 532 nm pulse the state of the ion is $\left(|0\rangle + ie^{i\phi_{\text{ac}}}|1\rangle\right)/\sqrt{2}$ where the accumulated phase is $\phi_{\text{ac}} = \delta_{\text{ac}}\tau_{532}$. We then apply another $\pi_\phi/2$ pulse whose axis of rotation is defined by the phase ϕ, followed by a measurement of the projection of the final state of the ion onto the initial state $|0\rangle$. The phase ϕ of the final $\pi_\phi/2$ pulse defines its axis of rotation, and hence the effective measurement basis. In the case of $\phi = 0$ we perform a measurement in the x basis, and in the case of $\phi = -\pi/2$ we perform a measurement in the y basis. In the case of an x basis measurement, the expectation value is $p_x = (1 - \cos\phi)/2$ and in the case of a y basis measurement, the expectation value is $p_y = (1 + \sin\phi)/2$. We can therefore use these two measurements to estimate the value of $\phi_{\text{ac}} = \arctan\frac{2p_y-1}{1-2p_x}$.

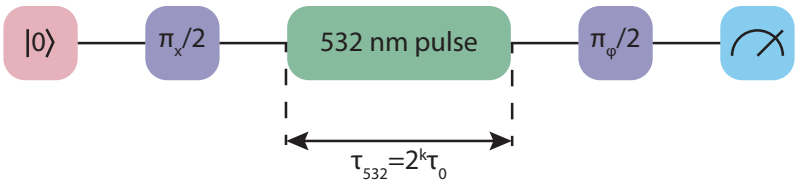


Figure 5.8: Robust phase estimation (RPE) sequence for estimating the phase introduced on the $|{\rightarrow}\rangle$ state as a result of the AC Stark shift introduced by a single 532 nm beam. The ion is prepared in the state $|0\rangle$ and the subsequent 1762 nm $\pi_x/2$ pulse prepares the state $|{\rightarrow}\rangle$. Then the 532 nm beam is turned on for a duration $\tau_{532} = 2^k\tau_0$ where τ_0 is the *base* duration defining the maximum observable AC stark shift of $\delta_{\text{ac, max}} = 1/\tau_0$. The final 1762 nm $\pi_\phi/2$ pulse defines the measurement basis. The measurement is performed for $\phi = 0, \pi/2, \pi, 3\pi/2$ to minimise the effect of polarising noise in one direction.

In this experiment we extend the protocol by measuring in the x, $-x$, y, and $-y$ bases by using $\phi = 0, \pi/2, \pi, 3\pi/2$, respectively. This helps reduce errors due to noise that might bias the measurement in one direction. We denote the

resulting expectation values by p_0, p_π, $p_{\pi/2}$, and $p_{3\pi/2}$ respectively. Then the phase estimate is obtained from $\phi_{\text{ac}} = \arctan \frac{p_{3\pi/2} - p_{\pi/2}}{p_\pi - p_0}$.

At each exponential multiplier $k > 0$, the phase is estimated by using the result from the previous multiplier $k - 1$ and constraining the shift of the new estimate from the previous one to $\pi/2^k$ as described in Algorithm 1 in the supplementary material to Ref. [94] and described in Ref. [93].

5.3.2 Knill Dynamical Decoupling Sequence

As outlined in Section 5.3.1, the RPE protocol used to calculate the AC Stark shift introduced by a single 532 nm beam relies on the long coherence time of the system such that the unitary whose phase we are trying to estimate can be repeated multiple times. At the time the data for this chapter was taken, the external magnetic field was set to $\approx 20\,\text{G}$. At this field, all accessible quadrupole transitions enabled by the beam geometry in Fig. 4.11 are first-order magnetically sensitive. Hence, the coherence time of the system is affected by any sources of magnetic field noise. The coherence time of the transition between the $|0\rangle$ and $|1\rangle$ states as defined in Section 5.3.1 was ~ 200 µs at the time of the measurements in this chapter. This is far too short to resolve differences in light shift spanning several orders of magnitude, which is required for estimating the cross-talk of the device. A well-known approach to extend the coherence time of the system is using a spin-echo experiment [95, 96]. This indeed increased the viable probe duration in our system to $\sim$ 500 µs, however, this was still not sufficient for the purposes of this measurement.

A common approach in reducing the system decoherence further is using dynamical decoupling (DD) sequences that aim to attenuate the system-environment interaction [97]. The general idea of DD is to apply a sequence of control fields in cycles of some period τ_c with delay between pulses τ_d. If the delay between consecutive pulses is shorter than the relaxation time of the environment, it is in principle possible to preserve a single-qubit state in any environment.

There have been several dynamical decoupling sequences proposed. All of them typically require a large number of pulses, where a faster environment relaxation time requires a larger number of pulses. Since the fidelity of these pulses if finite, in practice the performance of different DD sequences is also limited by the pulse errors, such as imperfections in the pulse amplitude or detuning. In addition to reducing the effectiveness of DD sequences, pulse errors can also make some sequences asymmetric, i.e. their performance will depend on the initial state of the system. Therefore here we are looking for a sequence that is robust to pulse imperfections. One way to reduce pulse area errors is by using composite pulses [98]. As shown in Ref. [99], the Knill pulse shown in Fig. 5.9(a) is the most robust to amplitude and detuning errors, typically the dominant error sources in single-qubit gates. Since this pulse consists only of π pulses, it can be used as a dynamical decoupling sequence by introducing delays between these pulses. In this work we choose to work with the Knill pulse. To further increase the robustness of the sequence, we extend the Knill DD (KDD) sequence as proposed in Ref. [99] by concatenating two Knill pulses, offset in phase by $\pi/2$, as shown in Fig. 5.9(b). From here onwards, we refer to this 10-pulse block as the KDD sequence and denote the number of blocks implemented in an experiment by N_{KDD}.

By using the KDD sequence we are able to split a long probe duration into small intervals, with DD pulses in between. The number of blocks N_{KDD} required depends on the delay between consecutive pulses τ_d and the total probe duration required by the RPE protocol. As outlined above, τ_d needs to be shorter than the relaxation time of the environment. In addition, when high-frequency noise, on a time scale faster than the environment relaxation time, is present in the system, choosing $\tau_d = 1/f_{\mathrm{noise}}$ amplifies the effect of these noise sources. We observed the best performance when working with a delay of 500 µs. This was the coherence time of the system when using a spin-echo and was far from noise peaks we had observed at lower durations. Using the KDD sequence we

were able to use probe durations of up to 200 ms, a factor of 400 improvement compared to the simple spin-echo sequence.

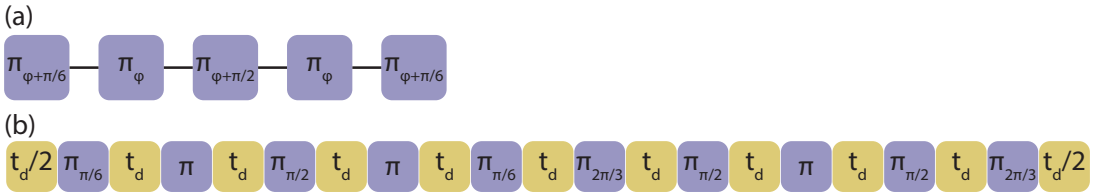


Figure 5.9: (a) A Knill composite pulse robust to amplitude and detuning errors. (b) A Knill dynamical decoupling (KDD) sequence block consisting of two Knill pulses offset in phase by $\pi/2$, with delays of length τ_d inserted between consecutive π pulses. Without loss of generality here we have set the initial phase $\phi = 0$. Depending on the choice of delay time and the total required probe duration, this block is repeated N_{KDD} times.

5.3.3 Measurement implementation

We combined the techniques discussed in sections 5.3.1 and 5.3.2 to evaluate the performance of our addressing system. Here, we were interested in measurements of the channel MFD and spacing (and thus confirming the magnification of the system) as well as the cross-talk between neighbouring channels. For this measurement, we used the 8-channel device shown in Section 5.2.2 with a target channel spacing of 8 µm and output mode field diameter of 2 µm.

The full experimental sequence is shown in Fig. 5.10. The ion is first Doppler cooled and prepared in the state $|0\rangle$ at the nominal trap centre. Afterwards, it is shuttled to the position of interest. Then, the RPE protocol is implemented as described in Section 5.3.1. Finally, the ion is shuttled back to the nominal trap centre for measurement.

The maximum observable light shift at the centre of the beams for the experiment parameters at the time the measurements were taken was 1.8 MHz. Therefore, we chose a base duration of $\tau_0 = 0.5$ µs to enable resolving shifts of up to 2 MHz. With the KDD sequence described in Section 5.3.2 we were able to probe for up to 100 ms (corresponding to a total delay of 200 ms) corresponding to a maximum duration multiplier in the RPE seqence of $k = 17$. Even at the minimum number of KDD pulses ($N_{\mathrm{KDD}} = 1$), the 532 nm pulse duration

τ_{532} would have to be split into 5 pulses of duration $\tau_d = \tau_{532}/5$. In the case of $k = 0$, this corresponds to a 50 ns pulse duration. At this point, the pulse duration becomes comparable to the timescale of the AOM dynamics, which can significantly reduce the accuracy of the experiment. To overcome this difficulty, we used an adaptive approach: in cases where the 532 nm probe duration was shorter than 500 µs (the coherence time achieved with a spin-echo sequence) we embedded the RPE protocol in a spin-echo sequence; in cases where this duration was longer than 500 µs we switched to a KDD sequence as shown in Fig. 5.10. We then repeated each measurement 20 times for each pulse duration, each basis measurement, and each ion position.

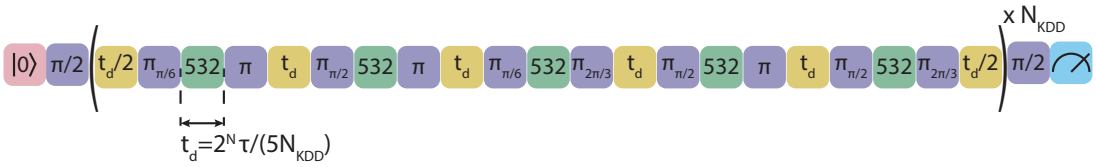


Figure 5.10: Sequence used to measure the AC Stark shift introcuced by a single 532 nm beam onto the $|0\rangle \leftrightarrow |1\rangle$ transition, combining a KDD sequence to extend the system coherence time and and RPE protocol for accurate phase estimation.

Ion shuttling

To measure beam intensity (i.e. the light shift) of the beams as a function of position, we must shuttle the ion along the trap axis. For the purposes of this measurement, we only need to shuttle the ion within the operation zone, i.e. Zone 5 in Fig. 4.2. To achieve this, we apply voltages to the electrodes defining Zones 4 and 6 of the trap that generate an electric field parallel to the trap axis. The dependence of the ion position as a function of the applied axial electric field is shown in Fig. 5.11. We observe a linear dependence of the axial position of the ion on the electric field, as expected. We measure an axial displacement of 69.7 nm/V/m.

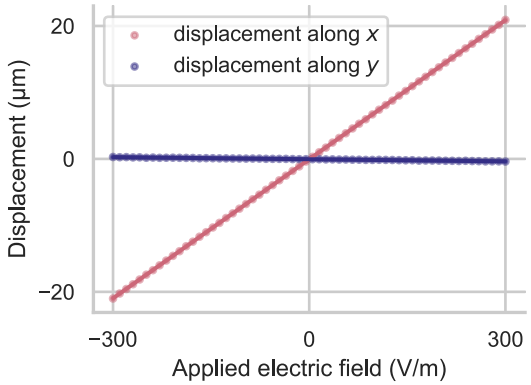


Figure 5.11: A measurement of the displacement of a single trapped ion along the x (axial) and y (radial) directions as denoted in Fig. 4.11 and linear fits in both directions. We measure an axial gradient of 69.7 nm/V/m. The small radial gradient (-1.1 nm/V/m) is caused by tilts in the imaging system and is not a displacement of the ion relative to the trap.

Quadrupole transition calibration

While the sequence used to estimate the AC Stark shift on the ion is robust to small gate errors, large gate errors due to amplitude and/or detuning errors cannot be completely compensated for and would reduce the effectiveness of the scheme.

Calibration of the parameters of the quadrupole pulses at the nominal trap centre is performed using well-known Bayesian techniques [100–103] to estimate the Rabi frequency and detuning of the transition relative to the laser frequency. This can then be used to tune the frequency of the laser to ensure we are on resonance with the transition and adjust the pulse durations to ensure we accumulate the required pulse area. However, due to the finite size of our 1762 nm beam at the trap there will be a variation in the Rabi frequency of the quadrupole transition as a function of the ion position. Furthermore, since the frequencies of the quadrupole transitions depend on the external magnetic field, any uncompensated changes in the magnetic field along the trap axis would cause the resonance frequencies of the transitions to change and therefore lead to detuning errors. In Fig. 5.12 we show plots of both the Rabi frequency and the detuning of the $|0\rangle \leftrightarrow |1\rangle$ transition from resonance (where we define 0 detuning at the nominal trap centre at $x = 0$). The Rabi frequency change is consistent

with a 62(1) µm waist diameter of the 1762 nm beam. The resonance frequency gradient is −174(2) Hz/ µm which is consistent with a magnetic field gradient of −343(4) mG/ µm across the trap axis[11]. These calibration curves are used to update the durations of the $\pi/2$ and π pulses and the detuning of the laser, as a function of the ion position, to ensure we accumulate the correct pulse area.

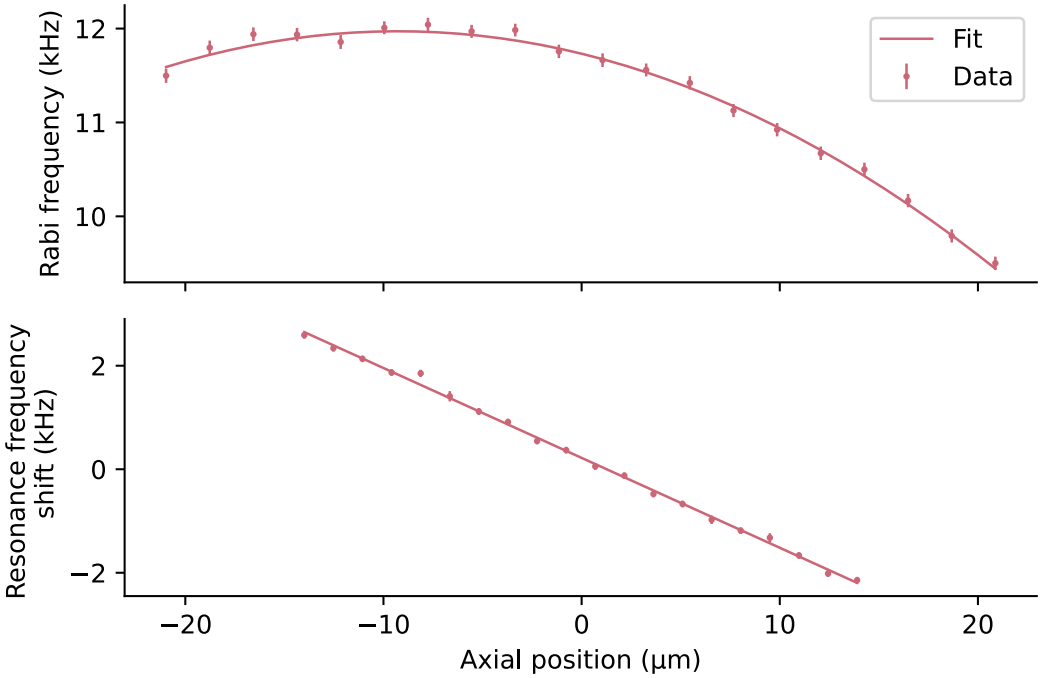


Figure 5.12: Dependence of the Rabi frequency (top) and the detuning (bottom) of the $|0\rangle \leftrightarrow |1\rangle$ transition as function of the ion's axial position.

Measurement results

The results of the light shift measurements are shown in Fig. 5.13. When measuring the performance of each channel, we keep all other channels off by turning off the corresponding fibre AOMs.

We measure a beam waist radius of 0.67(6) µm and a channel spacing of 3.95(3) µm. Both of these values are consistent with the measurement of the photonic chip output profile in Section 5.3 combined with the 2:1 demagnification of the optical system within the errorbars. For all channels, the intensity cross-talk

[11]This is not surprising: we have not taken any care to minimise this gradient in the system.

level is well below 10^{-3}. In some channels we observe a cross-talk of $\sim 10^{-5}$. This suggests that the cross-talk performance is not limited by the photonic chip design or manufacturing procedure. The variation in cross-talk channel-to-channel most likely comes from light leakage through the fibre AOMs as well as aberrations and/or scatter from the optics in the beam path that affect the channels non-uniformly.

To quantify the error introduced to the state on the nearest-neighbour ions due to cross-talk, we consider a pair of qubit states being driven on resonance with Rabi frequency Ω. Let the amount of intensity cross-talk at the nearest-neighbour positions be ϵ_I. Then the nearest-neighbour ion experiences a Rabi frequency $k\Omega$ where $k = \sqrt{\epsilon_I}$ if only one of the beams is focused onto the target ion and the other beam is global for the entire register, and $k = \epsilon_I$ if both beams are focused onto the target ion. Let's assume that the neighbouring ion is in the state $|\psi_0\rangle = a\,|0\rangle + be^{i\phi_0}\,|1\rangle$ where a, b, and ϕ_0 are real numbers, $a^2 + b^2 = 1$, and we are ignoring global phases. If we turn on the interaction on resonance with the qubit transition for some time τ we implement the unitary [38]:

$$U = \begin{pmatrix} \cos\left(\frac{k\Omega\tau}{2}\right) & i\sin\left(\frac{k\Omega\tau}{2}\right)e^{-i\phi} \\ i\sin\left(\frac{k\Omega\tau}{2}\right)e^{i\phi} & \cos\left(\frac{k\Omega\tau}{2}\right) \end{pmatrix}. \tag{5.1}$$

The resulting state is therefore

$$|\psi\rangle = U\,|\psi_0\rangle = \begin{pmatrix} a\cos\left(\frac{k\Omega\tau}{2}\right) + ib\sin\left(\frac{k\Omega\tau}{2}\right)e^{i(\phi_0-\phi)} \\ ia\sin\left(\frac{k\Omega\tau}{2}\right)e^{i\phi} + b\cos\left(\frac{k\Omega\tau}{2}\right)e^{i\phi_0} \end{pmatrix}. \tag{5.2}$$

The fidelity of the state of this ion after the operation is then given by

$$\mathcal{F} = |\langle\psi|\psi_0\rangle|^2 = \cos^2\left(\frac{k\Omega\tau}{2}\right) + 4\sin^2\left(\frac{k\Omega\tau}{2}\right)\cos^2\left(\phi_0 - \phi\right)a^2b^2, \tag{5.3}$$

and the corresponding error on the state is given by

$$\epsilon = \sin^2\left(\frac{k\Omega\tau}{2}\right)\left[1 - 4\cos^2\left(\phi_0 - \phi\right)a^2b^2\right]. \tag{5.4}$$

We see that the worst-case error occurs when $a = 0$, $b = 0$, or $\phi_0 - \phi = \pi/2$. Therefore the worst-case error is simply

$$\epsilon = \sin^2\left(\frac{k\Omega\tau}{2}\right). \tag{5.5}$$

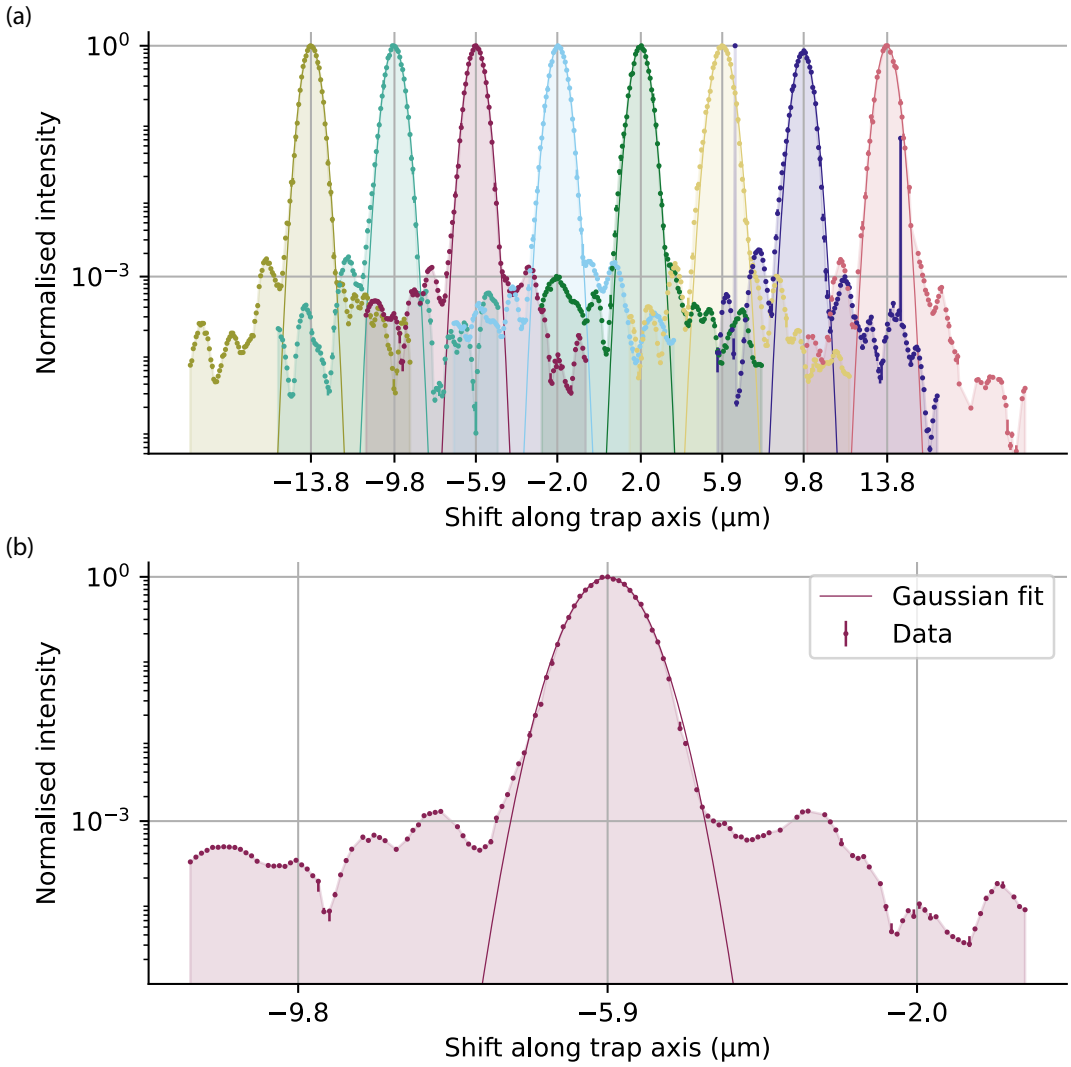


Figure 5.13: Intensity measurement of the 532 nm beams at the ion position obtained by shuttling a single ion across the trap axis in steps of 69.7 nm and measuring the AC Stark shift introduced on one of the quadrupole transitions. The data from each of the channels is normalised relative to the maximum intensity of that channel. (a) A measurement of all 8 channels. The two outliers in channel 7 (left-to-right) are caused by experiment instabilities, most likely jumps in the frequency of the 1762 nm laser. (b) A zoomed-in view of the third channel [left-to-right in (a)]. This is the channel we typically use when working with a single ion in the trap. In most cases, the errorbars are smaller than the marker size.

In the case of a π-pulse implemented on the target ion such that $\Omega\tau = \pi$, the error is $\epsilon = \sin^2(k\pi/2) \approx k^2\pi^2/4$ if $k \ll 1$. This corresponds to an error per π-pulse of $\epsilon = \epsilon_I\pi^2/4$ if only one of the beams is focused onto the target ion and the other beam is global for the register; and $\epsilon = \epsilon_I^2\pi^2/4$ if both Raman beams are focused on the target ion.

In Fig. 5.14 we plot the error on the qubit state in both of these cases for the data shown in Fig. 5.13. For each channel we sum the contributions from both neighbouring channels (except for the edge ones where there is a single contribution). In both cases this error is much lower than, or comparable to, the lowest achieved two-qubit gate error in a trapped-ion platform of 6×10^{-4}[4] and will therefore hardly limit the device performance. In addition, if further reduction of this error is necessary, well-known techniques such as coherent cancellation [104] or composite pulse sequences [98] can be trivially implemented.

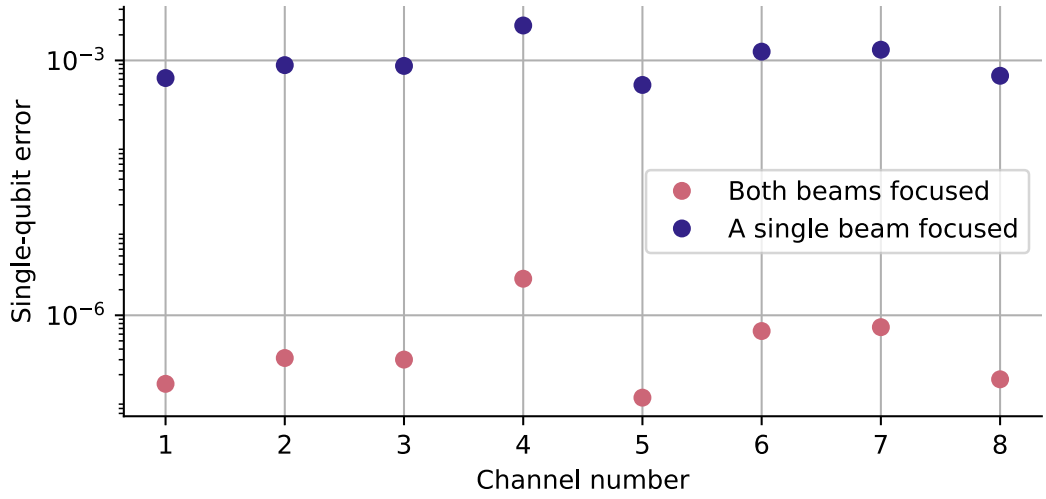


Figure 5.14: Estimated single-qubit error due to intensity cross-talk in the device accounting for contributions from all nearest-neighbour channels. The errorbars are smaller than the marker size.

5.4 Summary

We have presented a scalable and configurable way for individual optical addressing in chains of trapped atomic ions. We have designed and fabricated a novel waveguide chip that offers a very high refractive index contrast between the core and the cladding, achieving very low cross-talk at a small channel separation compared to previously demonstrated individual addressing approaches. The photonic chip is coupled with low loss to a fibre network that enables individual phase, frequency, and amplitude control of each channel, thus enabling parallel

operations. The photonic chip manufacturing procedure is very flexible in terms of the output pattern configuration. It can, therefore, be used to manufacture devices suited to most trapped-ion setups, including creating chips with a non-uniform channel spacing.

Measurements of the performance of our individual addressing system, optimised for 532 nm light, using a single trapped $^{137}Ba^{+}$ ion showed cross-talk well below 10^{-3} across the chip. The corresponding estimated worst-case nearest-neighbour error is $\sim 10^{-6}$ if both beams are focused on a single ion or $\sim 10^{-3}$ if one beam is focused on the ion and the second beam is global for the register. In the latter case, well-known techniques such as composite pulses or coherent cancellation can further reduce this error.

The procedure used to manufacture the SPIM-WGs is highly flexible, so it can be used to create devices with a large number of channels and devices with non-uniform channel spacing. The latter capability is essential for applications in ion traps that do not have sufficient degrees of freedom to generate anharmonic potentials to keep the ion spacing uniform across the ion chain or where harmonic potentials are desired. Preliminary results from characterising devices with 32 channels show no performance degradation in terms of cross-talk and minimal change in propagation losses. In addition, the same technique can be used to manufacture devices optimised for use at other optical wavelengths and, therefore, be integrated into setups using different ion species and gate mechanisms.

6

High-fidelity heralded state preparation and measurement

Contents

6.1 SPAM protocol overview 84
6.2 Protocol implementation in $^{137}Ba^{+}$ 87
6.3 High-fidelity population detection 89
6.3.1 Optical detection 90
6.3.2 Probability of decay of the $D_{5/2}$ level during detection 96
6.3.3 Total expected SPAM error 99
6.4 Experimental results 100
6.4.1 SPAM results 100
6.4.2 SPAM data rates 107
6.5 Summary of SPAM results 109

In this chapter, we describe a new protocol for the state preparation and measurement (SPAM) of qubits based on standard SPAM techniques combined with a series of mid-circuit measurements. These measurements detect errors during the SPAM protocol and reject the corresponding shots. As a result, any errors during SPAM result merely in a reduction of the experiment data rate. The protocol's success depends only on the presence of two distinct level manifolds and the ability to detect population in one of these manifolds with a very low error without changing populations in the other manifold. Therefore,

it applies to many quantum computing platforms, such as trapped ions, neutral atoms, or NV centres. Furthermore, the protocol only requires capabilities readily available in most quantum computing experiments: standard techniques for SPAM together with a set of population transfer pulses between the two manifolds. It can thus be implemented in existing setups with minimal hardware overhead. By implementing the protocol in a single trapped $^{137}Ba^{+}$ ion, we demonstrated the lowest ever SPAM errors of any qubit platform for all three qubit modalities in $^{137}Ba^{+}$ – optical (O): $7(4) \times 10^{-6}$, metastable (M): $5(4) \times 10^{-6}$, and ground (G): $8(4) \times 10^{-6}$.

In the case of a quantum computer based on an atomic system, such as a trapped ion, state preparation is typically performed by driving the electron population into a well-known state using optical pumping [48, 105–107]. The qubit state can be measured by applying a laser on the atomic system that couples only one of the qubit states, say $|0\rangle$, to a short-lived excited state. The number of photons scattered as a result of the decay of this excited state can then be used to infer whether the qubit was in the state $|0\rangle$ (presence of photons) or $|1\rangle$ (lack of photons).

Errors in SPAM can occur in several ways. First, we may prepare the wrong qubit state, for example, due to imperfect optical pumping in the case of a trapped ion. This could be due to polarisation impurities in the light used for optical pumping or off-resonant excitation through unwanted states, resulting in population outside of the desired qubit state. Second, we may detect the wrong qubit state. This could happen, for instance, due to a poor detection system resulting in a high signal-to-noise ratio or relaxation of the qubit during detection. Third, the population may be left outside the qubit manifold, for example, due to imperfect state preparation or population leakage during the quantum algorithm.

Several techniques have been employed in the past to overcome these sources of error. For example, Ref. [2] and Ref. [108] use a series of coherent population transfer pulses to improve upon the optical pumping fidelity during state preparation in a trapped-ion system. This technique has resulted in the previous record

SPAM error of $9.0(13) \times 10^{-5}$ [108]. Additionally, there have been demonstrations in a few different quantum computing platforms of using mid-circuit measurements to detect errors due to imperfect state preparation or population leakage outside of the qubit manifold and removing these errors by post-selecting the shots in which they occurred, converting them into erasures, or repeating the experiment until a successful outcome is obtained (repeat-until-success, or RUS) [39, 109–114]. For example, Ref. [39] uses a population detection step after state preparation to verify the successful preparation of a metastable level qubit in a trapped $^{171}Yb^{+}$ ion. Furthermore, Ref. [111] suggests using population detection to check for population leakage outside of the metastable level of a trapped-ion qubit due to a scattering event[1] through the $P_{3/2}$ level resulting in population in the $S_{1/2}$ or $D_{3/2}$ level. However, the method proposed in Ref. [111] does not remove errors resulting from population leakage within the ion's metastable level.

The SPAM protocol we developed in this work builds on the idea of using erasure conversion to detect and remove errors during SPAM. The protocol is designed to detect errors due to imperfect qubit state preparation or due to population leakage outside of the qubit manifold, irrespective of the population location within the quantum system or the source of the error. We have combined this with techniques analogous to those in Ref. [39] and Ref. [111] in a single protocol. We have extended the protocol to enable its use on different qubit types, for example, optical (states in the $S_{1/2}$ and $D_{5/2}$ manifolds), metastable (both states in the $D_{5/2}$ or in the $F_{7/2}$ manifold), and ground (both states in the $S_{1/2}$ manifold) qubits in the case of an atomic system [6]. As a result, our protocol overcomes more limitations in error source, error location, or qubit type than previous proposals.

[1]Such a scattering event can occur if two-photon Raman transitions are used for quantum logic gates.

6.1 SPAM protocol overview

Our protocol integrates conventional methods for SPAM with a sequence of mid-circuit measurements designed to detect population left outside the qubit manifold by imperfect initial state preparation, by imperfect population transfer pulses, or by decay of the qubit states. Implementing this protocol requires the quantum system to have at least two levels, A and B, which are connected by population transfer channels as shown in Fig. 6.1(a), and with the following properties. First, we must be able to detect population within A using some (dissipative) physical process that does not couple to states within B. Henceforth, we will call this process *population detection* to distinguish it from quantum measurement. We label the outcome of population detection as *bright* if we detect population within A and as *dark* otherwise. Second, we must be able to transfer population between states within A and states within B. Third, irrespective of the exact qubit encoding, we must be able to transfer the population from both qubit states to distinct states within B, $|0\rangle_B$, and $|1\rangle_B$ in Fig. 6.1(a), before the measurement process, unless the qubit is already encoded within B. Finally, we are ideally able to pump the majority of the population in A into some state $|SP\rangle_A$ via a physical process that also does not couple to states within B. In principle, the protocol would still work without this capability, but the number of rejected shots would increase linearly with the amount of population outside of the state $|SP\rangle_A$.

The full SPAM sequence, implemented in a quantum system with the properties described above, is shown in Fig. 6.1(b). We split the protocol into a *state preparation* part and a *measurement* part, in between which a desired quantum algorithm would be implemented. The state preparation part of the protocol starts with initialising the system in a well-known state $|SP\rangle_A$ using standard state preparation methods — for example, optical pumping in the case of an atomic system such as trapped ion [115]. This is followed by transferring the population from the state $|SP\rangle_A$ to any state in B that then allows us to prepare the desired qubit state using a series of population transfer pulses (if

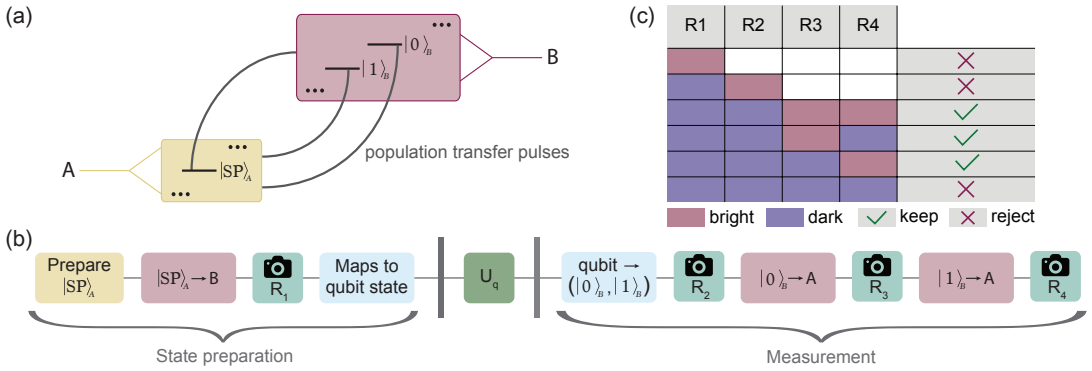


Figure 6.1: SPAM protocol. (a) Minimal requirements on the quantum system for the implementation of the protocol. (b) The complete sequence of the SPAM protocol. The camera icons represent the *population detection* steps with corresponding results R_1-R_4. The steps shaded in light blue are not necessary if the qubit is encoded exclusively in B. The step U_q corresponds to the implementation of the desired quantum algorithm. (c) A truth table showing whether a shot is accepted or rejected based on the results of the population detection steps $R_1 - R_4$ in (b).

necessary). At this point, we expect the population to be within B. To check, we perform a population detection step (R_1). We reject all bright outcomes signifying population left within A. This step thus eliminates errors stemming from both imperfect initialisation of the system in $|SP\rangle_A$ as well as imperfect population transfer from $|SP\rangle_A$ into B. We then perform a series of population transfer pulses, where necessary, to reach the desired qubit state.

We now describe the measurement part of the protocol. If the qubit is not fully encoded within B, we need to transfer the qubit into B to allow for population detection in A without destroying it. We then perform a population detection step, R_2 in Fig. 6.1(b) and reject any bright outcomes. At this stage, we expect all population to be in B, so detecting population in A, i.e. a bright measurement outcome, indicates an error. Thus, this step eliminates errors due to leakage outside the qubit manifold in A – for example, the spontaneous decay of B into A, as well as errors due to imperfect population transfer from states within A to states within B, if some of the qubit states are encoded within A. After this step, the population in the $|0\rangle_B$ state is transferred to a state in A and a population detection step, R_3 in Fig. 6.1(b), is carried out again. The outcome is used to infer the *quantum measurement result*, where a bright outcome corresponds to the

measurement result $|0\rangle$ and a dark outcome corresponds to the measurement result $|1\rangle$. Finally, the $|1\rangle_B$ state is transferred to a state in A and a final population detection step, R_4 in Fig. 6.1(b), is carried out. As the qubit has to be either in the state $|0\rangle$ or the state $|1\rangle$, either R_3 or R_4 has to result in a bright measurement outcome. Measuring dark both times suggests that the population is left in B due to imperfect population transfer to A, leakage out of the qubit manifold but still within B, or a total loss of the ion. In any case, we reject the shot.

The complete truth table for the decision to keep or reject a shot based on the outcomes of the four population detection steps is shown in Fig. 6.1(c). The multiple detection steps in this sequence allow us to detect and remove the most common sources of error in SPAM. This protocol eliminates the dependence of the SPAM fidelity on the population transfer and optical pumping quality and makes it dependent only on the population detection quality. This is very beneficial since, in most systems, detection can be made very robust.

There are only two types of physical errors that the protocol cannot detect. The first is preparation into the wrong qubit state, for example, preparing $|1\rangle$ instead of $|0\rangle$. This can happen due to off-resonant excitation during the population transfer pulses required to reach the desired qubit state after initialisation in $|SP\rangle_A$. However, this error can be made arbitrarily low by reducing the Rabi frequency of the radiation used for the population transfer pulses as well as choosing states that are sufficiently far in frequency. The second source of error is the transfer of population outside of the state $|0\rangle_B$ to A between R_2 and R_3 or during R_3. This would result in a bright outcome of R_3 leading us to wrongly conclude a measurement result $|0\rangle$. This can happen, for example, due to the decay of the state $|1\rangle_B$ to A. This error can be reduced by choosing a system with a long lifetime of B and by making the duration of the population detection pulses much shorter than the lifetime of B.

Above, we assumed that we would use the protocol with post-selection of the experiment shots based on the results of the mid-circuit measurements $R_1 - R_4$. In regular operation, this only slows down the data rate. However, the protocol

does not require post-selection and can be used as repeat-until-success (RUS), which is particularly useful in the case of large registers or sequences involving multiple mid-circuit SPAM steps, such as QEC.

The measurement part of the protocol cannot be implemented in an RUS fashion since its success is not known until the end of the algorithm after all detection steps have been performed. However, the state preparation part is compatible with RUS. In practice, this would mean performing the optical pumping and the optical detection step R_1 of the protocol and then repeating them only on qubits that resulted in a bright outcome after R_1. This could, in principle, result in an infinite number of repetitions. Therefore, the RUS attempts for state preparation must be truncated to some maximum value to give a deterministic run time with minimal truncation error. The rest of the protocol can then be used to detect errors and convert them into erasures, one of the most straightforward errors to correct when using QEC [33, 34].

6.2 Protocol implementation in $^{137}Ba^{+}$

We now describe the experimental demonstration of this protocol to prepare and measure an optical (O), a ground (G), and a metastable (M) qubit [6] in a single trapped $^{137}Ba^{+}$ ion. We used the following qubit states:

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

These qubit encodings are shown in Fig. 6.2. These states were selected for convenience, considering the transitions accessible in our setup. We note, however, that the protocol would be equally effective for any alternative pair of qubit states. In our setup, the levels A and B in Fig. 6.1 correspond to the levels $S_{1/2}$ and $D_{5/2}$ in $^{137}Ba^{+}$, respectively (see Fig. 3.2).

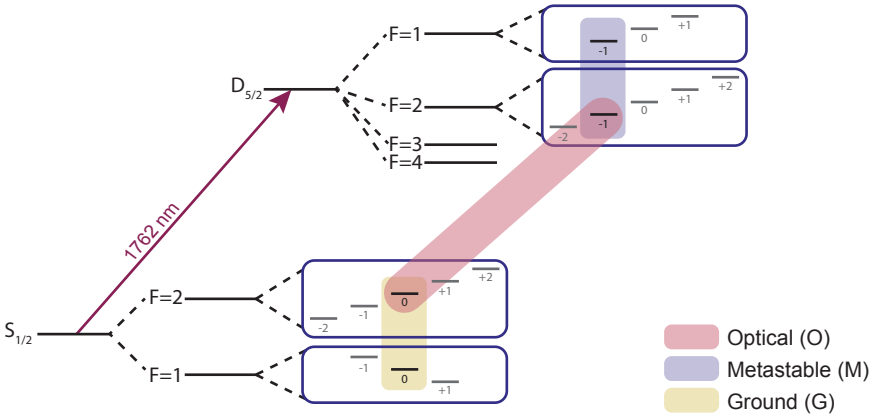


Figure 6.2: The pairs of states used as O, M, and G qubits in the demonstration of the SPAM protocol. The states in black are the qubit states. The states in grey are used during any additional population transfer pulses necessary for the implementation of the protocol. Population transfer between states within $S_{1/2}$ and states within $D_{5/2}$ is realised using coherent 1762 nm pulses, tuned on resonance with the desired transition.

The state preparation part of the protocol begins with optical pumping of the population into the state $|S_{1/2}, F = 2, m_F = 0\rangle$, as described in Section 3.3.3. Following this process, a 1762 nm pulse is used to transfer the population to a state in $D_{5/2}$ [2]. This is followed by a population detection step, R_1 in Fig. 6.1, and all measurements resulting in a bright outcome are discarded. Following this step, we use the 1762 nm laser to transfer the population to the desired qubit state where necessary. This concludes the state preparation part of the protocol.

The measurement process begins with the transfer of the population from all qubit states encoded in the ground $S_{1/2}$ level to distinct states in the metastable $D_{5/2}$ level. If a metastable qubit is used, such transfer pulses are not required. We then perform another population detection step, R_2 in Fig. 6.1, rejecting all bright results. Then the population from one of the qubit states, in our case $|0\rangle$, is transferred to a state in the ground level followed by another population detection step, R_3 in Fig. 6.1. The result of this measurement is used to determine the qubit state. Finally, we transfer the population from the $|1\rangle$ state to the ground level and perform a final population detection step, R_4 in Fig. 6.1. As detailed in Section 6.1, we reject shots where the results of both R_3 and R_4 are dark.

[2]Details on the transitions we use for each qubit encoding are provided in Section 6.4.1.

In addition to $R_1 - R_4$ in Fig. 6.1, here we perform two additional population detection steps, R_0 and R_5, at the beginning and end of the protocol to verify that the ion remained in the trap during the experiment. Hence, in addition to the rules in the post-selection table in Fig. 6.1(c), we reject any shots where the result of either R_0 or R_5 is dark. We note that the introduction of the detection steps R_0 and R_5 does not affect the SPAM fidelity – the loss of an ion would cause all detection steps $R_1 - R_4$ to result in a dark outcome, and the shot would be rejected according to the bottom row in Fig. 6.1(c). These steps were introduced to aid in understanding the individual sources of data rate reduction in the experiment analysis.

6.3 High-fidelity population detection

The success of our protocol relies only on the ability to detect population within the $S_{1/2}$ level with very high fidelity. In the case of a trapped ion, there are two contributions to the population detection fidelity that are different in nature. The first contribution stems from our ability to distinguish bright and dark detection outcomes and is mostly hardware-dependent. We refer to this as the *optical detection* fidelity, which can be increased by either higher photon collection efficiency or detection pulse duration. The second contribution is the probability of decay of the metastable $D_{5/2}$ level during the detection pulse, which would cause an ion originally encoded in the $D_{5/2}$ level to be detected as bright even if the optical detection were perfect. It should be noted that this is an asymmetric error that applies only to the qubit state which remains in the level B in Fig. 6.1(a) during the detection pulse R_3. This error is fundamentally dependent on the lifetime of the metastable level and increases with the duration of the detection pulse. Therefore, the duration of the detection pulse is a compromise between the optical detection fidelity and the probability of decay of the $D_{5/2}$ level during population detection.

6.3.1 Optical detection

We detect population within the $S_{1/2}$ level of $^{137}Ba^{+}$ by continuously applying the 493 nm laser to the ion and collecting the scattered photons, as described in Section 3.3.4. An optical detection error occurs if a bright ion (i.e. encoded within $S_{1/2}$) is measured as dark and vice versa, a dark ion (i.e. encoded in $D_{5/2}$ and remained in $D_{5/2}$ during the population detection pulse, or equivalently, the lack of an ion) is measured as bright. This error depends on the fluorescence rate from the ion, the detection efficiency of the imaging system, and the noise model of the detector, which, in our case, is an sCMOS camera.

The collection efficiency of our imaging system, limited by the NA of the imaging objective, is fixed by the system design, which is described in Section 4.2. Hence, there are three types of experiment parameters we can tune to enhance the optical detection fidelity: the laser beam parameters, which can be optimised to maximise the signal from the ion, the imaging system magnification, which changes the camera noise model, and the camera readout parameters, detection time and region of interest (ROI).

Laser beam parameters

We used the same 493 nm laser parameters for optical detection as in Section 3.3.4. We used a laser detuning of ≈ -5 MHz relative to the measured line centre to maximise the ion fluorescence signal while ensuring we did not heat up the ion enough to cause ion loss. We used a laser intensity of $\sim 100 I_{\mathrm{sat}}$. This intensity was chosen by observing the signal from the ion at a fixed laser detuning of -5 MHz and finding the point at which the signal saturated, i.e. there was no appreciable increase in the fluorescence rate with the laser power.

Imaging system magnification

When performing optical detection using a camera, the counts from each pixel are read out independently, and the ion signal is obtained by summing over the pixel counts from a chosen ROI. A larger ROI increases the amount of signal collected

from the ion but also increases the amount of noise in the measurement. This is because the pixel readout noise is broadly uniform over all pixels, so the amount of noise we add scales approximately linearly with the number of pixels in our ROI.

It is, therefore, advantageous to focus the ion image on the smallest possible number of pixels, resulting in a large signal-to-noise ratio. In our experiment, we use a microscope objective with an EFL of $f = 4$ mm to collect fluorescence from the ion (see Section 4.2). The ion image is then refocused onto the camera. Here, we chose to do this using an achromatic lens with $f = 35$ mm[3]. This is the shortest focal length commercially available lens compatible with the camera housing. A multi-lens system could be used instead if a tighter focus is required. However, as we will see, the current setup is not limiting the experiments discussed in this thesis.

Camera region of interest

To choose the camera ROI, we measured the signal from the ion using the 493 nm laser parameters given above at a camera exposure time of 500 µs. We then compared the signal-to-noise ratios for a few different ROI sizes and shapes, as shown in Fig. 6.3. We also tried weighting the counts from each ion pixel as proposed in Ref. [75]. We found that the highest signal-to-noise ratio is achieved using an ROI size of 3×3. We found that changing the pixel weights had only a small ($< 10\%$) effect on the optical detection error. These weights may need to be re-optimised if the ion image lands on different camera pixels, for example, due to mechanical drifts in the imaging system. Therefore, they may add experiment calibration complexity for minimal gain. Thus, for the rest of the measurements here, we chose to work with an ROI size of 3×3 and equal pixel weights (that is, we simply add up the counts from all 9 pixels in the ROI).

[3]Thorlabs AC254-030-A

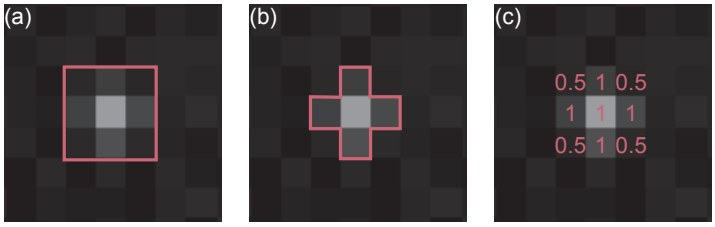


Figure 6.3: Choice of camera ROI for readout. (a) A $n \times n$ square region centred around the ion. We considered $n = 1, 3, 5$ and found that we achieve the highest signal-to-noise ratio for $n = 3$. (b) A 'cross' region. This is the only size of the cross region we considered. A smaller size would be equivalent to $n = 1$ in (a). A larger size (extending by a pixel horizontally and vertically) would be worse compared to $n = 3$ in (a) since the corner pixels in the 3×3 square region will contain more information as the ion image is equivalent to a $2D$ Gaussian distribution. (c) Using pixel weights to try and improve the signal-to-noise ratio. In the example presented here, we sum the counts from the centre cross region and then add the counts from the corners divided by a factor of two. We found that changing the pixel weights had only a small ($< 10\%$) effect on the optical detection error. Hence, for the experiments in this chapter, we chose to work with the ROI shown in (a).

Camera detection time

When using a camera for optical detection, there are two durations that we need to consider. The first is the camera *exposure time*, which is the time for which the camera measures the fluorescence from the ion. It is a compromise between the optical detection fidelity and the probability of decay of the metastable level (see Section 6.2). The second is the *frame time*[4], which is the minimum possible duration between consecutive camera triggers and is longer than the exposure time. Physically, this corresponds to the total time taken for exposure to be completed and the image pixels to be read out and stored in the camera buffer before the acquisition of the next frame can begin. Thus, the frame time increases with the exposure time[5] and the size of the image we wish to acquire. The amount of signal we collect from the ion is determined by the camera exposure time, but the probability of decay of the ion during population detection is determined by the camera frame time.

[4]This is inverse of the camera *frame rate*, measured in frames per second, fps.

[5]This is unless the exposure time is much shorter than the minimum attainable frame time for the camera in use.

Our camera supports three acquisition modes: fast, standard, and ultra-quiet. Each of these modes is associated with a different level of pixel noise and a different frame time for a fixed exposure time. The fast mode has the shortest frame time but a higher pixel noise. The ultra-quiet mode has the lowest pixel noise but requires a longer amount of time for the pixels to be read out, resulting in a much longer frame time. We empirically found that the overhead in pixel readout time in the standard and ultra-quiet modes overcompensated for any gains in the camera exposure time compared to the camera fast mode to achieve the same optical detection error. Since the error due to the decay of the metastable level depends on the frame time, we chose to work with the fast mode. We chose to work with a duration of 400 µs – this was the shortest duration at which the optical detection error we measured was much lower than the probability of decay of the metastable level, as we will shortly see.

Finally, for a given camera mode, the frame time is strongly dependent on the number of horizontal lines (i.e. the height) in an image and largely independent of the number of vertical lines (i.e. the width of the image)[6]. Since we use a 3×3 ROI per ion, and the number of horizontal lines is restricted to multiples of four for our camera, we choose to work with 8 lines. That way, we ensure we can tolerate small drifts in the ion imaging system, resulting in the ion image landing on different pixels, without having to mechanically adjust the imaging setup before each experiment. The frame time of our camera at this image size and 400 µs exposure time is $\approx$ 458.6 µs.

Optical detection error

We wish to measure the optical detection error independently of the decay of the metastable level, as this is an important component of the SPAM error model. We expect this to be much smaller than the optical pumping error.

[6]This is very useful: we only need a small number of horizontal lines to ensure we capture the ion. However, as we add more ions to the chain, we will need to use more vertical lines to capture the entire ion chain. The fact that the acquisition duration is independent of the number of vertical lines means that we expect the SPAM error using this protocol to be independent of the number of ions in the chain.

Hence, we need to design a measurement scheme carefully to ensure we are only sensitive to the optical detection error.

We will call the probability to classify an ion in the $S_{1/2}$ state as bright $P(\text{bright}|S_{1/2})$ and the probability to classify an ion in the $D_{5/2}$ level (that remained in the $D_{5/2}$ level during the population detection step) as dark $P(\text{dark}|D_{5/2})$.

To evaluate the error of a bright outcome, i.e. $1 - P(\text{bright}|S_{1/2})$, we perform three consecutive population detection steps. The error of a bright outcome is given by the probability of a dark result of the second step, provided the first and third steps result in a bright outcome. This ensures that the ion remained in the $S_{1/2}$ level and stayed in the trap during the second measurement.

To evaluate the error of a dark outcome, i.e. $1-P(\text{dark}|D_{5/2})$, we first prepare the ion into the $|S_{1/2}, F = 2, m_F = 0\rangle$ state and then transfer the population to the $|D_{5/2}, F = 2, m_F = -1\rangle$ state. We perform three consecutive population detection steps as before. The error of a dark outcome is given by the probability of a bright result after the second step, provided the first and third steps result in a dark outcome. By demanding that the first and third steps result in a dark outcome, we ensure that the ion remained in the $D_{5/2}$ level during the second step and are thus able to decouple the error due to the decay of the metastable level (see Section 6.3.2) from the optical detection error.

The results from these experiments are shown in Fig. 6.4. The threshold for discriminating between bright and dark is set by fitting Gaussian functions to the bright and dark distributions and minimising the overlap between the two fitted curves, as discussed in Ref. [75]. We observe an error on the bright state of $3.2(28) \times 10^{-6}$ and an error on the dark state of $0.0(11) \times 10^{-6}$. We believe the slightly higher error on the bright state is due to low-energy ion collisions during the measurement, causing the ion to decrystallise[7].

In Fig. 6.4, we see that the bright histogram is significantly wider than the dark histogram. The widths of the two histograms are set by all of the noise sources present in the system. We split these noise sources into two categories:

[7] At the time this data was taken we were seeing a collision once every ~ 30 min.

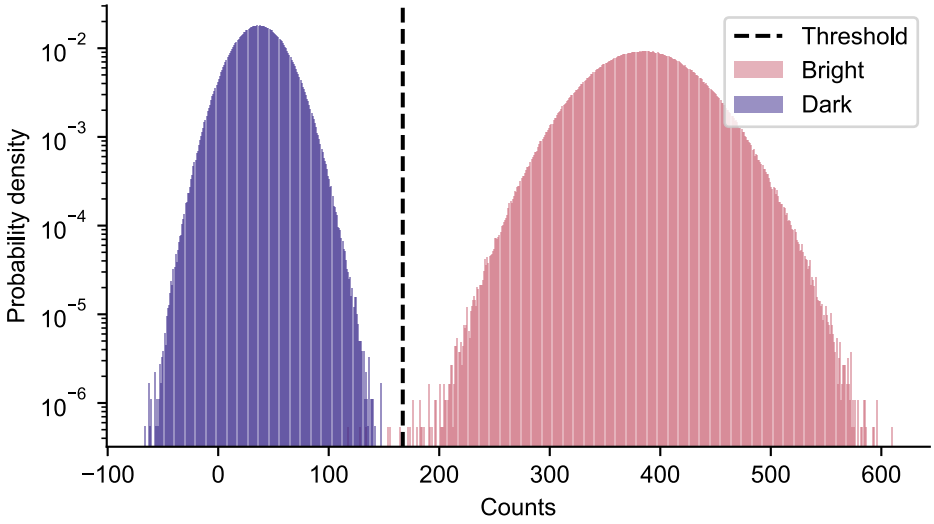


Figure 6.4: Optical detection error for discriminating between a bright and a dark state. We observe an error on the bright state of $3.2(28) \times 10^{-6}$ and an error on the dark state of $0.0(11) \times 10^{-6}$. The bright and the dark data contain $1,850,000$ shots each (before post-selection). An uncalibrated nominal camera dark count offset of 100 counts per pixel has been subtracted from each raw count value in this and following histograms, leading to occasional negative count readings.

noise introduced by the camera readout electronics, which is the same for both histograms and statistical shot noise from the collected photons, which increases with the number of photons and is, therefore, higher for the bright histogram. We now estimate the expected shot noise from the collected photons, resulting in the generation of photoelectrons. The conversion factor between the measured number of counts and the number of photoelectrons for our camera is 0.24 electrons/count. By adding this correction to the data to find the number of electrons generated and fitting a normal distribution to each histogram, we find:

- Bright mean: 92.420(8) electrons
- Bright STD: 10.400(8) electrons
- Dark mean: 8.898(4) electrons
- Dark STD: 5.322(4) electrons

If the data were limited by photon shot noise, we expect the standard deviation to be approximately equal to the square root of the mean number of counts.

Therefore, we would expect an STD of 3.0(7) for the dark histogram and 9.6(4) for the dark histogram. We see that the measured STD of the dark histogram is slightly higher than expected, while the measured STD for the bright histogram is consistent with the measured one within errorbars. This suggests that the dark histogram width is limited by camera noise, while the bright histogram width is limited by electron shot noise.

6.3.2 Probability of decay of the $D_{5/2}$ level during detection

The decay of the metastable level can cause a bright measurement outcome even if a dark state was prepared, and the optical detection was perfect. This is not a problem for R_1 and R_2 in Fig. 6.1 as it would merely result in more rejected shots. However, it can cause reading out the wrong qubit state in R_3 if the decay occurs between R_2 and R_3. The probability of decay during some time period t is given by $\epsilon_{\text{decay}} = 1 - e^{-t/\tau}$ where τ is the lifetime of the metastable $D_{5/2}$ level.

The natural lifetime of the metastable level of $^{137}Ba^{+}$ is 30.14 s. However, the observed lifetime in an experiment can be reduced by leakage of the 614 nm light used to pump population outside of the $D_{5/2}$ level.

Optimisation of the metastable level lifetime

An initial measurement of the metastable level lifetime resulted in a lifetime of ~ 100 ms. This is much shorter than the natural lifetime of the metastable level. At the time, the maximum attainable 614 nm beam power at the trap was 100 µW. The power we measured at the trap with the 614 nm beam nominally off was $\approx 40\,\text{nW}$ corresponding to a leakage of $\approx -34\,\text{dB}$. This is a very high level of leakage and it could be explained by either reflections within the laser beam path, causing light from unwanted AOM orders coupling into the fibre used to deliver light to the trap, or by RF switch leakage[8].

[8]Here we are referring to the RF switch on the Urukul board used to drive the double-pass AOM in the 614 nm beam path.

Upon inspection of the laser beam path, we found that there was a significant amount of back reflection of the zeroth order of the AOM from the surface of one of the lenses. To mitigate this, we installed a beam block in front of the lens in question, which blocked the AOM zeroth order and only allowed the first order to pass through. This increased the measured metastable level lifetime to $\sim 1\,\mathrm{s}$.

To further reduce the effect of 614 nm leakage, we reduced the maximum amount of 614 nm light going towards the trap. The power was chosen to maximise the measured metastable level lifetime while also ensuring that we could pump the population out of the $D_{5/2}$ level, when necessary, within a reasonable amount of time ($\sim$ 1 ms). For the experiments in this chapter, the maximum 614 nm power at the trap was set to 10 µW, resulting in a 1 ms deshelving duration. This deshelving duration is very long compared to our quantum logic gates. However, it is only necessary at the very end of an experiment shot: turning on the 614 nm beam otherwise would completely destroy the metastable level qubits. Since each experiment shot, in general, consists of Doppler cooling, sideband cooling, several quantum logic gates, and multiple detection steps to implement the protocol described in this chapter, we find that the 1 ms deshelving duration only adds $\approx 5\%$ to the total experiment runtime[9].

Metastable level lifetime measurement

To measure the metastable level lifetime after the optimisation above, we first initialise the ion into the state $|S_{1/2}, F = 2, m_F = 0\rangle$ and then transfer the population to the state $|D_{5/2}, F = 2, m_F = -1\rangle$ followed by a series of consecutive population detection steps, each with an exposure time of 20 ms. Once a bright measurement is obtained, the duration for which the ion was dark is recorded and the procedure is repeated again. Performing the measurement in this way allows us to measure the metastable level decay probability in the time domain. The results from this experiment are shown in Fig. 6.5. We observed a lifetime of

[9]We note that here we opted for reducing the maximum 614 nm power at the trap as it was the easiest way to reduce the effects from leakage. In future implementations, if a faster deshelving duration is needed, we could further improve back-reflections in the beam paths and improve the isolation of the RF chain.

$\tau = 27.2(2)$ s. This is slightly lower than the natural lifetime of the $D_{5/2}$ level in $^{137}Ba^{+}$ of 30.14 s. We believe that this is still limited by leakage of the 614 nm laser light used to drive the $D_{5/2} \leftrightarrow P_{3/2}$ transition.

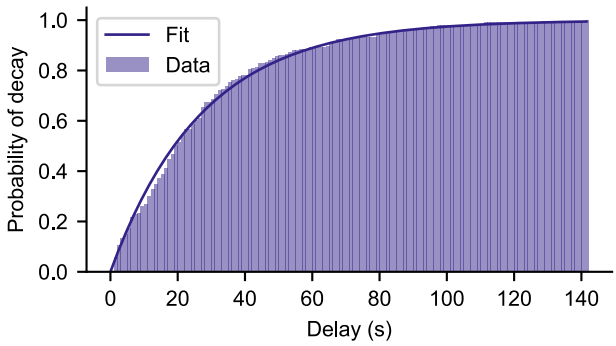


Figure 6.5: Probability of decay of the metastable level as a function of the delay time between preparation in the metastable level and measurement. The solid line is an exponential fit to the data, $\epsilon_{\text{decay}} = 1 - e^{-t_d/\tau}$. The fit corresponds to a decay time of $\tau = 27.2(2)$ s.

The population detection step R_2 in our protocol checks for decay of the metastable level during the algorithm and we reject shots where such a decay has occurred. However, a decay during the steps R_2 or R_3, which results in a bright outcome of R_3, leads to a readout error on the dark qubit state. Let's assume that for a given population detection step of length t_d (equal to the camera frame time), if the ion decays within some time t_0 from the start of the exposure, the result would be a bright measurement. If the ion decays during the last $(t_d - t_0)$ of the exposure, the result would still be a dark measurement. During the measurement part of our protocol, we implement the detection step R_2, followed by a 1762 nm population transfer pulse, followed by the detection step R_3. We will assume that the duration of the 1762 nm pulse is much shorter than the camera frame time[10]. Then the duration over which we record a decayed ion as bright during R_3 is $(t_d - t_0) + t_0 = t_d$. The first term in the brackets corresponds to the part of the detection step R_2 in which the ion would decay, but R_2 would still result in a dark outcome, and hence it would not be post-selected. The second term

[10]In our case, the 1762 nm pulses have lengths of ~ 20 µs, over an order of magnitude shorter than the camera frame time.

corresponds to the duration over which the ion can decay during R_3, and R_3 would result in a bright outcome. Hence, the contribution to the SPAM error from the decay of the metastable level is approximately equal to the probability of decay during the detection step, or $1.684(10) \times 10^{-5}$ given the camera frame time and our measurement of the metastable level lifetime in Fig. 6.5.

To independently verify this decay probability, we perform the following experiment. We first prepare the ion in the state $|\mathrm{D}_{5/2}, F = 2, m_F = -1\rangle$ as above. We then perform two consecutive detection pulses. The decay probability is given by the probability of a bright outcome from the second pulse if the first pulse results in a dark outcome. The result from this measurement is a decay probability of $1.4(6) \times 10^{-5}$. This agrees with the measurement of the metastable level lifetime and our frame time within the errorbars.

6.3.3 Total expected SPAM error

The total measurement error is composed of the optical detection error and the decay probability of the metastable level, where the latter is only relevant for the dark state. In this section, we estimate the total expected measurement error using our protocol based on the independent measurements from Section 6.3.1 and Section 6.3.2.

The error in the readout of an ion that is prepared into a bright state is simply given by the optical detection error of the bright state, i.e. $1 - P(\mathrm{bright}|\mathrm{S}_{1/2})$, and from the measurement in Section 6.3.1 this is $3.2(8) \times 10^{-6}$. The error in the readout of an ion that was prepared into the dark state is given by $\Big(1 - P(\mathrm{dark}|\mathrm{D}_{5/2}) + \epsilon_{\mathrm{decay}} \times P(\mathrm{bright}|\mathrm{S}_{1/2})\Big)$ and in our case this amounts to a total error of $1.4(6) \times 10^{-5}$. Thus, the average measurement error of the two states is given by $8.6(34) \times 10^{-6}$.

6.4 Experimental results

6.4.1 SPAM results

Metastable level qubit

The metastable level qubit in $^{137}Ba^{+}$ is naturally suited to this protocol as it requires the lowest number of population transfer pulses. The detailed experiment sequence for SPAM of the metastable level qubit, as defined in Fig. 6.2, is:

1. Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$.
2. Population transfer using a 1762 nm pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to either $|0\rangle = |D_{5/2}, F = 2, m_F = -1\rangle$ or $|1\rangle = |D_{5/2}, F = 1, m_F = -1\rangle$.
3. Population detection step R_1.
4. Population detection step R_2.
5. Population transfer using a 1762 nm pulse from $|D_{5/2}, F = 2, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
6. Population detection step R_3.
7. Population transfer using a 1762 nm pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
8. Population detection step R_4.

In Fig. 6.6, we plot the observed counts when preparing the states $|0\rangle$ (R_3 bright) and $|1\rangle$ (R_3 dark) following implementation of all steps of the protocol. The histograms demonstrate that we are able to achieve excellent population detection in our experiment, the core (and only) requirement for the success of our protocol. The SPAM error of the dark state is $1.1(7) \times 10^{-5}$, and the SPAM error of the bright state is $0.0(20) \times 10^{-6}$, resulting in a net SPAM error of $5(4) \times 10^{-6}$. This is the lowest SPAM error ever demonstrated in any quantum computing platform.

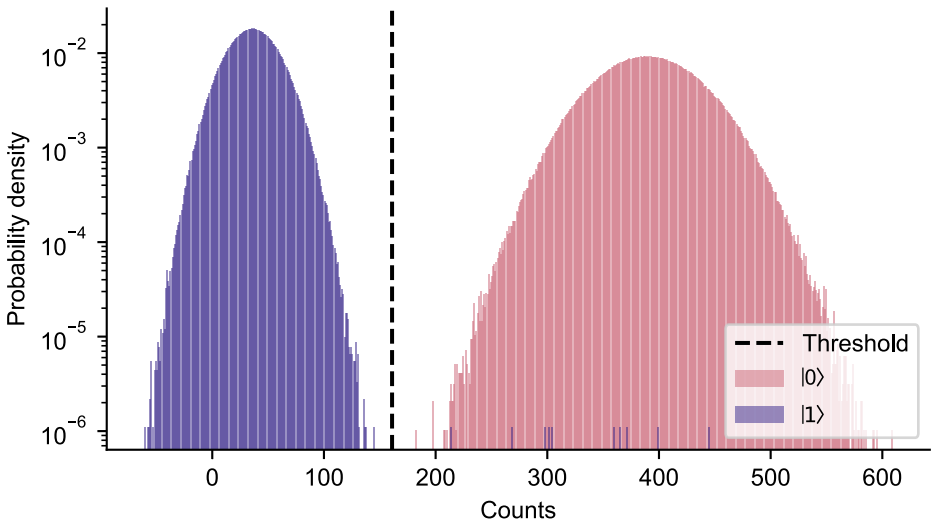


Figure 6.6: SPAM of the metastable level qubit in $^{137}Ba^{+}$. The histograms show the number of counts from the detection pulse R_3 in Fig. 6.1(b) observed after preparation of the $|0\rangle$ (pink) state and the $|1\rangle$ (blue) state after implementation of the full protocol. The data for both the $|0\rangle$ and the $|1\rangle$ states contain $1,000,000$ shots (before post-selection). The error due to the decay of the metastable level during the detection duration is present only in the measurement of the dark ($|1\rangle$, blue) state. The measured SPAM error for the bright ($|0\rangle$, pink) state is $0.0(20) \times 10^{-6}$, and the measured SPAM error for the dark state is $1.1(7) \times 10^{-5}$, giving an average error of $5(4) \times 10^{-6}$. The uncertainties represent one Wilson interval.

The measured SPAM error agrees with our initial estimate of $8.6(34) \times 10^{-6}$ within the error bars. The measured $0.0(20) \times 10^{-6}$ error of the bright state is also expected: the dominant contribution to the SPAM error is the decay of the metastable level and the bright state is not affected by this decay. If we wanted to make this error symmetric for the two qubit states, we could alternate the definition of bright and dark in each shot of the experiment.

In Fig. 6.7(a), we show how the SPAM error decreases after each post-selection step, and in Fig. 6.7(b), we plot the reduction in SPAM success probability (or equivalently, data rate) after each post-selection step. The data rate and the amount of error reduction after each step are dominated by the fidelities of the 1762 nm population transfer pulses, as we will see shortly in Section 6.4.2. We see that, while as a result of the protocol, the data rate is reduced by $< 10\%$, the SPAM error is reduced by four orders of magnitude.

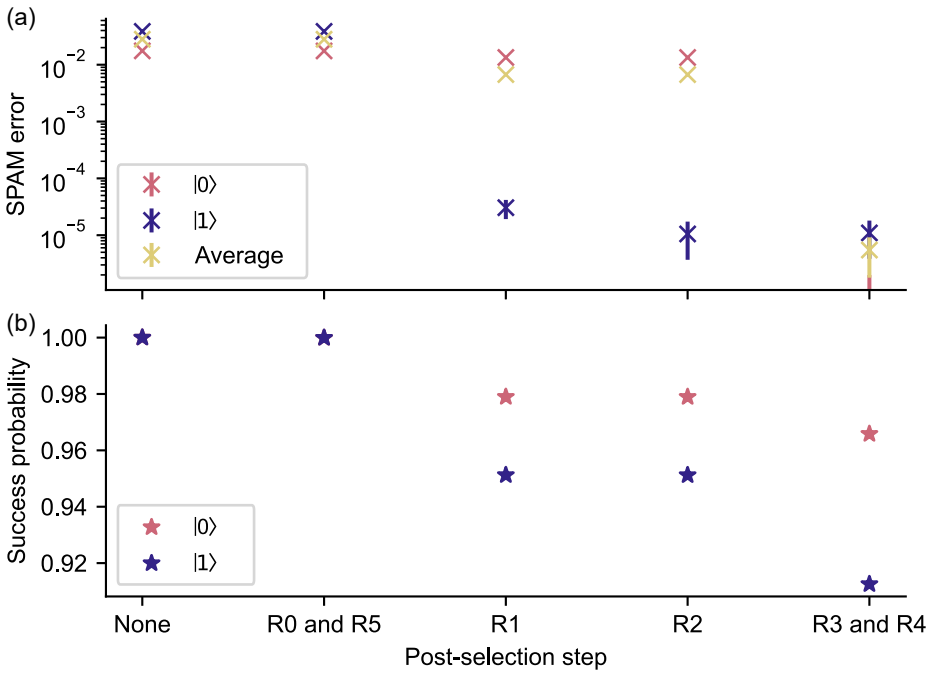


Figure 6.7: State preparation and measurement of the metastable qubit in $^{137}Ba^{+}$. (a) Reduction in SPAM error after each post-selection step. (b) Reduction in the SPAM data rate as a function of the post-selection step. If we perform no post-selection, the average SPAM error is 2.807(23)%. The post-selection steps R_0 and R_6 are used at the very start and very end of the protocol to ensure that the ion remained in the trap. We expect reasonably infrequent ion loss, thus they have a negligible effect on the data. The step R_1 checks for successful state preparation and the step R_2 checks for successful mapping of the qubit states to the metastable level and for decay of the metastable level. Finally, the steps R_3 and R_4 are used to check that at least one of the two qubit states result in a bright measurement and are hence plotted together. Thus, they ensure that the qubit states were successfully transferred to $S_{1/2}$ for measurement and that there was no leakage outside of the qubit manifold. We see that, while as a result of the protocol, the data rate is reduced by $< 10\%$, the SPAM error is reduced by four orders of magnitude. The differences between the two states arise from different fidelities of the 1762 nm pulses used to transfer the qubit states to and from $S_{1/2}$ (see Section 6.4.2).

Ground level qubit

To demonstrate the universality of our protocol to the qubit encoding, we also implemented it on the $^{137}Ba^{+}$ ground-level qubit. The exact experimental sequence for implementing the protocol on the ground level qubit, as defined in Fig. 6.2, is:

1. Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$.

2. Population transfer using a 1762 nm pulse from $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ to $|\mathrm{D}_{5/2}, F = 2, m_F = -1\rangle$.

3. Population detection step R_1.

4. Population transfer using a 1762 nm pulse from $|\mathrm{D}_{5/2}, F = 2, m_F = -1\rangle$ to either $|0\rangle = |\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ or $|1\rangle = |\mathrm{S}_{1/2}, F = 1, m_F = 0\rangle$.

5. Population transfer using a 1762 nm pulse from $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ to $|\mathrm{D}_{5/2}, F = 2, m_F = +1\rangle$.

6. Population transfer using a 1762 nm pulse from $|\mathrm{S}_{1/2}, F = 1, m_F = 0\rangle$ to $|D_{5/2}, F = 1, m_F = -1\rangle$.

7. Population detection step R_2.

8. Population transfer using a 1762 nm pulse from $|\mathrm{D}_{5/2}, F = 2, m_F = +1\rangle$ to $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$.

9. Population detection step R_3.

10. Population transfer using a 1762 nm pulse from $|\mathrm{D}_{5/2}, F = 1, m_F = -1\rangle$ to $|\mathrm{S}_{1/2}, F = 1, m_F = 0\rangle$.

11. Population detection step R_4.

This sequence is quite a bit longer compared to the sequence for SPAM of the metastable level qubit due to the additional population transfer pulses required before R_1 to verify successful optical pumping, after R_1 to reach the desired qubit state, and before R_2 to ensure the qubit is hidden in the $\mathrm{D}_{5/2}$ level before measurement.

The results from the SPAM of the ground-level qubit are shown in Fig. 6.8. The measured SPAM error is $8(4) \times 10^{-6}$, which agrees with the SPAM error we measured for the metastable-level qubit within the errorbars. This is expected: as we noted in Section 6.1, the SPAM fidelity using our protocol is determined purely by the population detection error and is insensitive to the qubit encoding

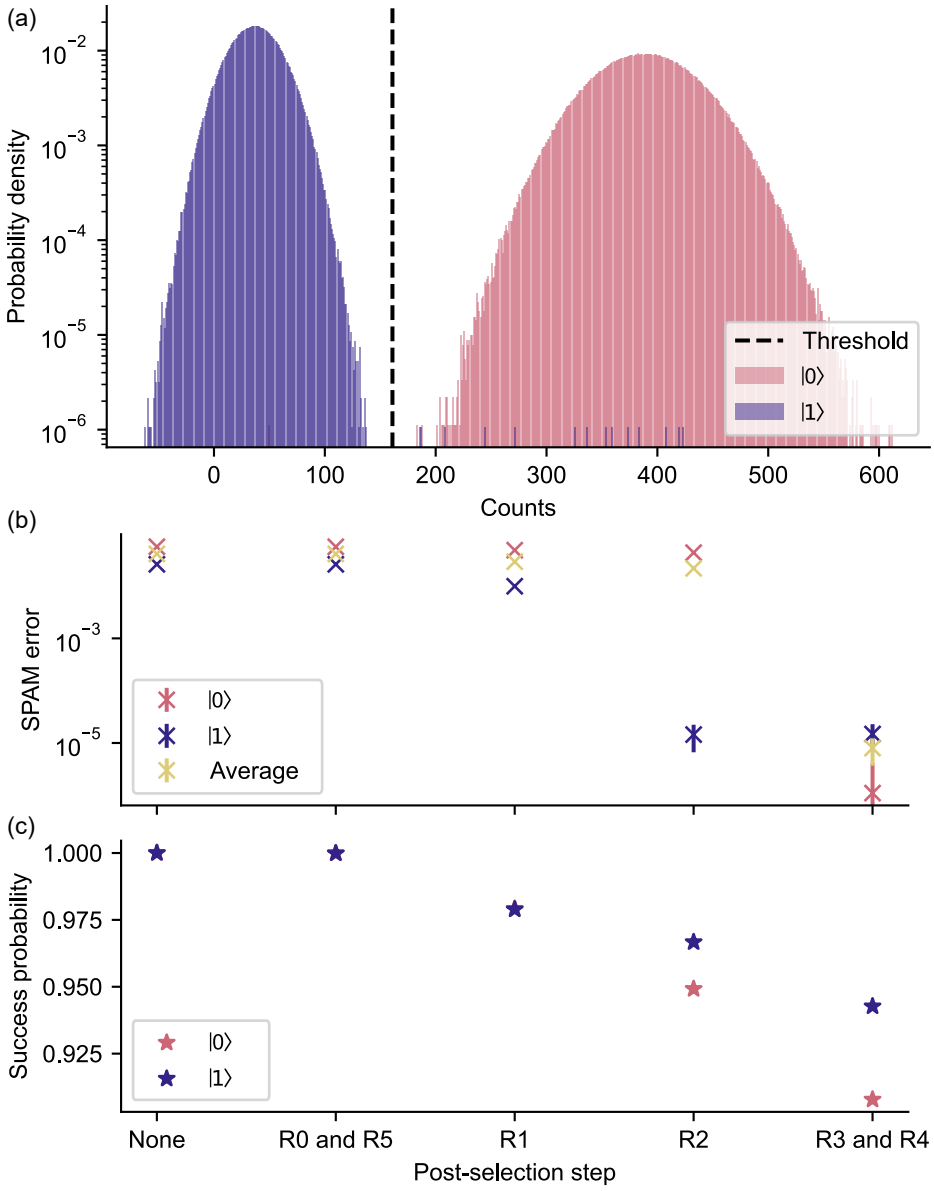


Figure 6.8: Ground level qubit SPAM data. The SPAM error is $8(4) \times 10^{-6}$. (a) The number of counts from the detection pulse R_3. The data contain $1,000,000$ shots for each state. (b) Reduction in the SPAM error after each post-selection step. (c). Reduction in the SPAM data rate after each post-selection step.

or the quality of the required population transfer pulses. As with the metastable level qubit, errors that occur during the sequence merely result in a reduction in the data rate, as can be seen from Fig. 6.8(c).

Optical qubit

Finally, we also implement the protocol on the optical qubit in $^{137}Ba^{+}$, as defined in Fig. 6.2. The detailed experiment sequence is:

1. Optical pumping into $|S_{1/2}, F = 2, m_F = 0\rangle$.
2. Population transfer using a 1762 nm pulse from $|S_{1/2}, F = 2, m_F = 0\rangle$ to either $|0\rangle = |D_{5/2}, F = 2, m_F = -1\rangle$ or $|D_{5/2}, F = 1, m_F = -1\rangle$.
3. Population detection step R_1.
4. If preparing $|1\rangle = |S_{1/2}, F = 2, m_F = 0\rangle$, population transfer using a 1762 nm pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
5. If preparing $|1\rangle$, population transfer using a 1762 nm pulse from the state $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{5/2}, F = 1, m_F = -1\rangle$.
6. Population detection step R_2.
7. Population transfer using a 1762 nm pulse from $|D_{5/2}, F = 2, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
8. Population detection step R_3.
9. Population transfer using a 1762 nm pulse from $|D_{5/2}, F = 1, m_F = -1\rangle$ to $|S_{1/2}, F = 2, m_F = 0\rangle$.
10. Population detection step R_4.

Here, only one of the qubit states requires additional population transfer pulses before and after R_1 and before R_2. The results from the SPAM of the optical qubit are shown in Fig. 6.9. The measured SPAM error is $7(4) \times 10^{-6}$. As before, the SPAM error we measure is consistent with the SPAM error of the metastable level

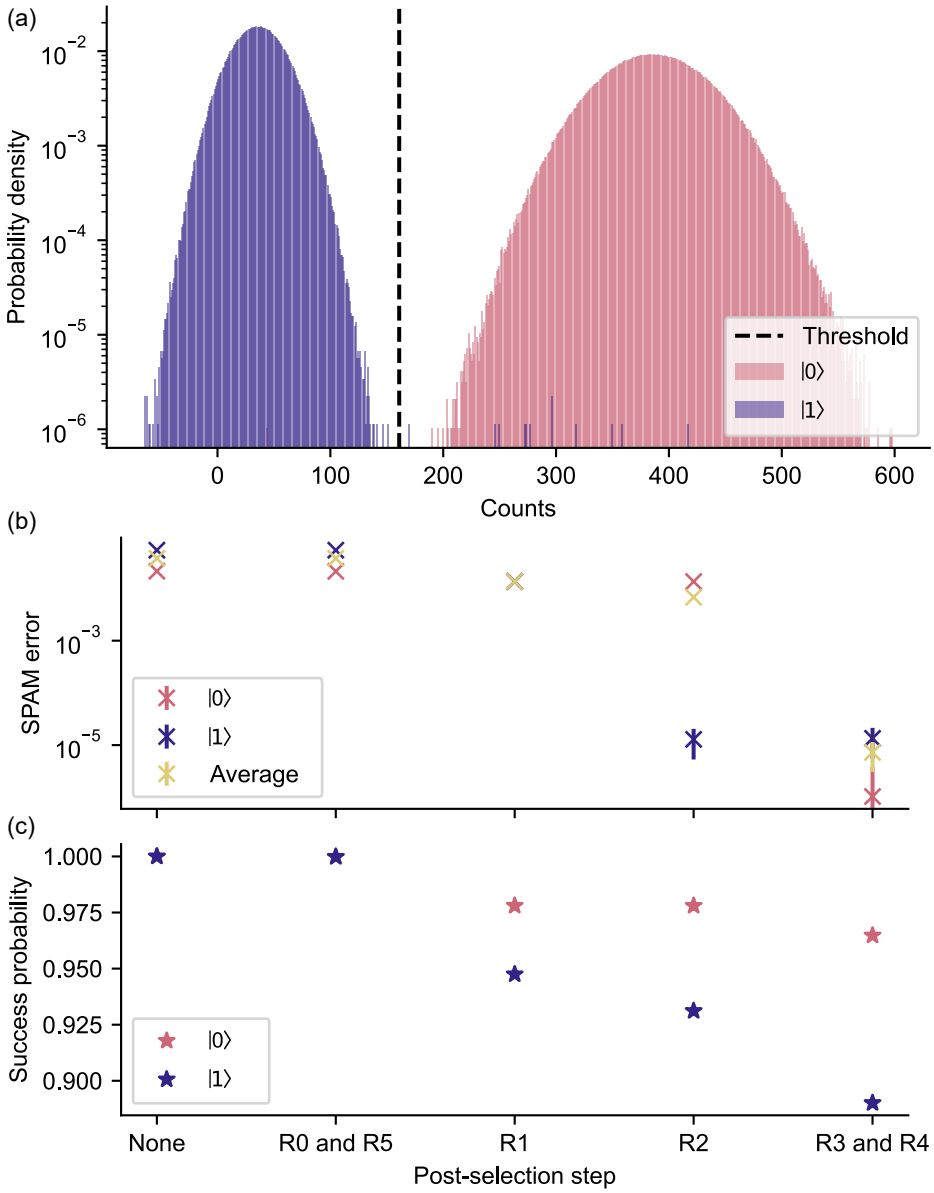


Figure 6.9: Optical qubit SPAM data. The SPAM error is $7(4) \times 10^{-6}$. (a) The number of counts from the detection pulse R_3. The data contain $1,000,000$ shots for each state. (b) Reduction in the SPAM error after each post-selection step. (c). Reduction in the SPAM data rate after each post-selection step.

qubit within the errorbars, and any errors during optical pumping or population transfer are removed by the multiple detection shots.

6.4.2 SPAM data rates

The protocol presented in this paper effectively eliminates errors acquired during standard SPAM procedures at the expense of a reduced data rate. The data rate is determined by the errors in the operations in all steps of the protocol, such as optical pumping during the state preparation part and the $S_{1/2} \leftrightarrow D_{5/2}$ population transfer pulses. To ensure we understand where the errors we post-select on come from, we want to measure the assumed error sources independently.

To measure the error of the optical pumping step, we first pump the population into the $|S_{1/2}, F = 2, m_F = 0\rangle$ state. We then perform four consecutive π pulses using the 1762 nm laser from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{1/2}, F = 1, m_F = -1\rangle$, $|D_{5/2}, F = 1, m_F = +1\rangle$, $|D_{5/2}, F = 2, m_F = -1\rangle$, and $|D_{5/2}, F = 2, m_F = +1\rangle$. Finally, we perform a detection step and measure the probability of a bright outcome. The multiple 1762 nm pulses ensure that the error due to mapping to $D_{5/2}$ is $\sim \epsilon_{1762}^4$ where ϵ_{1762} is the error of a single pulse. The optical pumping error, ϵ_{op}, is $0.80(6)\%$.

To measure the errors of the 1762 nm population transfer pulses (π pulses), we first prepare the metastable level state. That is, we implement the protocol up to and including R_1. We then perform a population transfer pulse from that state to a state in the $S_{1/2}$ level and measure the probability of a dark outcome.

The errors from all these processes are shown in Table 6.1. In what follows, we calculate the expected fraction of rejected shots at each of the protocol steps. As we can see from Table 6.1, the post-selected fraction is dominated by the errors of the 1762 nm population transfer pulses and the optical pumping error. Therefore, in estimating the expected fraction of rejected shots, we will neglect the population detection error, which, according to Section 6.3.3, is a factor of $\sim 10^4$ lower. We will denote the optical pumping error by ϵ_{op} and the error of a 1762 nm π-pulse by ϵ_{1762}.

Error type	Error (%)
Optical pumping (ϵ_{op})	0.80(6)
$\|F=2,m_F=0\rangle \leftrightarrow \|F=2,m_F=-1\rangle$	1.38(7)
$\|F=2,m_F=0\rangle \leftrightarrow \|F=1,m_F=-1\rangle$	4.73(14)
$\|F=2,m_F=0\rangle \leftrightarrow \|F=2,m_F=+1\rangle$	3.10(11)
$\|F=1,m_F=0\rangle \leftrightarrow \|F=2,m_F=-1\rangle$	1.11(7)
$\|F=1,m_F=0\rangle \leftrightarrow \|F=1,m_F=-1\rangle$	0.98(6)

Table 6.1: A breakdown of the errors during each of the steps used in the SPAM sequences for all three qubit types. The first row gives the error of preparing into $|S_{1/2}, F=2, m_F=0\rangle$ using optical pumping. This measurement is performed by first performing optical pumping, then doing four consecutive π pulses from $|S_{1/2}, F=2, m_F=0\rangle$ to $|D_{1/2}, F=1, m_F=-1\rangle$, $|D_{5/2}, F=1, m_F=+1\rangle$, $|D_{5/2}, F=2, m_F=-1\rangle$, and $|D_{5/2}, F=2, m_F=+1\rangle$, then performing a detection step and measuring the probability of a bright outcome. The multiple 1762 nm pulses ensure that the error due to mapping to $D_{5/2}$ is $\sim \epsilon_{1762}^4$ where ϵ_{1762} is the error of a single pulse. The remaining rows give the errors on the π-pulses used to map between states in $S_{1/2}$ and $D_{5/2}$. The first state is the state in $S_{1/2}$ and the second state is the state in $D_{5/2}$. The measurements of the π-pulse errors are performed by first preparing the metastable level state (using post-selection) and then performing a mapping pulse from that state to $S_{1/2}$ and measuring the probability of a dark outcome.

The fraction of shots rejected after R_1 is given by the optical pumping error, leaving population within $S_{1/2}$ but outside of $|S_{1/2}, F=2, m_F=0\rangle$, and the error of the 1762 nm pulse used to transfer the population from the $|S_{1/2}, F=2, m_F=0\rangle$ state to the required state in $D_{5/2}$.

If we prepare a metastable level qubit or the metastable level state of an optical qubit, the fraction of shots rejected in R_2 is simply given by the probability of decay of the $D_{5/2}$ level during population detection. If we prepare a ground level qubit or the ground level state of an optical qubit, the fraction of shots rejected in R_2 is equal to the error of the 1762 nm π-pulse, ϵ_{1762}, used to map the qubit state to the metastable level.

The fraction of shots rejected as a result of R_3 and R_4 is equal to the error of the 1762 nm π-pulse, ϵ_{1762}, used to map the qubit state we have prepared from

Qubit type	State	Shots rejected (%)	Expected shots rejected (%)
O $\|0\rangle$	$\|D_{5/2}, F = 2, m_F = -1\rangle$	3.51(4)	3.56(15)
O $\|1\rangle$	$\|S_{1/2}, F = 2, m_F = 0\rangle$	10.98(6)	15.0(4)
M $\|0\rangle$	$\|D_{5/2}, F = 2, m_F = -1\rangle$	3.41(4)	3.56(15)
M $\|1\rangle$	$\|D_{5/2}, F = 1, m_F = -1\rangle$	8.74(6)	10.26(29)
G $\|0\rangle$	$\|S_{1/2}, F = 2, m_F = 0\rangle$	9.21(6)	9.76(27)
G $\|1\rangle$	$\|S_{1/2}, F = 1, m_F = 0\rangle$	5.72(5)	5.25(17)

Table 6.2: A comparison of the measured and expected total fraction of rejected shots during the SPAM protocol. We see a good agreement between the two. Any discrepancies are attributed to drifts in system performance between measurements.

the $D_{5/2}$ level back to the $S_{1/2}$ level before population detection.

Following the reasoning above and the errors provided in Table 6.1, in Table 6.2 we compare the expected and measured total fraction of rejected shots. We see that the measured values agree very well with our independent measurements of the fidelities of the individual operations in Table 6.1. Any differences are attributed to drifts in system performance between measurements, in particular, we have observed significant changes in the behaviour of the 1762 nm pulses in our setup.

6.5 Summary of SPAM results

The results from implementing the protocol on all three qubit types in $^{137}Ba^{+}$ are shown in Fig. 6.10. We see that we achieve the same errors for all qubit types within the error bars. This is expected: regardless of the qubit type being prepared, we expect the SPAM error to be purely determined by the population detection error. As we can see in Fig. 6.10, the SPAM errors agree with our estimate of the expected error obtained in Section 6.3.3. The only difference between the different types of qubits is the additional pulses required to prepare and measure the O and G qubits, introducing further errors at the different stages

of the protocol. However, these errors are all eliminated by post-selection, and hence, they result merely in a reduction in the data rate. This protocol has allowed us to demonstrate the lowest ever SPAM errors of any qubit platform for all three qubit types in $^{137}Ba^{+}$ – O: $7(4)\times10^{-6}$, M: $5(4)\times10^{-6}$, and G: $8(4)\times10^{-6}$.

This demonstrates that our protocol is agnostic to the qubit encoding and only depends on our ability to detect population with high fidelity. Furthermore, the individual steps in the protocol require capabilities that are usually readily available in quantum computing setups: standard techniques for qubit state preparation together with a set of population transfer pulses between two level manifolds of the quantum system. It can, therefore, be implemented in any existing setup with no experimental hardware overhead. Finally, the only requirement on the quantum system for implementing the protocol is the presence of two distinct level manifolds and the ability to detect population in one of these manifolds with a very low error. Thus, it applies to many quantum computing platforms, such as trapped ions, neutral atoms, or NV centres.

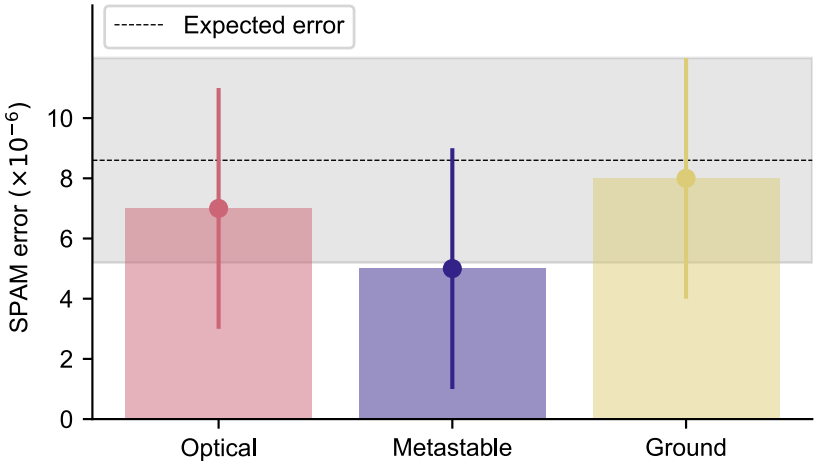


Figure 6.10: Measured SPAM errors for all three qubit types in $^{137}Ba^{+}$ — O: $7(4)\times10^{-6}$, M: $5(4)\times10^{-6}$, and G: $8(4)\times10^{-6}$. The dashed line and the shaded region represent the expected error and its uncertainty based on the independent measurements of the optical detection error and metastable level decay probability in Section 6.3.3. The errors for all three qubit types agree within the error bars and also agree with the expected error from Section 6.3.3.

7

Exploring a metastable level clock qubit

Contents

7.1 Choice of metastable level clock qubit 112
7.2 Experimental characterisation of the metastable level clock qubit 114
7.2.1 Field-insensitive point and qubit frequency 114
7.2.2 Qubit coherence time 116
7.3 Isolation of the metastable level qubit from dissipative light 122
7.4 Summary 125

In this chapter, we explore a metastable level clock qubit in $^{137}Ba^{+}$. Clock qubits are first-order insensitive to fluctuations in the external magnetic fields, one of the most common sources of error in quantum computing setups, which is generally hard to reduce. As a result, clock qubits exhibit longer coherence times. Clock qubits are thus widely used in trapped-ion and neutral atom systems, both for the purposes of precision measurements and quantum computing [1, 116].

Metastable clock qubits are not commonly used due to the complex structure of the metastable level in atomic systems that exhibit hyperfine splitting. Additionally, in most cases the system coherence time is limited by the natural lifetime of the metastable level, T_1, resulting in maximum obtainable net dephasing time of $2T_1$ [38]. However, the SPAM protocol we described in Chapter 6

detects and corrects for cases where a decay occurs, therefore enabling achieving metastable clock qubit coherence times that extend beyond the metastable level lifetime, as we will see.

7.1 Choice of metastable level clock qubit

The metastable $D_{5/2}$ level of $^{137}Ba^{+}$ has 24 states and, as a result, there are numerous possibilities for the qubit encoding, and also many pairs of states that form clock qubits at certain magnetic fields. A full list of the metastable level clock qubits obtained using numerical simulations [43] in $^{137}Ba^{+}$ is provided in Appendix A. When choosing which clock qubit to work with, there are several desirable properties we aim to satisfy:

- The qubit transition: we must make sure that the qubit transition is allowed by atomic physics selection rules.
- The required magnetic field: it is preferable to work with a low quantisation field. Higher magnetic fields increase the span of transition frequencies, resulting in more difficult fundamental ion operations such as cooling, state preparation, and readout. In addition, lower magnetic fields (few G) are much easier to generate experimentally.
- The qubit frequency: lower qubit frequencies (RF versus microwave) require lower-frequency control fields that are typically easier to generate.
- Spectator transitions: these are transitions from either one of the qubit states to another state outside of the qubit manifold. The further detuned these transitions are from the qubit transition, the lower the probability of driving population outside of the qubit manifold.
- Field sensitivity: while clock qubits are insensitive to the applied magnetic field to first order, they still have a quadratic dependence on changes in the magnetic field. It is, therefore, preferable to minimise the second-order magnetic field sensitivity to maximise the qubit coherence time.

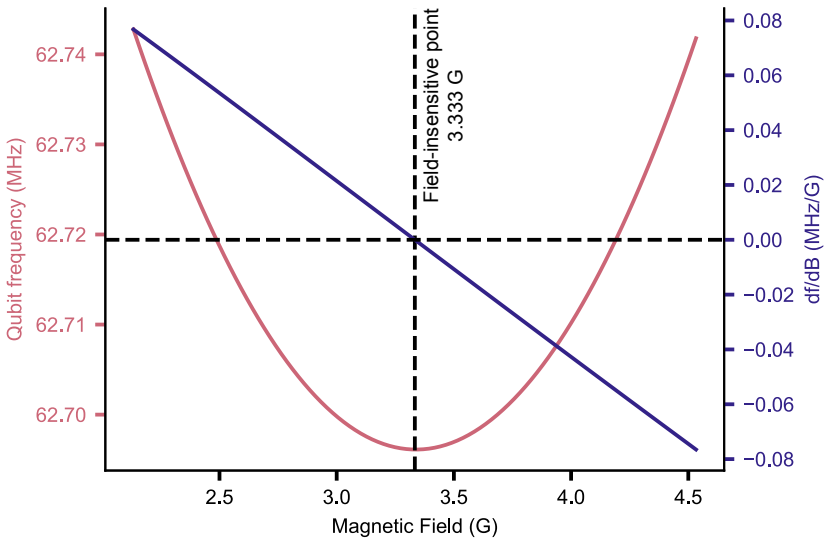


Figure 7.1: Evolution of the frequency splitting between the states $|F = 2, m_F = +2\rangle$ and $|F = 3, m_F = +2\rangle$ in the $D_{5/2}$ level of $^{137}Ba^+$. The field-insensitive point, at which $df/dB = 0$ is at $B_c \approx 3.333\,\mathrm{G}$. The corresponding clock qubit frequency is $f_q \approx 62.696$ MHz, and the second-order magnetic field sensitivity is $d^2f/dB^2 \approx 64.312\,\mathrm{kHz/G^2}$.

In Appendix A, we only list clock qubits with transitions allowed by atomic physics selection rules. Of these qubits, the clock qubit at the lowest magnetic field is formed by the states $|F = 2, m_F = 0\rangle \leftrightarrow |F = 3, m_F = -1\rangle$[1] at a quantisation field of $B_c \approx 1.693\,\mathrm{G}$ and a qubit frequency of $f_q \approx 63.054$ MHz. However, the nearest spectator transition to this qubit is detuned by only ≈ 0.2 MHz and is of the same polarisation as the qubit itself. This limits the speeds of both single- and two-qubit gates we can implement.

The second lowest-field metastable clock qubit is the $|F = 2, m_F = +2\rangle \leftrightarrow |F = 3, m_F = +2\rangle$ qubit at $B_c \approx 3.333\,\mathrm{G}$ and a qubit frequency of $f_q \approx 62.696$ MHz. The nearest spectator transition to this qubit is detuned by ≈ 2.4 MHz and has σ polarisation, contrary to the qubit transition, which has π polarisation. Hence, off-resonant effects can be further suppressed through the choice of the laser and/or microwave polarisations used for driving the qubit transition. The frequency

[1]We have labelled the states with the quantum numbers that they would have at a zero magnetic field for clarity. It should, however, be noted that due to the very small hyperfine splitting between the $F = 3$ and $F = 4$ manifolds (see Fig. 3.7), at this value of B_0 we are already in the intermediate-field regime and hence F is no longer a good quantum number.

and first-order sensitivity of this qubit as a function of the applied magnetic field are shown in Fig. 7.1[2].

7.2 Experimental characterisation of the metastable level clock qubit

7.2.1 Field-insensitive point and qubit frequency

We wish to measure the frequency of this qubit as a function of the applied magnetic field to verify that it matches the numerical simulations from Fig. 7.1 and to experimentally determine the field-insensitive point, at which $df/dB = 0$.

Magnetic field generation

The magnetic field in our experiment is generated using a current-carrying coil placed near the trap chamber, hereby referred to as the 'main' coil. We can change the magnetic field at the ion position by changing the amount of current applied to this coil. For the state preparation scheme described in Section 3.3.3, we must ensure that the magnetic field we generate is perpendicular to the 493 nm beam in Fig. 4.11. Any misalignment between the field produced by the main coil and the beam and any stray magnetic fields in the lab would degrade the optical pumping quality. To mitigate this effect, we installed a permanent magnet and an additional coil, placed such that the fields they generate at the ion positions are perpendicular to the trap table and parallel to the 493 nm beams, respectively (see Fig. 4.11). The position of the permanent magnet and the current applied to this additional coil are then tuned to maximise the optical pumping fidelity. This results in a very well-defined direction of the magnetic field in our experiment, which remains unchanged to first order for small changes in the current applied to the main coil.

We measured the applied magnetic field by measuring the transition frequencies of the quadrupole transitions from the $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ state to the states $|\mathrm{D}_{5/2}, F = 1, m_F = -1\rangle$, $|\mathrm{D}_{5/2}, F = 1, m_F = +1\rangle$, $|\mathrm{D}_{5/2}, F = 2, m_F = -1\rangle$, and

[2]This figure is the same as Fig. 3.9 and it has been reproduced here for completeness.

$|\mathrm{D}_{5/2}, F = 2, m_F = +1\rangle$, and using numerical simulations to fit the resulting magnetic field [43]. We used several transitions to ensure the results were self-consistent and thus verify the numerical simulations. The magnetic field as a function of the coil current, as measured by the ion, is shown in Fig. 7.2. As we would expect for a current-carrying coil, we observe a linear dependence of the magnetic field as a function of the applied current, with a gradient of 0.87418(25) G/A.

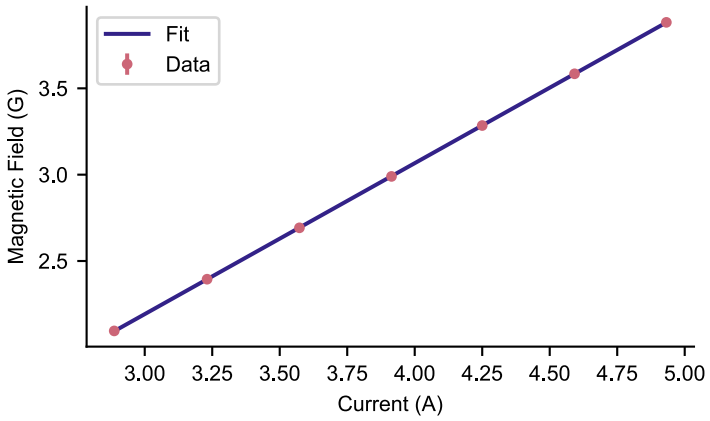


Figure 7.2: Magnetic field at the ion position as a function of the current applied to the coil. A linear fit to the data results in a gradient of 0.87418(25) G/A. The fitted magnetic field offset is −0.4302(10) G. We expect that this offset is due to stray magnetic fields in the lab in the direction of the field produced by the 'main' coil. The errorbars are smaller than the marker size.

Qubit frequency measurement

To measure the qubit frequency, we used an an RF tone[3] applied to an antenna[4] placed near the trap chamber to drive the qubit transition directly.

At the start of each experiment, the ion was prepared in the $|F = 2, m_F = 0\rangle$ state in the $\mathrm{S}_{1/2}$ level (see Section 3.3.3). The geometry of the 1762 nm beam in our setup suppresses quadrupole transitions with $\Delta m = \pm 2$, so we cannot transfer population directly from the $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ state to either of the metastable

[3]The RF tone is generated using a Sinara Urukul board.

[4]Here an antenna is a fancy name for a hand-made coil out of copper wire glued to a post near the trap chamber. This antenna was installed only for the purposes of this chapter, in particular, for the coherence time measurement in Section 7.2.2. It is only used for diagnostics and is not a permanent part of the setup.

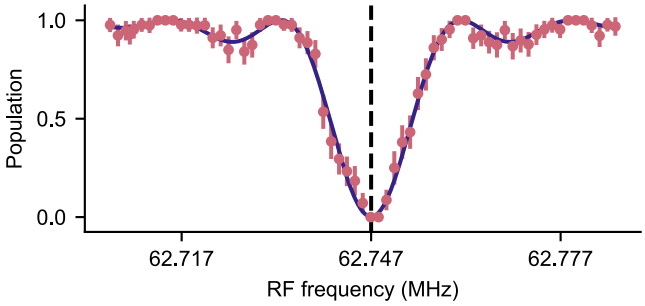


Figure 7.3: A detuning scan of the qubit transition obtained by applying a π-pulse to the qubit at varying RF frequencies. This scan was taken at a magnetic field of 2.0951(5) G. The qubit frequency is obtained by fitting a sinc^2 function to the data. The measured qubit frequency in this scan is 62.74703(11) MHz.

level qubit states. Hence, after state preparation, the population was transferred from the $|S_{1/2}, F = 2, m_F = 0\rangle$ state to the $|S_{1/2}, F = 1, m_F = +1\rangle$ state with the 532 nm Raman laser. We used the ground and the co-propagating Raman beam configuration shown in Fig. 4.11 to ensure good population transfer irrespective of the ion temperature. We then used a 1762 nm π-pulse to transfer the population from the $|S_{1/2}, F = 1, m_F = +1\rangle$ state to the $|D_{5/2}, F = 3, m_F = +2\rangle$ state. Once the qubit state was prepared, we probed the qubit transition using a Rabi π-pulse, scanning the frequency of the RF tone. The qubit frequency was extracted using a resonance curve fit to the detuning scan. An example of this measurement is shown in Fig. 7.3.

We performed the measurement described above at all coil current values in Fig. 7.2. The results are presented in Fig. 7.4 and are compared to the numerical simulation. The measured values of the field-insensitive point and the qubit frequency were obtained by fitting a quadratic function to the data. The measured field-insensitive point is 3.33(5) G and the corresponding qubit frequency is 62.6970(16) MHz. These results agree with our numerical simulations in Section 7.1 within the errorbars.

7.2.2 Qubit coherence time

The transition frequencies of clock qubits are first-order insensitive to variations in the applied magnetic field. As a result, they are much more robust to magnetic

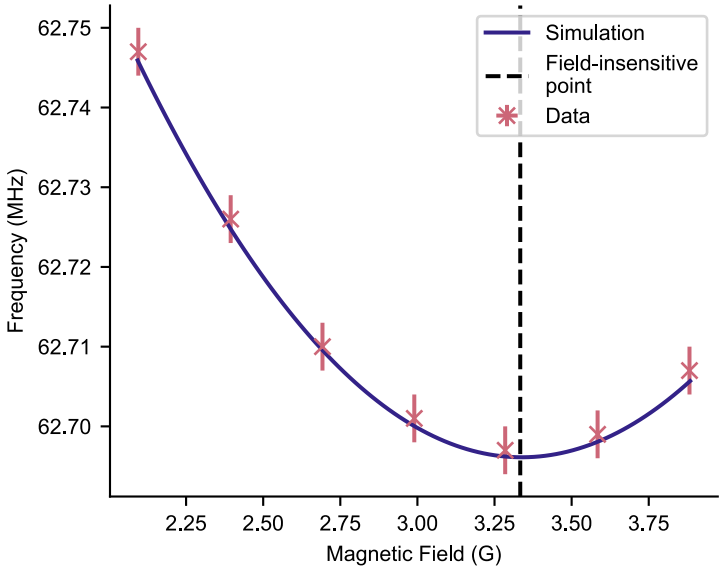


Figure 7.4: Metastable level qubit frequency as a function of the quantisation field. The blue line is obtained using numerical simulations of the level structure [43]. Using a quadratic fit to the data we infer a measured field-insensitive point of 3.33(5) G and a corresponding qubit frequency of 62.6970(16) MHz. These values agree with the simulated values within the errorbars.

field noise compared to Zeeman qubits, one of the main sources of error in trapped-ion systems, and hence exhibit much longer coherence times. In this section, we explore the coherence time of the metastable level clock qubit we used in Section 7.2.1.

The most standard method for measuring qubit coherence time is the Ramsey sequence, as shown in Fig. 7.5(a). The qubit is first prepared in the state $(|0\rangle + |1\rangle)/2$ and then left idle for some time τ. Any variation in the external magnetic field ΔB will cause the qubit frequency to change by $\Delta\omega$ and hence result in an acquired phase $\Delta\phi = \tau\Delta\omega$. In the case of a clock qubit, $\Delta\omega \propto \Delta B^2$. Finally, another $\pi_\phi/2$ pulse with some phase ϕ is applied to the qubit, and its state is measured. In the absence of magnetic field noise, such that $\Delta\phi = 0$, scanning the phase of the final $\pi/2$ pulse would result in fringes with a period of 2π and contrast 1, as shown in Fig. 7.5(b). Noise in the external magnetic field, resulting in a random variation of the qubit frequency and hence a random acquired phase $\Delta\phi$ would cause the contrast of the Ramsey fringes to drop. The qubit coherence time is then the time T_2 for which the Ramsey fringe contrast is reduced to $1/e$.

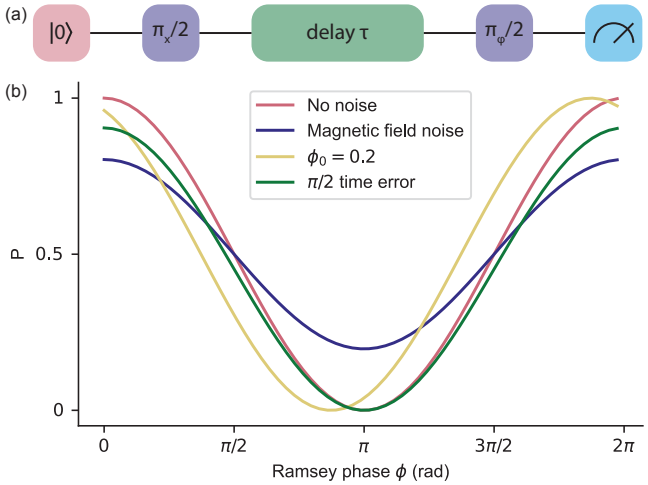


Figure 7.5: (a) A simple Ramsey sequence used for measuring the qubit coherence time. To obtain the Ramsey contrast we perform two measurements with analysis phases of $\phi = 0.0$ and $\phi = \pi$, and the contrast is calculated by taking the difference between the two points. (b) Simulation of the observed Ramsey fringes for $\tau = 1\,\mathrm{s}$ in the following scenarios: a perfectly calibrated experiment in the absence of noise (red), magnetic field noise resulting in $T_2 = 1\,\mathrm{s}$ (blue), coherent qubit frequency offset of 0.2 Hz (yellow), and $\pi/2$ pulse duration error of 20% (green). During a two-point measurement, we only probe at $\phi = 0$ and $\phi = \pi$, therefore, all of these effects would result in a reduction of the measured contrast.

The behaviour of the Ramsey fringes in a real experiment is affected by several factors:

- External magnetic field noise resulting in a random acquired phase during the Ramsey delay. This results in a reduction of the fringe contrast.
- Mis-calibrated qubit frequency. This would result in an offset of the Ramsey fringes. However, this is a coherent effect, which would not result in contrast reduction.
- Mis-calibrated $\pi/2$ pulse duration. This would reduce the contrast of the fringes even in the absence of any other sources of noise.
- Decay of the metastable level resulting in qubit erasure.

Simulations of these effects[5] are provided in Fig. 7.5(b).

[5]Simulation performed using Qiskit and Qiskit-Aer.

The qubit frequency and π time cannot be adjusted during a computation. Therefore, to characterise the behaviour of the metastable level clock qubit in an equivalent scenario, we used a two-point measurement to measure the contrast of the fringes, as in Ref. [2, 7]. That is, we performed the experiment in Fig. 7.5(a) with $\phi = 0$ and $\phi = \pi$. The contrast is then the difference in the populations between the two measurements. When measured in this way, we are sensitive to loss of contrast both due to qubit decoherence and drift of the qubit frequency.

The goal of this measurement was to characterise the performance of the metastable level qubit as a clock qubit, and not to characterise the metastable level lifetime of $^{137}Ba^{+}$. For a metastable level lifetime of T_1, the maximum net qubit coherence time T_2^* that we can measure, in the limit of no external noise, is $T_2^* = 2T_1$ [38]. To enable measurement beyond this limit, we embedded the Ramsey experiment shown in Fig. 7.5(a) in the SPAM protocol described in Chapter 6. This protocol detects and removes shots where there was a qubit decay before the quantum measurement is performed. Hence, we were able to measure the metastable qubit coherence time, i.e. its T_2 time [38], independently of decay to the ground level.

Experiment calibration

We use a two-point measurement to measure the Ramsey fringe contrast. As a result, in addition to qubit decoherence, we are also sensitive to mis-calibration of the transition frequency and $\pi/2$ duration.

The $\pi/2$ pulse duration is calibrated as follows. We use Ramsey sequence in Fig. 7.5(a) with $\tau = 0$ and scan the $\pi/2$ pulse duration, which is equivalent to performing a π-pulse on the qubit. The optimal duration is achieved when we observe full population transfer of the qubit from the state $|0\rangle$ to the state $|1\rangle$.

The qubit frequency is calibrated using the RPE protocol described in Section 5.3.1. The accuracy of the qubit frequency calibration and the frequency resolution of the RF source can both result in a lower contrast measurement.

Suppose that we acquire some frequency offset $\Delta\phi$ during the Ramsey experiment. Then, the two-point measurement will result in a contrast of $C = \sin(\Delta\phi + \pi/2) - \sin(\Delta\phi)$. The contrast is $C = 1/e$ when $\Delta\phi \approx 0.5224\,\text{rad}$. Therefore, for a given frequency offset $\Delta\omega$ the maximum coherence time we can measure is given by $\tau_{\text{max}} = 0.5224/\Delta\omega$.

The accuracy of the RPE frequency calibration is limited by the qubit coherence time itself, and does not pose a limit to the measurement range of the two-point measurement.

Another source of qubit frequency error is the finite resolution of the RF sources we control the qubit with. The DDS chip[6] we use has a frequency resolution of 0.23 Hz. Therefore, even with perfect calibration, the maximum coherence time that can be obtained using a two-point measurement is ≈ 0.361 s. To increase the measurement range further, we first use RPE to calibrate the centre frequency of the transition and hence set the DDS frequency. We then run RPE again to calibrate the remaining frequency offset (that is below the DDS frequency resolution). To compensate for this offset, we add an additional phase offset to the second $\pi/2$ pulse in the Ramsey sequence corresponding to the phase accumulated due to the frequency offset during the Ramsey delay time. For a given frequency offset of $\Delta\omega$ and a Ramsey delay time of τ, the acquired phase is given by $\tau\Delta\omega$. Therefore, to probe at a nominal analysis phase of 0, we instead probe at a phase $\tau\Delta\omega$, and to probe at a nominal analysis phase of π, we instead probe at a phase of $\pi + \tau\Delta\omega$.

Experimental results

We measured the qubit coherence time relative to both an RF source, driving the qubit using an antenna close to the trap, as we did in Section 7.2.1, and relative to a Raman laser beam beatnote. For this measurement we used the Raman beams in the metastable co-propagating configuration shown in Fig. 4.9 and

[6]We use a Sinara Urukul board as an RF source. The DDS chip used by this board is AD9910. This board was used to drive both the antenna we used in the measurement in Section 7.2.1 and the AOMs in the Raman beam paths shown in Fig. 4.9.

Fig. 4.11[7]. Microwave sources generally have much lower phase noise than laser sources. Hence, we measure the coherence time using the antenna as a benchmark for both the qubit performance and the performance of our Raman beam setup.

The results from these measurements are shown in Fig. 7.6. The measured qubit coherence time using the RF antenna is 31(5) s, which is long enough to not dominate the device performance. This coherence time is longer than the natural lifetime of the metastable level of $^{137}Ba^{+}$, 30.14 s. The ability to achieve a coherence time longer than the metastable level lifetime is important: while we can detect and correct for the decay of the metastable level, as we do in the SPAM protocol in Chapter 6, we cannot trivially detect and correct for qubit dephasing.

Let's estimate what the measured coherence time means in terms of magnetic field noise in the lab, assuming this is the only noise source limiting the measured coherence time. We can write the qubit frequency as a function of the magnetic field near the field-insensitive point as $f = f_{\text{clock}} + \frac{1}{2}\alpha\,(B - B_0)^2$, where f_{clock} is the frequency of the clock qubit, B_0 is the field-insensitive point, and α is the second-order magnetic field sensitivity. For the clock qubit used here, $\alpha \approx 64.312\,\text{kHz/G}^2$. For a given change in the magnetic field ΔB, the change in frequency, to second order, is thus given by $\Delta f \approx \frac{1}{2}\alpha\,(\Delta B)^2$. The measured coherence time of $\tau = 31(5)$ s corresponds to a frequency variation of $\sim 1/\tau \approx 0.032(5)\,\text{Hz}$. Thus, an upper bound on the magnetic field fluctuations in the lab is 1.00(8) mG, which is of the expected order of magnitude.

The measured qubit coherence time using the Raman lasers is 9.1(10) s. As expected, this is lower than the coherence time measured using a microwave source. We expect the coherence time here to be limited by phase noise between the two Raman beams. However, this is a reasonable starting point, considering we have not taken much care in our setup to minimise noise in the Raman beams. This could be improved using standard techniques for reducing laser phase noise [117, 118].

[7]For these measurements, we use global Raman beams with waist radii of ~ 5 µm rather than the individual addressing setup presented in Chapter 5.

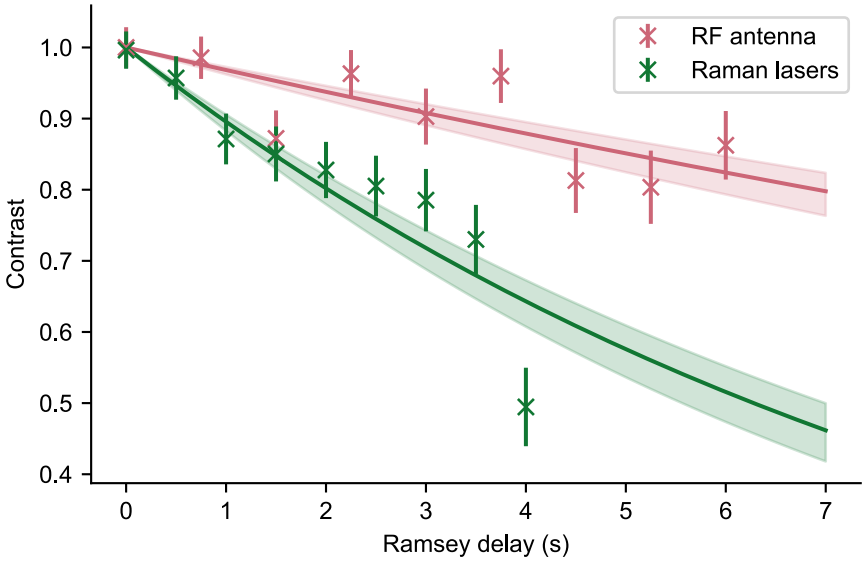


Figure 7.6: Coherence time measurement of the metastable level qubit using the Ramsey experiment shown in Fig. 7.5. The red curve shows the decay in contrast of the metastable level qubit measured by directly driving the qubit transition using an RF antenna placed near the trap. We use a decay model of the contrast as a function of Ramsey delay time of $C = e^{-\tau/T_2}$. The fit corresponds to a coherence time of $T_2 = 31(5)$ s. The green curve shows the decay in contrast of the metastable level qubit measured using the Raman lasers. The fit corresponds to a coherence time of $T_2 = 9.1(10)$ s. The shaded regions correspond to a 1σ uncertainty interval.

7.3 Isolation of the metastable level qubit from dissipative light

The main advantage of using the metastable level for both gate operations and memory is that we can implement mid-circuit dissipative operations such as cooling, measurement, and qubit reset in the ground level, while preserving the coherence of the metastable level qubits. In this section, we quantify this statement by measuring the qubit coherence time in the presence of dissipative light. While these lasers do not drive transitions within the metastable level, they can cause a differential light shift between the two qubit states, resulting in a qubit phase shift. This phase shift can be compensated for by advancing the DDS phase of the following operation, as we did above to compensate for the finite frequency resolution of the DDS. Therefore, as long as this phase shift is stable, it can be calibrated and compensated out. However, any intensity noise in the applied laser beams will turn into noise on the generated phase shift, leading to decoherence.

To evaluate the coherence of the metastable level qubit during dissipative operations, we perform the same measurement as in Section 7.2.2, except this time, instead of a simple delay, we turn on the 493 nm and the 650 nm beams, used for cooling and measurement, for the duration τ. These laser beams are used at several detunings and powers during different operations. For this measurement, we set them to the parameters used during qubit readout, as it typically requires the highest beam intensities to maximise the ion fluorescence. The parameters we used for this measurement were a 493 nm beam power of 150 µW and a 650 nm beam power of 1 mW. The measurement was performed using the RF antenna to probe the qubit directly, as it resulted in a longer qubit coherence time compared to probing with the Raman beams, as shown in Fig. 7.6.

We start by measuring the light shift produced by the two beams on the metastable level qubit using RPE as in Chapter 5. The resulting light shift is $0.871^{+0.005}_{-0.004}$ Hz[8]. As before, this is compensated for by advancing the phase of the DDS for the second $\pi/2$ pulse. We then use the two-point measurement described in Section 7.2.2 to measure the qubit coherence time during dissipative operations.

The results from this measurement are shown in Fig. 7.7. The measured coherence time in the presence of the dissipative beams is 21.9(21) s. We see that while there is some reduction in the qubit coherence time, this effect is small enough to not limit the operation of the device. The measured coherence time corresponds to a variation of the light shift experienced by the qubit of ≈ 0.05 Hz, or equivalently, $\approx 6\%$. This is within the accuracy we would expect from our beam power stabilisation system.

Discussion

The coherence time in the presence of dissipative operations of 21.9(21) s that we measured is a reasonable starting point, and we do not expect it to dominate the device performance in the near term. However, as the number of qubits and the lengths of the circuits increase, we may wish to increase this coherence time further

[8]The confidence intervals here are obtained by bootstrapping the RPE data, hence the asymmetry in the error bounds.

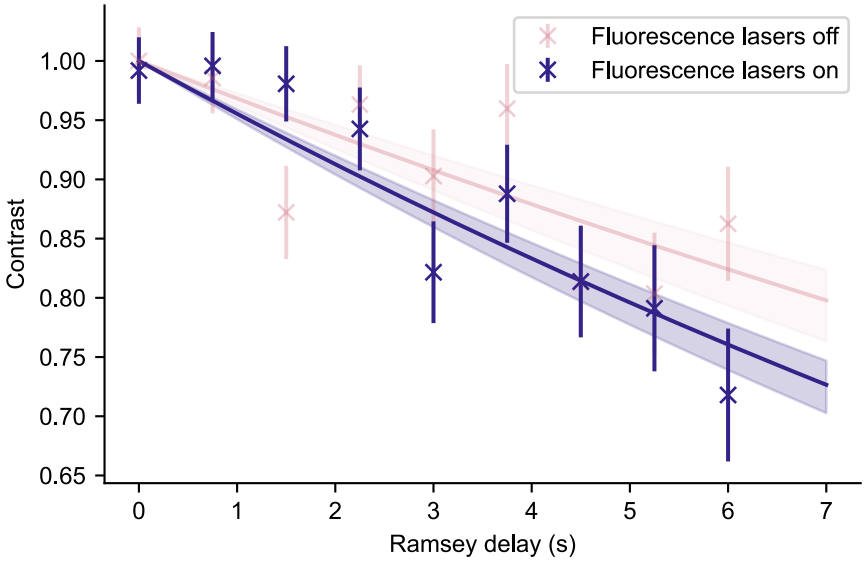


Figure 7.7: Coherence time measurement of the metastable level qubit in the presence of the 493 nm and the 650 nm beams set to the parameters used during readout (blue data). The coherence time we measure is $T_2 = 21.9(21)$ s. The pink curve shows the measured qubit coherence time (31(5) s) in the absence of the dissipative beams – this is the same measurement as in Fig. 7.6 and is plotted here again for easier comparison. We see a small reduction in the coherence time due to the dissipative beams. We expect that this is due to intensity noise in these beams. The measured coherence time corresponds to an intensity noise level of $\approx 6\%$. The shaded regions represent a 1σ uncertainty interval.

beyond the lifetime of the metastable level. The decay of the metastable level can be detected and corrected as is done in the SPAM protocol in Chapter 6. On the other hand, qubit dephasing, regardless of its source, cannot be easily detected.

There are two approaches to mitigating the effect of light shift noise: engineering and physics. The engineering approach involves improving the power stabilisation system to reduce the intensity noise of the beams. However, the performance of the power stabilisation system will always limit the coherence time of the qubit.

An alternative approach is to set the beam polarisations, such that there is no net light shift of the qubit transition. The polarisation of the 493 nm beam in this experiment is constrained by the scheme used for state preparation (see Section 3.3.3). However, in the case of the metastable level qubit, we expect the majority of the light shift to be produced by the 650 nm beam: this beam is both closer in detuning to the $D_{5/2} \leftrightarrow P_{3/2}$ transition and has a factor of 10 more

intensity compared to the 493 nm beam. The only constraint for the polarisation of the 650 nm beam is the ability to drive population out of all states within the $D_{3/2}$ level. In other words, it must have non-zero σ_+ and σ_- components. Numerical simulations of the light shift experienced by the metastable level clock qubit show that the light shift is cancelled out when the 650 nm beam has equal components in each of the σ_+, σ_- and π polarisations. This polarisation of the 650 nm beam is indeed able to repump population out of all states within the $D_{3/2}$ level and could be implemented for future experiments.

7.4 Summary

The results in this section demonstrate that we can implement clock qubits in the metastable level of $^{137}Ba^+$ and achieve a qubit coherence time that is longer than the natural lifetime of the metastable level and, therefore, does not dominate the system's performance.

We measured a qubit coherence time of 31(5) s. This is a good enough coherence time to not limit coherent operations significantly. The coherence time could be improved with improvements to the experimental setup, such as using magnetic field shielding to further reduce the effects of magnetic field noise or improving the isolation of the qubit drive. Such techniques are discussed in detail in Ref. [1].

One of the main advantages of working with metastable level qubits is the implementation of mid-circuit operations, such as cooling and measurement. We evaluated the isolation of the metastable level qubit in the presence of the laser beams necessary for these operations. We measured a coherence time of 21.9(21) s, consistent with $\approx 6\%$ intensity noise in these beams. However, our numerical calculations showed that it is possible to remove the majority of this effect, irrespective of the laser intensity noise, by adjusting the polarisation of the 650 nm beam.

The results above show that the metastable level of $^{137}Ba^+$ is suitable for encoding qubits for quantum information processing and storage. This encoding

also enables the implementation of mid-circuit dissipative operations with minimal effects on the computational (i.e. metastable level) qubit coherence.

I really miss the days when the weirdest thing science ever created was me.

— Captain America, *Avengers: Age of Ultron*

8 Two-qubit gates

Contents

8.1 Mølmer-Sørenson gate overview 129
8.2 Raman beam setup 131
8.2.1 Raman beam configuration 131
8.2.2 Raman beam power 132
8.2.3 Raman beam polarisation 133
8.3 Motional mode structure 138
8.3.1 Choice of gate mode 139
8.3.2 Motional heating rates 139
8.3.3 Lamb-Dicke parameter 140
8.4 Gate implementation 140
8.4.1 Qubit encoding 140
8.4.2 General gate sequence 141
8.5 Gate calibration 141
8.5.1 Equal ion illumination 141
8.5.2 Equal sideband strength 142
8.5.3 Calibration of the gate detuning and duration 142
8.6 G-G gates 144
8.7 M-M gates 144
8.8 G-M gates 146
8.8.1 Experiment sequence 148
8.8.2 Gate calibration 148
8.8.3 G-M gate results 150
8.9 Discussion 150
8.10 Summary 155

In this chapter, we describe the experimental implementation of two-qubit entangling gates between both ground and metastable level qubits in $^{137}Ba^{+}$. Two-qubit entangling gates are required to build a complete set of unitary operations necessary for implementing arbitrary quantum algorithms [119].

In the architecture proposed in Chapter 2, we implement all quantum information processing in the metastable level and as such, we only require gates between metastable level qubits (M-M gates).

In addition to implementing a complete set, two-qubit gates are required to implement quantum non-demolition (QND) measurements. QNDs are crucial for many quantum algorithms and quantum error correction [24, 25, 41, 42]. While QNDs collapse the state of the qubit, they do not destroy it, and thus, the qubit can be used in subsequent operations.

Typically, QNDs are implemented by using two qubits, a computational qubit and an ancilla qubit. The states of the two qubits are first entangled. Then, the state of the ancilla qubit must be measured in a way that does not destroy the computational qubit [41, 42].

In an architecture that uses multiple qubit encodings, QND measurements can be implemented in two ways. One way is to use a computational qubit and an ancilla qubit of the same type; for example, encode them both in the ground level for the entangling gate. Then, the computational qubit must be mapped to the metastable level before the ancilla measurement, such that it is not destroyed by the measurement. While this approach works in principle, it requires coherent mapping of the qubit state between the ground and the metastable level, imposing strict requirements on the coherence of the laser used to drive the transitions between the two levels.

An alternative approach is to perform a direct entangling gate between a metastable level (i.e. computational) qubit and a ground level (i.e. ancilla) qubit (a G-M gate), the equivalent of the mixed-species gate [7]. The ability to perform this type of gate removes the need for any coherent mapping pulses between the ground and the metastable level. Thus, these pulses are only necessary for

population transfer for state preparation and measurement, neither of which require good coherence of the laser. Hence, performing G-M gates makes our architecture completely independent of the coherence of our quadrupole laser (or any method used to transfer population between the $S_{1/2}$ and the $D_{5/2}$ levels).

8.1 Mølmer-Sørenson gate overview

Implementing entangling gates in trapped-ion systems requires applying a spin-dependent force on the ions to excite their collective motion, which is coupled via their Coulomb repulsion, except for Rydberg ions where the dipole-dipole interactions of ions in highly excited states is used instead. Two-qubit gates based on applying spin-dependent forces vary in the way and the basis in which the force is implemented. For example, light-shift gates drive a force in the z-basis [31], and Mølmer-Sørenson gates drive a force in the x-basis [38]. For this work, we chose to work with Mølmer-Sørenson (MS) gates since they can easily be directly applied to clock qubits [38].

The theory of MS gates has been described in detail elsewhere [38, 44]. In this section, we only give a brief overview of the general principle of the gate to introduce notation.

Performing an MS gate requires simultaneously driving both the red and blue motional sidebands of the two qubits. This is achieved by applying two Raman beatnotes symmetrically detuned by some amount δ from each of the two motional sidebands, the blue sideband tone (BSB) and the red sideband tone (RSB), as shown in Fig. 8.1.

For a pair of ions, in the interaction picture this generates the Hamiltonian:

$$\hat{H}_I(t) = \hbar\Omega_{\mathrm{MS}} \left(\hat{\sigma}_{x,1} \pm \hat{\sigma}_{x,2}\right) \left(\hat{a}^\dagger e^{i\delta t} + \hat{a} e^{-i\delta t}\right) \tag{8.1}$$

where Ω_{MS} is the gate Rabi frequency, δ is the detuning from the motional mode as shown in Fig. 8.1, t is the instantaneous time in the interaction, and $\hat{a}$ and $\hat{a}^\dagger$ are the annihilation and creation operators, respectively. The signs in the first pair of brackets correspond to the relative motion of the ions in the mode

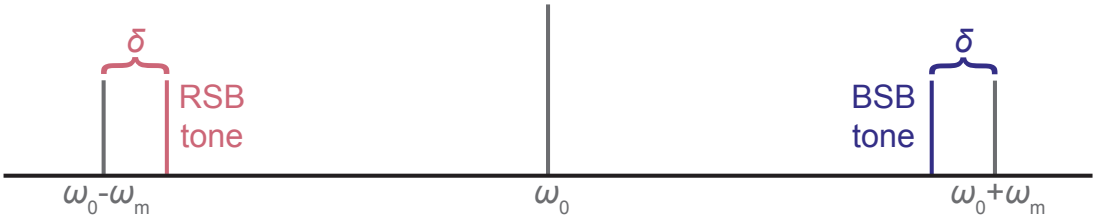


Figure 8.1: Blue (BSB) and red (RSB) sideband tones applied for driving an MS gate. The qubit frequency is ω_0. The motional mode frequency is ω_m. The two tones are detuned by δ from either of the motional modes, that is, we apply two beatnotes with frequencies $\omega_0 + \omega_\mathrm{m} - \delta$ and $\omega_0 - \omega_\mathrm{m} + \delta$.

we choose to couple to, where a positive sign corresponds to the centre-of-mass (COM) mode (in-phase motion), and a negative sign corresponds to the radial rocking mode (out-of-phase motion). We can then integrate the Hamiltonian to obtain the corresponding propagator:

$$\hat{U}(t) = \hat{D}\left(\alpha(t)\right) e^{-i\phi(t)} \tag{8.2}$$

where $\hat{D}$ is the displacement operator. The displacement $\alpha(t)$ is given by

$$\alpha(t) = \frac{\Omega_\mathrm{MS}}{\delta}\left(\hat{\sigma}_{x,1} \pm \hat{\sigma}_{x,2}\right)\left(e^{-i\delta t} - 1\right) \tag{8.3}$$

and the phase is given by

$$\phi(t) = \frac{\Omega_\mathrm{MS}^2}{\delta^2}\left(\hat{\sigma}_{x,1} \pm \hat{\sigma}_{x,2}\right)^2 \left(\delta t - \sin(\delta t)\right). \tag{8.4}$$

The propagator in Eqn. 8.2 corresponds to a spin-dependent displacement in phase space and a spin-dependent geometric phase. The gate Rabi frequency, detuning, and duration must be chosen such that the required geometric phase is acquired and that the spin and the motional states of the ions are decoupled again at the end of the gate. Any residual coupling between the spin and the motion of the ions after the gate leads to subsequent errors when the motion is traced out, e.g. to read out the spin.

Let's denote the total gate duration by τ_MS. To ensure that the spin and the motion are decoupled at the end of the gate, we require that $\alpha\left(\tau_\mathrm{MS}\right) = 0$. From Eqn. 8.3 this results in

$$\tau_\mathrm{MS}\delta = 2n\pi. \tag{8.5}$$

for some integer n. To create the maximally entangled state, we further require that the geometric phase in Eqn. 8.4 satisfies $\phi(\tau_{\mathrm{MS}}) = \pi/2$. Combined with Eqn. 8.5 this results in

$$\delta = \pm 4\Omega_{\mathrm{MS}}. \tag{8.6}$$

Under these conditions, a system that is initially in $|00\rangle$ ends in the maximally entangled state $(|00\rangle + i\,|11\rangle)/\sqrt{2}$ after the application of the MS gate.

8.2 Raman beam setup

Two-qubit gates in our experiment are performed using two-photon Raman transitions. Our Raman laser source is described in detail in Chapter 4. In this section, we describe the beam parameters used for implementing two-qubit gates.

8.2.1 Raman beam configuration

The available Raman beam configurations in our apparatus are shown in Fig. 4.11. Since implementing two-qubit gates requires coupling to the motion of the ions, here we work with the counter-propagating beam configuration.

The BSB and RSB tones are applied to the counter-propagating beam as shown in Fig. 8.2. To generate a gate drive on the ground level qubit, we also turn on the ground beam in Fig. 4.11, and to generate a gate drive on the metastable level qubit, we turn on the metastable beam instead. In the case of a G-M gate, both the ground and the metastable beams are on in addition to the counter-propagating beam.

One of the main sources of error in MS gates is off-resonant coupling to nearby carrier transitions [7]. A well-known way to mitigate this effect is by shaping the amplitude of the gate pulse. We implement pulse shaping using a custom circuit board[1], which generates a Gaussian shaping of the edges of the pulse with a bandwidth of 238 kHz [120]. This board is used to shape the

[1]This board was built by Peter Drmota.

counter-propagating beam pulse only, as shown in Fig. 8.2, resulting in shaping the Raman Rabi frequency experienced by the ion.

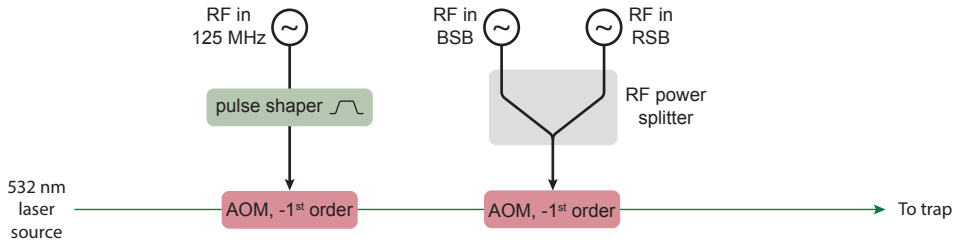


Figure 8.2: Beam path configuration of the counter-propagating beam used for implementing two-qubit gates (see also Section 4.3.4). We apply a static RF tone at 125 MHz to the first quartz AOM. The RF tone passes through a pulse shaper board to smooth the rising/falling edges of the pulse, effectively shaping the Rabi frequency experienced by the ions. This minimises off-resonant coupling to nearby transitions during the gate. The BSB and RSB tones are generated separately, combined using an RF power splitter and then applied to the second quartz AOM. The frequencies of the BSB and RSB tones are 125 MHz $\pm\, \omega_m \mp \delta$, respectively, where ω_m is the frequency of the mode we wish to couple to and δ is the MS gate detuning.

As in Chapter 7, the measurements here were taken with global beams with a waist radius of $\sim$ 5 µm that illuminate both ions at the same time, rather than the addressing system presented in Chapter 5.

8.2.2 Raman beam power

The powers of the Raman beams determine the gate Rabi frequency and, thus, the gate duration. Faster gates are desirable because they help to reduce any effects of the ions' interaction with the environment, such as motional heating or decoherence. However, a high gate Rabi frequency increases the off-resonant coupling to nearby transitions.

For this initial demonstration, we chose to maximise the available power, thus maximising the gate Rabi frequency, while staying comfortably within the power ratings of the components in each beam path:

- Ground beam power: 1 mW
- Metastable beam power: 5 mW
- Counter-propagating beam power: 300 mW

The counter-propagating beam has a much higher power due to the use of quartz AOMs in its beam path instead of TeO_2 AOMs, and the use of a photonic crystal fibre to deliver the light to the trap (see Chapter 4). These beam powers could be increased further in the future by either increasing the spot sizes along the laser beam paths, thus reducing the beam intensity, or by using components with a higher power tolerance (for example, replacing the TeO_2 AOMs).

8.2.3 Raman beam polarisation

Choosing the Raman beam polarisations for our experiment is a more complex task than for experiments involving a single qubit encoding. We need to carefully configure the polarisation to efficiently drive both the ground and metastable level qubits while minimising undesired effects like significant light shifts of the qubit transition or off-resonant excitation of spectator transitions.

Transitions within the $F = 2$ manifold in $S_{1/2}$

When we first set up the counter-propagating beam and attempted driving the ground level qubit transition ($|F = 2, m_F = 0\rangle \leftrightarrow |F = 1, m_F = 0\rangle$), we noticed that we were driving a significant amount of the population ($\approx 30\%$) outside of the qubit manifold.

To understand the source of this effect, we performed the following experiment. We initialised the ion in the state $|S_{1/2}, F = 2, m_F = 0\rangle$ and applied either the ground or the counter-propagating beam for some time t. We then transferred the population in the $|S_{1/2}, F = 2, m_F = 0\rangle$ state to the $|D_{5/2}, F = 2, m_F = -1\rangle$ state and measured the remaining population within $S_{1/2}$ and $D_{3/2}$. In the absence of scattering and off-resonant excitation, the result of this measurement would always be 0. This is indeed what we measured when turning on the ground beam only. However, when turning on the counter-propagating beam only, we saw the same loss of population of $\approx 30\%$ as in the case of both beams being on.

Driving population outside of the qubit manifold with only the counter-propagating beam could be either due to Raman scattering or some coherent

effect due to the high intensity of the beam. The Raman scattering probability in barium due to 532 nm light has been studied and experimentally measured in detail before [121, 122]. These results show that the expected Raman scattering probability for the ground level would be $\sim 10^{-6}$, which is far too small to explain the $\approx 30\%$ population loss we observed.

To understand this effect further, we individually measured the populations in each state within the ground level. We did this by initialising the ion in the state $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ and then turning on the counter-propagating beam for some duration t. We then transferred the population from all but one state within $\mathrm{S}_{1/2}$ to $\mathrm{D}_{5/2}$ and measured the remaining population within $\mathrm{S}_{1/2}$. We found that the population was only distributed within the $F = 2$ manifold within the $\mathrm{S}_{1/2}$ level. This would not be the case if the reason for population loss were Raman scattering, resulting in population within the $\mathrm{D}_{3/2}$ level and the $F = 1$ manifold in $\mathrm{S}_{1/2}$ as well. The results from this measurement for the states in $F = 2$ in $\mathrm{S}_{1/2}$ are shown in Fig. 8.3. We see that the single intense Raman beam is causing coherent population transfer between states in the $F = 2$ manifold of the $\mathrm{S}_{1/2}$ level.

To understand this effect, we model the ion as a five-level system acted on by a DC electric field, following the treatment in Chapter 5 of Ref. [123]. The atomic Hamiltonian is given by:

$$\hat{H}_{\text{ion}} = \begin{pmatrix} -2\omega_0 & 0 & 0 & 0 & 0 \\ 0 & -\omega_0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \omega_0 & 0 \\ 0 & 0 & 0 & 0 & 2\omega_0 \end{pmatrix} \tag{8.7}$$

where ω_0 is the frequency splitting between neighbouring states within $F = 2$ in $\mathrm{S}_{1/2}$ and we have taken the zero of energy to be the state $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$. The field Hamiltonian is given by

$$\hat{H}_{\text{field}} = \begin{pmatrix} 0 & \Lambda & 0 & 0 & 0 \\ \Lambda & 0 & 1.22\Lambda & 0 & 0 \\ 0 & 1.22\Lambda & 0 & 1.22\Lambda & 0 \\ 0 & 0 & 1.22\Lambda & 0 & \Lambda \\ 0 & 0 & 0 & \Lambda & 0 \end{pmatrix} \tag{8.8}$$

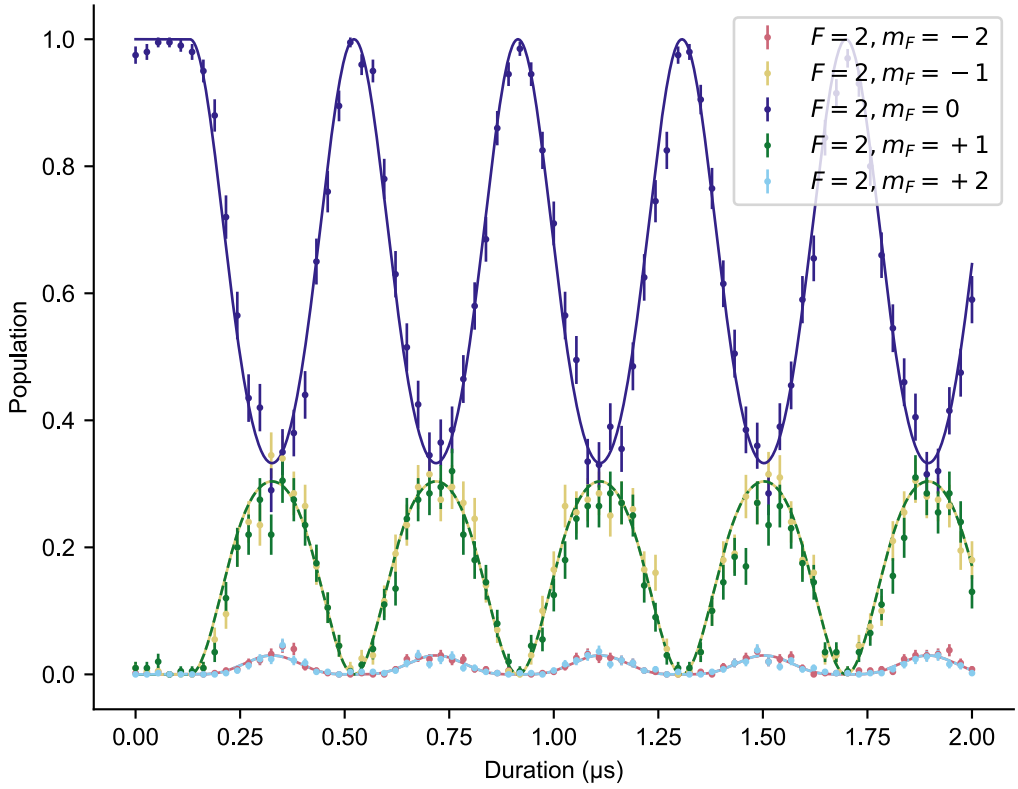


Figure 8.3: Population transfer within the $F = 2$ manifold of the $S_{1/2}$ level due to the intense counter-propagating Raman beam. The lines are numerical simulations of the 5-level system formed by the states in $F = 2$ in $S_{1/2}$. The parameters of the numerical simulation (frequency splitting and field Rabi frequency) are extracted by fitting the data for the $|F = 2, m_F = 0\rangle$ states, and are used to subsequently numerically simulate the evolution of the remaining states.

where Λ represents the transition coupling strength, and we have assumed that we can only drive transitions between states connected by dipole-allowed transitions, such that $|\Delta m_F| \leq 1$. The factor of 1.22 in some of the terms of this Hamiltonian corresponds to the ratio of the Raman transition matrix elements for transitions between $m_F = 0$ and $m_F = \pm 1$ states compared to transitions between $m_F = \pm 1$ and $m_F = \pm 2$ states, calculated numerically [43]. From symmetry, the sign of m_F does not change the interaction strengths. The total Hamiltonian is then the sum of the Hamiltonians in Eqn. 8.7 and Eqn. 8.8:

$$\hat{H}_{\text{tot}} = \hat{H}_{\text{ion}} + \hat{H}_{\text{field}}. \tag{8.9}$$

We can simulate the time evolution of the system under this Hamilonian

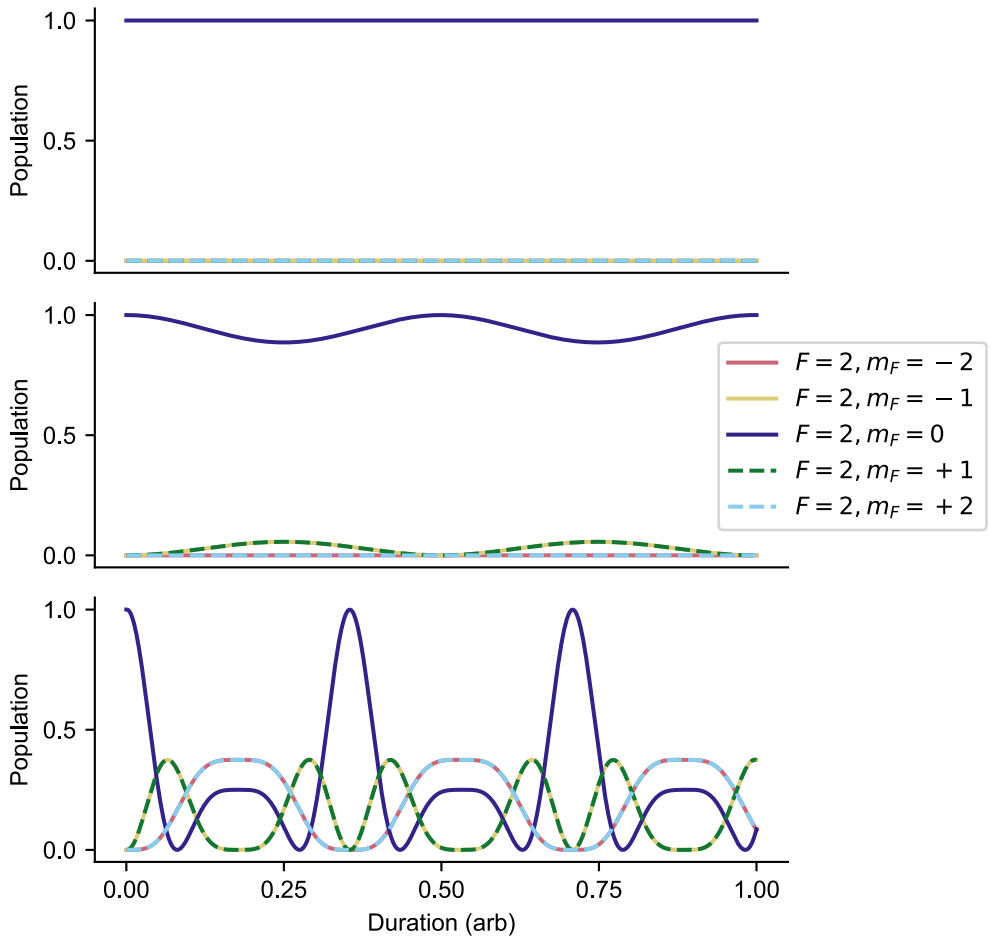


Figure 8.4: Simulation of the time evolution of each of the states in a five-level system acted upon with a DC electric field. The frequency splitting between states in these simulations, in arbitrary units, was chosen to be $\omega_0 = 2$. We see that as the coupling strength Λ becomes comparable to the frequency splitting, we are able to coherently transfer population between the $F = 2$ states in $S_{1/2}$.

numerically[2]. The results of the numerical simulations for a few different values of the coupling strength Λ are shown in Fig. 8.4. We see that when the coupling strength becomes comparable to the frequency splitting, we are able to drive the transitions between the states within $F = 2$ in $S_{1/2}$.

We used the numerical simulations together with the measured population of the $|S_{1/2}, F = 2, m_F = 0\rangle$ state to fit the frequency splitting between states within $F = 2$ in $S_{1/2}$ and the coupling strength Λ in Eqn. 8.8. We obtained a frequency splitting of 2.451(3) MHz and a coupling strength of 0.705(5) MHz. We can think of this effect as two identical photons from the counter-propagating

[2]Numerical simulations performed using the python Qutip package.

beam driving the (nearly degenerate) transitions within $F = 2$ off-resonantly, with a detuning equal to the frequency splitting between neighbouring states. In this case, the population in the $|S_{1/2}, F = 2, m_F = 0\rangle$ state is given by $P = 1 - \frac{\Omega^2}{\Omega^2+\delta^2} \sin^2\left(\frac{\sqrt{\Omega^2+\delta^2}t}{2}\right)$, where Ω is the Rabi frequency on resonance and $\delta \approx 2.451(3)$ MHz is the detuning [38]. By fixing the detuning and fitting this model to the $|S_{1/2}, F = 2, m_F = 0\rangle$ data with the only free parameter being the Rabi frequency, we find $\Omega \approx 0.705(4)$ MHz. This is consistent with the coupling strength $\Lambda \approx 0.705(5)$ MHz we found before, as expected. The expected frequency splitting at the given magnetic field, calculated numerically, is ≈ 2.33 MHz. This is consistent with the fitted value, and we believe that the difference of ≈ 0.12 MHz is due to a Stark shift on these transitions from the beam. To check our simulations further, we used the fit parameters obtained above to also plot the numerically simulated evolution of the remaining states in $F = 2$ in $S_{1/2}$ in Fig. 8.3. We see an excellent agreement between our numerical simulations and the experimental data for all five states, thus verifying our theoretical model.

This effect violated our general assumption at the time that a single far-detuned Raman beam can only generate a Stark shift but should not be able to transfer population between states. However, as we can see from Fig. 8.3, our theoretical model explains the experimental data very well, so how can we physically explain this?

Physically, we can think of this effect in two ways. First, as explained above, we can think of this as two identical photons off-resonantly driving a Raman transition between states within $F = 2$ in $S_{1/2}$. Second, we can think of this in the dressed state picture. The strong applied field results in mixing (dressing) of the states within $F = 2$ in $S_{1/2}$ such that they are, in general, no longer eigenstates of the Hamiltonian, except for specific durations, corresponding to the zeroes of the $|F = 2, m_F = 0\rangle$ state evolution.

To mitigate this effect without reducing the beam intensity, thus sacrificing the gate speed, we can change the beam polarisation. Experimentally, we found that this effect is no longer observed when the beam polarisation vector is set

to $\vec{e} = \vec{e}_{\sigma_-}/2 - \vec{e}_{\sigma_+}/2 - \vec{e}_{\pi}/\sqrt{2}$. This makes sense: numerical simulations of the Raman transition between the state $|S_{1/2}, F = 2, m_F = 0\rangle$ and the state $|S_{1/2}, F = 2, m_F = +1\rangle$[3] show that the transition strength is proportional to $\vec{e}_{\sigma_+}\vec{e}_{\pi} + \vec{e}_{\pi}\vec{e}_{\sigma_-}$. Thus, the opposite signs of the $\vec{e}_{\sigma_+}$ and the $\vec{e}_{\sigma_-}$ components in the chosen beam polarisation result in a zero net transition strength.

Final beam polarisation choice

To avoid the population transfer within $F = 2$ in $S_{1/2}$ that we observed in Fig. 8.3, the polarisation of the high-intensity counter-propagating beam in our experiment is constrained to $\vec{e}_{\mathrm{counter-propagating}} = \vec{e}_{\sigma_-}/2 - \vec{e}_{\sigma_+}/2 - \vec{e}_{\pi}/\sqrt{2}$. There are no such constraints on the ground and the metastable beams since they are a factor of ~ 100 lower in intensity and thus this effect is negligible, as we saw in Fig. 8.4.

The ground and the metastable level beams are delivered to the trap using the same optical fibre and are designed to have the same polarisation. Their polarisation vector was set to $\vec{e}_{\mathrm{G-M}} = \vec{e}_{\sigma_-}/2 + \vec{e}_{\sigma_+}/2$. This resulted in a good compromise of the single-qubit π-pulse duration between the ground and the metastable level qubits.

8.3 Motional mode structure

Driving two-qubit gates requires coupling to the ions' motion; hence, we must accurately measure the ions' motional mode frequencies. In addition, before the gate, we cool all modes near the motional ground state. A cold gate mode makes the gate more robust to small parameter imperfections, such as the gate detuning. Cooling the remaining modes that have a projection on the gate drive minimises coupling to them during the gate, making decoupling the spin and the motion at the end of the gate easier.

[3]The state $|S_{1/2}, F = 2, m_F = -1\rangle$ would behave in the same way from symmetry.

8.3.1 Choice of gate mode

The Raman beams in our setup are incident parallel to the trap axis in order to only couple to the radial motion of the ions. There are two sets of motional modes whose directions are at $\approx 45^\circ$ to the Raman beams, given our trap geometry (see Section 4.1). In each direction, there are two normal modes: a centre-of-mass (CoM) mode (in-phase motion) and a rocking mode (out-of-phase motion). For the gate experiments described in this chapter, the radial motional mode frequencies were:

- Lower frequency (LF) modes: 1.65 MHz (rocking) and 1.72 MHz (CoM).
- Higher frequency (HF) modes: 1.91 MHz (rocking) and 1.98 MHz (CoM).

We cool all four of these modes using sideband cooling with the Raman beams on the ground level qubit. For all four modes, we are able to achieve a temperature of $\bar{n} < 0.1$.

In general, out-of-phase modes have lower heating rates since they are insensitive to common-mode noise. Additionally, the motional heating rate generally decreases with the motional mode frequency [124]. Therefore, we chose to work with the HF rocking mode at 1.91 MHz for two-qubit gates ($\omega_m = 1.91$ MHz using the notation in Section 8.1).

8.3.2 Motional heating rates

We measured the heating rate of the ions by cooling them down to their motional ground state, waiting for some duration τ, and then measuring the resulting ion temperature as a function of τ [60]. We measured a heating rate of the LF CoM mode (1.72 MHz) of ≈ 150 quanta/s and of the HF CoM mode (1.98 MHz) of ≈ 30 quanta/s. We also tried measuring the heating rates of the two rocking modes. When probing for up to $\tau = 100$ ms, we observed no measurable increase in the ion temperature. Longer probing durations were not possible due to frequent ion loss caused by background collisions. This suggests that the rocking mode heating is $\ll 10$ quanta/s and will not limit the gate quality.

8.3.3 Lamb-Dicke parameter

The Lamb-Dicke (LD) parameter characterises the strength of the coupling between the applied qubit drive and the ion's motion. Hence, for a fixed single-qubit Rabi frequency it sets the two-qubit gate speed.

The Lamb-Dicke parameter is equal to the ratio between the ion's ground state wavefunction extent z_0 and the wavelength of the light. For a laser beam with a given $\vec{k}$-vector incident at an angle θ relative to the mode direction, it is thus given by $\eta = \left|\vec{k}\right| z_0 \cos(\theta)$ [59]. In the case of two counter-propagating Raman beams, the effective $\vec{k}$-vector is given by $\vec{k} = 2\vec{k}_{532}$, where $\vec{k}_{532}$ is the $\vec{k}$-vector of a single 532 nm beam. The extent of the ion's ground state wavefunction is given by $z_0 = \sqrt{\frac{\hbar}{2m\omega_m}}$, where m is the mass of the ion and ω_m is the frequency of the motional mode. The frequency of the mode used for the two-qubit gates in this chapter is 1.91 MHz, resulting in $z_0 \approx 4.4$ nm. Given our trap geometry (see Section 4.1), we expect this mode to be at an angle of $\approx 45°$ relative to the Raman beams. Thus, we expect an LD parameter of $\eta \approx 0.073$.

To experimentally measure the LD parameter, we compared the π pulse duration of the qubit transition and the π pulse duration of the blue sideband (BSB) transition immediately after sideband cooling. In the limit of $\bar{n} = 0$, the ratio between the two π times is equal to the LD parameter [60]. We measured a carrier π time of ≈ 5.22 µs and a BSB π time of ≈ 68 µs, resulting in an LD parameter of ≈ 0.0768. This is consistent with our calculation above. Using the experimentally measured LD parameter, we calculated the angle between the gate mode and the Raman beams to be $\approx 42.7°$.

8.4 Gate implementation

8.4.1 Qubit encoding

The magnetic field for the gate experiments was set to the field-insensitive point for the metastable level qubit we found in Chapter 7, i.e. ≈ 3.333 G. Thus, the

metastable level qubit encoding that we used for both the M-M gates and the G-M gates was $|0_M\rangle = |\mathrm{D}_{5/2}, F = 2, m_F = +2\rangle$ and $|1_M\rangle = |\mathrm{D}_{5/2}, F = 3, m_F = +2\rangle$.

At this magnetic field, the qubit encoding in the ground level resulting in the lowest magnetic field sensitivity is $|0_G\rangle = |\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ and $|1_G\rangle = |\mathrm{S}_{1/2}, F = 1, m_F = 0\rangle$[4]. This is the ground level qubit encoding we used for both the G-G gates and the G-M gates.

8.4.2 General gate sequence

We begin the experiment by Doppler cooling both ions. We then sideband cool all four radial modes as described in Section 8.3. Following the cooling process, we implement the SPAM protocol described in Chapter 6 with U_q in Fig. 6.1 being the MS gate pulse. The ground level qubit is prepared following the sequence in Section 6.4.1. The metastable level qubit is prepared following the sequence in Section 7.2.1.

8.5 Gate calibration

This was the first implementation of two-qubit gates in our experiment, and hence, all gate calibrations were performed manually to enable easier debugging and understanding of the system's behaviour. However, all of the calibrations below could be automated.

8.5.1 Equal ion illumination

As a first step, we must ensure that the Rabi frequencies on the two ions are balanced. We did this by measuring the light shifts produced by each of the beams on one of the quadrupole transitions and fine-tuning the positions of the beams until they were equal between the two ions. By manually applying this procedure we were able to balance the Rabi frequencies between the ions to within $\sim 5\%$.

[4]This qubit is a clock qubit at zero magnetic field.

8.5.2 Equal sideband strength

The strengths of the BSB and RSB tones must be equal during the gate. Due to the nonlinearity of the AOM response to the applied RF power, simply applying equal RF signals or ensuring that the Rabi frequencies produced when just one of these tones is switched on are equal is not enough to guarantee that the two tones will be balanced when they are applied simultaneously.

To balance the two tones, we perform a simple Rabi flop experiment, with the following beam parameters:

- Counter-propagating beam with both the BSB and RSB tones applied, as shown in Fig. 8.2.
- The ground level beam on, detuned by either $+\omega_m$ or $-\omega_m$ relative to the frequency we use to drive the ground level qubit on resonance.

When the ground beam detuning is set to $+\omega_m$, the ground beam and the RSB tone drive the ground level qubit on resonance, in the presence of an additional tone detuned by $2\omega_m$ relative to the qubit transition. Similarly, when the ground beam detuning is set to $-\omega_m$, we instead drive the ground level qubit transition using the BSB tone. In the case of $\omega_m \approx 1.9$ MHz, the additional tones are detuned by $2\omega_m \approx 3.8$ MHz and do not couple to any transitions within the ground level. Thus, they only introduce a small light shift to the qubit transition.

For each of the two ground beam detunings, we perform a Rabi flop experiment, scanning the pulse duration. We adjust the RF strengths of the BSB and the RSB tones until the Rabi frequencies in the two scans are balanced. By manually applying this procedure, we are able to balance the two tones to within $\sim 1\%$.

8.5.3 Calibration of the gate detuning and duration

To calibrate the gate detuning (δ in Fig. 8.1) and duration (τ_{MS} in Eqn. 8.5), we followed the procedure described in Ref. [125, 126]. This is an iterative procedure, where each parameter is scanned individually and optimised, until the parameter values converge.

For the purposes of this section, we define the *asymmetric detuning* as the carrier detuning from ω_0 in Fig. 8.1, causing the BSB and RSB tones to be asymmetrically detuned from the carrier. We will also call the gate detuning, δ, the *symmetric detuning*.

To obtain initial rough estimates of the gate parameters, we use the carrier Rabi frequency, Ω_{C}, measured during the calibration in Section 8.5.2 and the measured Lamb-Dicke parameter in Section 8.3.3, η. An initial guess for the gate Rabi frequency is then given by $\Omega_{\mathrm{MS}} = \eta\Omega_{\mathrm{C}}$. From this, we can estimate the gate duration and the gate detuning using Eqn. 8.5 and Eqn. 8.6. Once we have initial guesses for the gate parameters, the calibration procedure goes as follows.

1. Scan the asymmetric detuning using a duration much longer than the gate duration, with a detuning that is a little higher than the expected gate detuning (to minimise excitation of the motion). The optimal value is the one that maximises the p_{11} signal (both ions bright).

2. Scan the gate detuning (the symmetric detuning) at approximately the gate duration. From this scan, determine the detuning at which the populations p_{00} (both ions dark) and p_{11} (both ions bright) are equal.

3. At that value of the asymmetric detuning, scan the gate duration. Choose the shortest duration at which the p_{01+10} signal (one ion dark, one ion bright) is at zero.

Steps 2 and 3 of this procedure are repeated iteratively until the values for the gate detuning and the gate duration converge. In what follows the calibration method described above is performed manually. However, it could be automated, for example, by means of Bayesian optimisation and feedforward to iteratively find the best parameters as demonstrated in Ref. [126].

8.6 G-G gates

Gates between ground level hyperfine qubits are well-understood and widely used in quantum computing. Although G-G gates are not crucial for our architecture outlined in Chapter 2, the ground level offers a simpler structure compared to the metastable level, making the gate implementation simpler. As a result, we first implemented G-G gates to validate our setup and to refine our gate implementation and calibration procedures.

The results from implementing the G-G gates following the procedure above are shown in Fig. 8.5. In Fig. 8.5(a) we show a time scan of the gate dynamics. The measured gate duration after all calibration procedures in Section 8.5 were performed was 207.4 µs.

During the calibration of the sideband strengths described in Section 8.5.2, we measured a carrier π pulse duration, when both sideband tones are on, of ≈ 7.6 µs. Given the measurement of the Lamb-Dicke parameter in Section 8.3.3, we expect the gate duration to be ≈ 198 µs. We believe that the slight discrepancy between the expected and measured gate durations is due to slow power drifts in the counter-propagating beam. At the time these measurements were taken, we were seeing $\sim 20\%$ fluctuations in the beam power over several hours.

To quantify the quality of the gate, we performed a parity scan at the gate duration. This was done by performing a $\pi/2$ carrier pulse at the end of the gate and measuring the p_{01+10} population. The result of this scan is shown in Fig. 8.5(b). The resulting fringe contrast is 0.864(19). We used partial state tomography (PST), as described in Ref. [38] to calculate the gate fidelity, resulting in a fidelity of 90.8(11)%. The dominant source of error here was qubit decoherence, and a detailed discussion is provided in Section 8.9.

8.7 M-M gates

The M-M gates are implemented completely analogously to the G-G gates, except in this case, we turn on the metastable beam in Fig. 4.11.

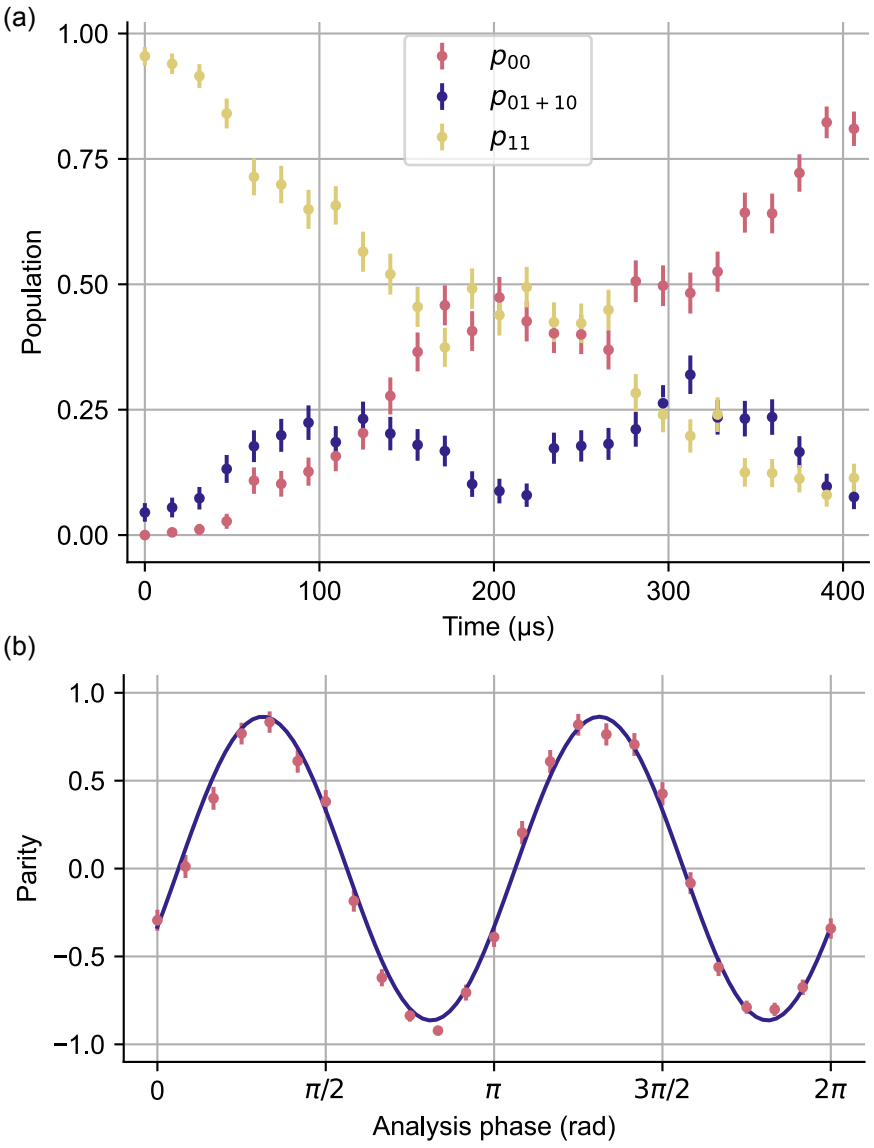


Figure 8.5: G-G gate results. The measured contrast of the pariy scan is 0.864(19) and the populations at the gate duration are 0.951(12) resulting in a gate fidelity of 90.8(11)%.

The calibration procedure we apply in this case follows exactly the procedure described in Section 8.5. Since the BSB and RSB tones are applied to the counter-propagating beam, common to both gate types, calibration of their strengths using the ground beam is sufficient[5]. Similarly, since the ground and the metastable beam are delivered to the trap using the same optical fibre, we can use the ground beam to ensure that the Rabi frequencies of the two ions are balanced.

Hence, the only calibration we need to perform for the implementation of the M-M gate is the calibration of the gate duration and detuning. For this we follow the exact same procedure as in Section 8.5 but with the metastable beam.

A duration scan of the gate after optimisation is shown in Fig. 8.6(a) and the corresponding parity scan is shown in Fig. 8.6(b). The gate duration was 500 µs. This is roughly consistent with the G-G gate duration in Section 8.6 and the ratio of the single-qubit Rabi frequencies between the ground and metastable level qubits for the chosen beam parameters (≈ 66 kHz and ≈ 30 kHz, respectively).

We used PST to obtain a measured gate fidelity of 87.9(15)%. As with the ground level gate, we believe that this is limited by qubit decoherence and a detailed discussion is provided in Section 8.9[6].

8.8 G-M gates

In this section, we describe the implementation of the G-M gates. This is the first ever demonstration of this type of gate.

The G-M gates are implemented by simply turning on the control fields for the G-G gates and the M-M gates at the same time. The only differences in the implementation of the G-M gates, compared to G-G and M-M gates, are the state preparation sequence and some parts of the calibration procedures.

[5]Of course, we can use the analogous calibration with the metastable beam. The point we wish to make here is that the beam configuration we have chosen in the experiment ensures that for most calibrations, calibrating only one type of gate is sufficient to calibrate them all.

[6]As a sneak peek, the qubit coherence time here is significantly reduced compared to the one measured in Chapter 7 due to the light shift introduced by the intense counter-propagating beam on the qubit transition.

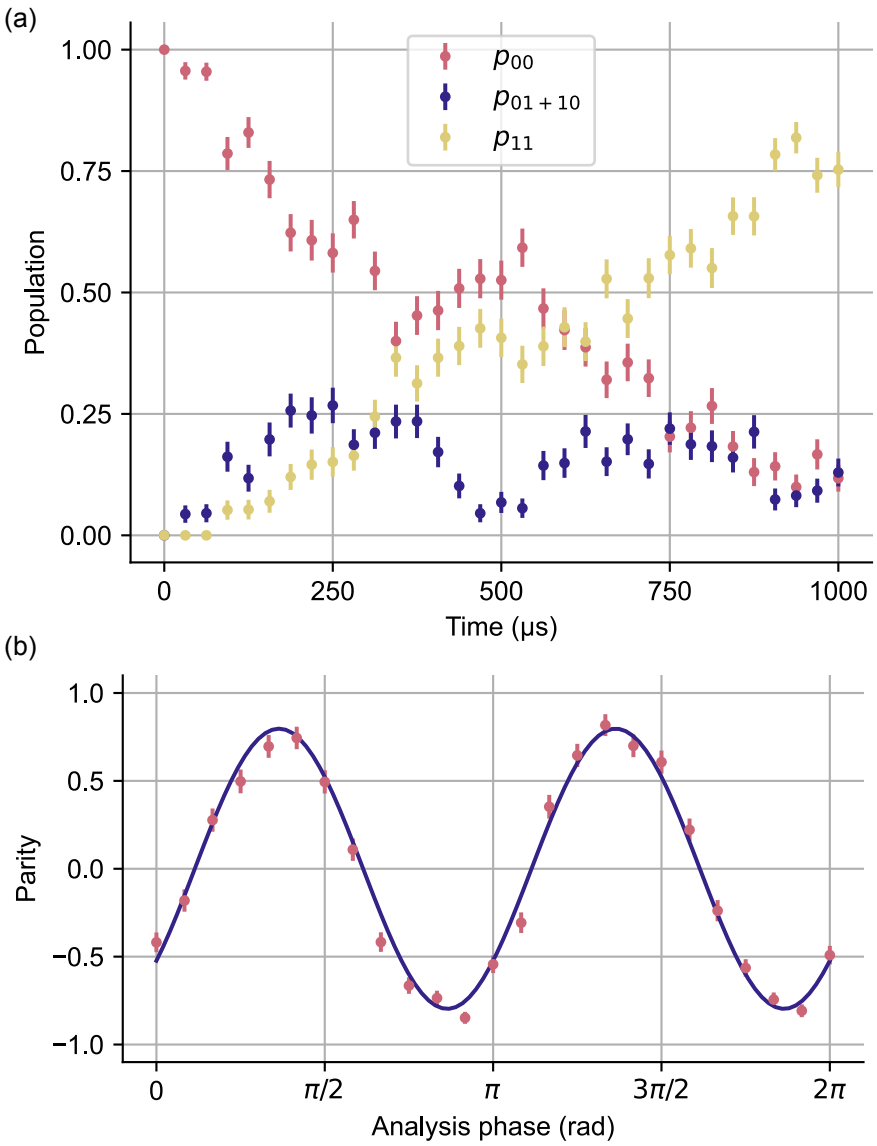


Figure 8.6: M-M gate results. The measured contrast of the parity scan is 0.831(26) and the populations at the gate duration are 0.927(14) resulting in a gate fidelity of 87.9(15)%.

8.8.1 Experiment sequence

The only difference in the experiment sequence of the G-M gates compared to the G-G and M-M gates is in the state preparation stage. As before, we start by Doppler cooling the two ions. Next, we must prepare one ion in the ground level qubit state $|0\rangle_G$ and the other ion in the metastable level qubit state $|0\rangle_M$. For the work in this chapter, we used global Raman beams that illuminate both ions simultaneously. Therefore, to prepare the two ions in these states, we first initialise both of them in the state $|\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$ and then transfer the population to the state $|\mathrm{S}_{1/2}, F = 1, m_F = +1\rangle$ using the ground and co-propagating Raman beams. This is followed by a 1762 nm $\pi/2$ pulse between the $|\mathrm{S}_{1/2}, F = 1, m_F = +1\rangle$ state and the $|0\rangle_M = |\mathrm{D}_{5/2}, F = 3, m_F = +2\rangle$ state. Next, we perform a population detection step and only keep outcomes in which the first ion is bright (remained in $\mathrm{S}_{1/2}$) and the second ion is dark. As a result, the second ion is initialised in the state $|0\rangle_M$.

Since the measurement on the ion that we wish to encode in the ground level is bright, the population is scrambled within the ground level and thus the qubit state must be initialised again. However, the scattered photons during the measurement will have heated up the ions. Hence, we implement another round of Doppler cooling, followed by sideband cooling. Cooling in this case is only performed on the G qubit, which in turn sympathetically cools the M qubit as well. Following the cooling steps, we initialise the G qubit in the state $|0_G\rangle = |\mathrm{S}_{1/2}, F = 2, m_F = 0\rangle$. To do this, we use the same steps for the preparation of the ground level qubit as in Section 6.4.1.

As before, the rest of the experiment is embedded in the SPAM protocol in Chapter 6 with U_q being the gate itself.

8.8.2 Gate calibration

Equal sideband strength

The BSB and RSB are applied to the counter-propagating beam, common to all types of gates. Hence, as before, this calibration is done with the ground beam

and no modifications are necessary for the G-M gate.

Equal ion Rabi frequency

Similar to ensuring that the ions are equally illuminated by the two beams as described in Section 8.5, here we need to ensure that the Rabi frequencies for driving the ground and the metastable level qubit transitions are equal. To do this, we initialise one ion in the ground level and one ion in the metastable level as described in Section 8.8.1. We then turn on the ground level beam, the metastable level beam, and the counter-propagating beam, with the ground and metastable level beam detunings set on resonance with the respective qubit types. We then perform a carrier time scan and adjust the powers of the ground and the metastable beams until the Rabi frequencies of the two ions match.

Calibration of the asymmetric detuning

A non-zero asymmetric detuning correction arises due to a differential light shift on the carrier qubit transition due to the applied BSB and RSB tones. In the case of a G-M gate, the amount of light shift is different for the two ions due to the different level structure of the ground and the metastable levels, resulting in different qubit frequencies.

To calibrate the asymmetric detunings for this gate, we calibrate the asymmetric detuning for the G-G gate while leaving the metastable beam on. We then do the same for the M-M gate, but this time in the presence of the ground beam. That way, we can calibrate the asymmetric detuning irrespective of the different LS experienced by the two ions while also taking into account the light shift on the ground level qubit due to the metastable beam and vice versa.

Calibration of the gate detuning and duration

Once all other calibrations have been performed, the calibration of the gate detuning and duration proceeds in exactly the same way as described in Section 8.5. The gate detuning is set by the BSB and RSB tones applied to the counter-propagating beam and is thus easily tuned simultaneously for both

qubits. Similarly, the gate duration is set by turning the beams on and off and is defined by the rising/falling edge of the pulse shaper. Since the pulse shaping is also implemented on the counter-propagating beam, the gate duration is set simultaneously for both qubits.

8.8.3 G-M gate results

A scan of the gate duration after optimisation is shown in Fig. 8.7. The gate duration was 500 µs, equal to the M-M gate duration. This is expected: for the Raman beam parameters we chose, the metastable level qubit carrier Rabi frequency is lower compared to the ground level qubit and thus determines the duration of the G-M gate.

However, we see a slightly worse decay of the populations at times close to twice the gate duration compared to both the M-M and the G-G gates. As we will see shortly in Section 8.9, this makes sense. The gate duration here was longer than the G-G gate duration, but the coherence time of the ground level qubit was shorter than the metastable level qubit. Thus, this gate performed slightly worse compared to both the G-G and the M-M gates.

Unfortunately, the ions were lost during the duration scan presented in Fig. 8.7. The performance suddenly degraded after this scan, and we could not get the gates working again in the couple of days I had before I had to leave the lab. We were, therefore, unable to obtain a parity scan within the timeline of this thesis. We suspect that the source of the sudden degradation in performance is degradation in the Raman laser setup, for example, an increase in the laser beam phase or intensity noise.

8.9 Discussion

We have demonstrated the implementation of two-qubit gates between ground and metastable level qubits in $^{137}Ba^{+}$. In principle, there is no fundamental limit to the performance in any of these gate types, except for Raman scattering errors that are present in all Raman-driven gates.

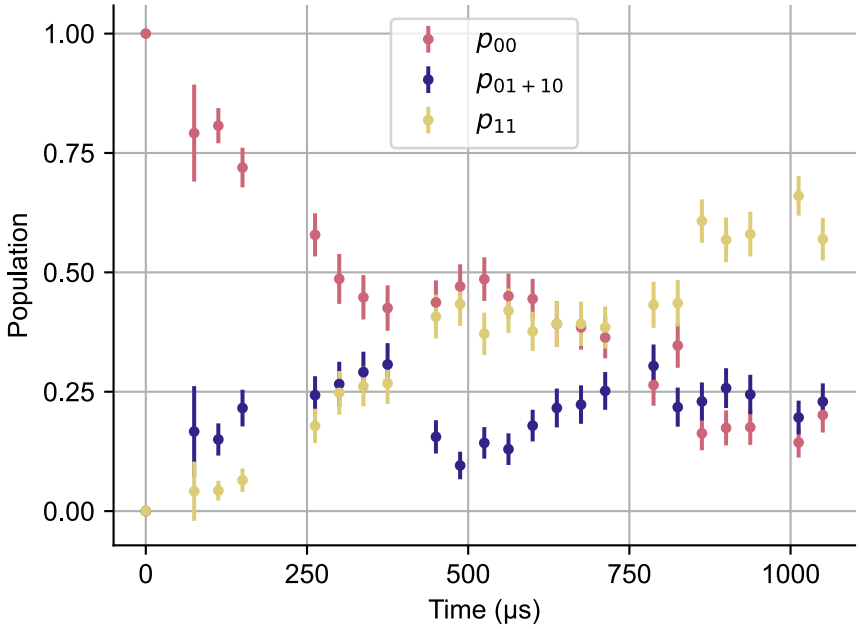


Figure 8.7: G-M gate implementation. We see that the gate duration scan is comparable to the M-M gate duration scan. This is expected since the duration of this gate is determined by the M-M gate duration, and, as we will see in Section 8.9, the gate quality is limited by the qubit coherence, common to both the M-M and the G-M gates. Thus, we expect the fidelity of this gate to be comparable to the M-M gate fidelity.

Fidelity measurements of the G-G gates and the M-M gates showed gate errors of $\sim 10\%$ in both cases, with the M-M gates performing slightly worse than the G-G gates. Following the treatment in Ref. [31], we estimated that the motional mode temperatures, the motional heating rates, and the accuracy of our calibrations would result in gate errors much lower than the ones we measured. Since the gate fidelity was lower for the M-M gate, which is longer in duration than the G-G gate, we suspected that the fidelity may be limited by qubit coherence.

The magnetic field used in this chapter was set to the field-insensitive point for the metastable level qubit we found in Chapter 7. In Chapter 7, we measured the coherence time in this qubit to be 31(5) s relative to an RF source and 9.1(10) s relative to the Raman beatnote formed by the metastable and co-propagating beams. This is clearly much longer than the M-M gate duration and would not explain the M-M gate fidelity we measured. However, in Chapter 7, we also saw that applying laser beams to the ion resulted in a reduction of the

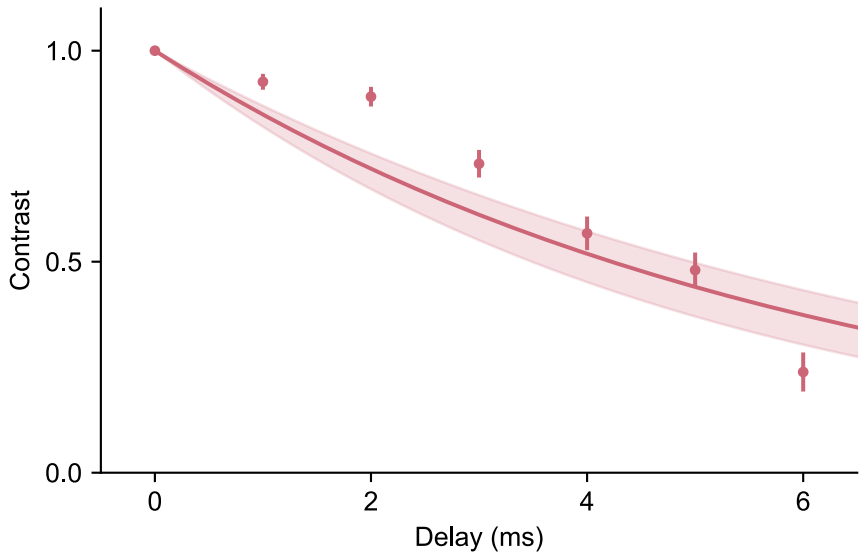


Figure 8.8: Coherence time of the metastable level qubit measured using the RF antenna in the presence of the counter-propagating Raman beam. The measured coherence time is 6.1(11) ms.

qubit coherence from 31(5) s to 21.9(21) s. In that measurement, we only applied the 493 nm and the 650 nm beams, which are very low in intensity. When we perform two-qubit gates, we apply the counter-propagating Raman beam, which has a much higher intensity compared to the 493 nm and the 650 nm beams, as specified in Section 8.2.2.

To evaluate the effect of the counter-propagating Raman beam on the coherence of the metastable level qubit, we performed the same two-point Ramsey measurement as in Section 7.3, except this time turning on the counter-propagating beam during the Ramsey delay time. The results from this measurement are shown in Fig. 8.8. We see that in the presence of this beam, the coherence time of the metastable level qubit was reduced to 6.1(11) ms. This is ~ 10 times longer than the M-M gate duration, thus accounting for the majority of the error that we see.

The most robust way to reduce the effects of beam intensity noise on the qubit is to work with a beam polarisation that results in no net light shift of the qubit transition. As we saw in Chapter 7, numerical simulations suggest that for the particular metastable level encoding here this is achieved by using a beam with equal components in all three polarisation states. While such a beam would cause no net light shift of the qubit, we must also ensure that it is able to drive

the qubit transition in both the ground and the metastable level qubits. The numerically calculated Raman matrix element for driving the metastable level qubit at the field-insensitive point, in arbitrary units, is given by

$$\mu_M = -5.84\sigma_-\sigma_- + 5.86\pi\pi + 0.16\sigma_+\sigma_+. \tag{8.10}$$

Therefore, the optimal beam polarisations to drive this qubit, while cancelling out the light shift, are:

$$\vec{e}^{\,1}_{\mathrm{M}} = \left(\vec{e}_{\sigma_-} + \vec{e}_{\sigma_+} + \vec{e}_{\pi}\right)/\sqrt{3} \tag{8.11}$$

$$\vec{e}^{\,2}_{\mathrm{M}} = \left(-\vec{e}_{\sigma_-} + \vec{e}_{\sigma_+} + \vec{e}_{\pi}\right)/\sqrt{3}. \tag{8.12}$$

We now check whether it is possible to drive the ground level qubit with these polarisations. The Raman matrix element for driving the ground level qubit at this magnetic field, in arbitrary units, is given by

$$\mu_G = -5.23\sigma_-\sigma_- + 5.23\sigma_+\sigma_+. \tag{8.13}$$

We see that the choice of polarisation in Eqn. 8.11 and Eqn. 8.12 allow for efficiently driving this qubit as well.

We have, therefore, found a set of beam polarisations that work for driving both the ground and the metastable level qubit efficiently, even if not in the most optimal way power-wise, that minimise the light shift of the metastable level qubit. These simulations show that the qubit encodings we have chosen here are compatible with each other and can be driven simultaneously while minimising the effects on the metastable level qubits (i.e. our computational qubits).

Finally, to verify that the G-G gates are also limited by the qubit coherence, we measured the coherence of the ground level qubits with the ground and co-propagating Raman beams following the same method as in Section 7.2.2. The results from this measurement are shown in Fig. 8.9. The measured coherence time is 3.55(18) ms. Given the measured coherence time of the metastable level qubit, and the ratio of the numerically simulated magnetic field sensitivities of the ground and metastable qubits, we expect the coherence time of the ground

level qubit, in the absence of laser noise, to be ~ 300 ms. The measured ground qubit coherence time in Fig. 8.9 is thus limited by laser noise. This is not unexpected. The Raman beatnote used for driving the metastable level qubit is generated using an AOM network from a single laser source, as described in Section 4.3.4. However, the Raman beatnote for driving the ground level qubit is generated from two separate laser sources, so it is not too surprising that it has a higher phase noise. The coherence time of the ground level qubit could thus be increased by improving the phase noise of the Raman beatnote and by improving the magnetic field noise in the lab.

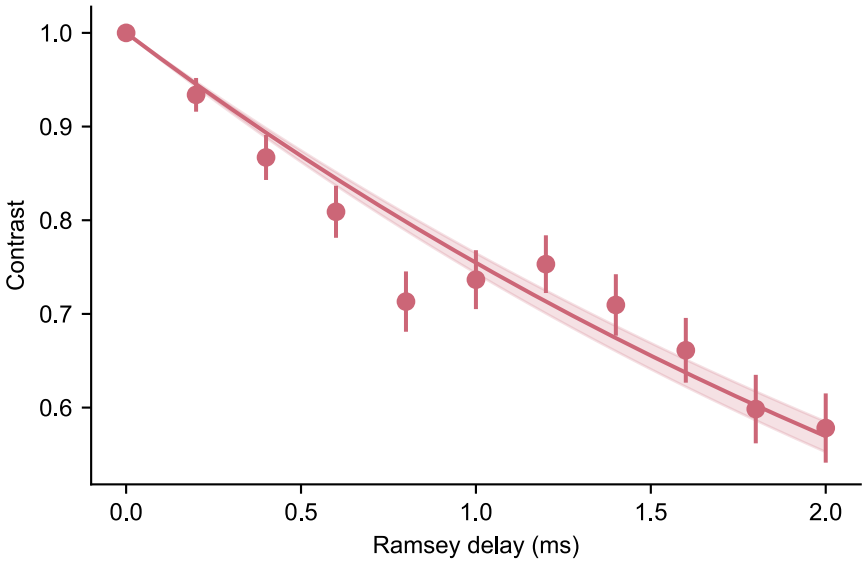


Figure 8.9: Coherence time measurement of the ground level qubit using the co-propagating Raman beams. The measured coherence time is 3.55(18) ms.

The coherence time of the ground level qubit also accounts for the slightly worse performance of the G-M gate compared to the M-M gate. When we implemented the G-G gate, the gate duration was 207.4 µs and the resulting gate fidelity was 90.8(11)%. However, the speed of the G-M gate is determined by the M-M gate speed and is ≈ 500 µs. The coherence time of the ground level qubit we measured above is shorter than the coherence time of the metastable level qubit in the presence of the counter-propagating beam we measured in Fig. 8.8 and thus limits the gate fidelity. Therefore, based on the measurements in this section, we expect the G-M gate fidelity for the gate data presented in Fig. 8.7 to be $\approx 83\%$.

8.10 Summary

In this chapter, we demonstrated the implementation of two-qubit entangling gates between two ground level qubits (G-G gates), two metastable level qubits (M-M gates) and a ground level qubit and a metasable-level qubit (G-M gates). The fidelities of the G-G and M-M gates we measured were 90.8(11)% and 87.9(15)%, respectively.

We identified the main contribution to the gate error to be qubit decoherence. In the case of M-M gates, we found that this is solely due to intensity noise in the counter-propagating beam, which reduced the qubit coherence time from 31(5) s measured in Chapter 7 to 6.1(11) ms. By performing numerical simulations, we designed a way to set up the beam polarisations such that the light shift on the metastable qubit is cancelled out, while still being able to drive both the metastable level qubit and the ground level qubit efficiently. This shows that the qubit encodings we have chosen here are compatible, and the demonstrations in this chapter could be performed with much lower error in future experiments.

The experimental results and numerical analysis in this chapter demonstrate that we are able to leverage all of the capabilities of a mixed-species crystal using a single ion species with multiple qubit encodings. We can encode the computational qubits in the metastable level and we have shown that we are able to perform entangling gates between two such qubits. We have further shown that we can encode any ancilla qubits in the ground level, and we have, for the first time ever, demonstrated an entangling gate between a ground level qubit and a metastable level qubit, the equivalent of the mixed species gate. The ability to drive this gate directly makes the system completely insensitive to the coherence of any population transfer pulses between the ground and the metastable level. This is the final building block necessary for implementing the full architecture described in Chapter 2 in a real system.

9 Outlook

In this thesis, we have developed and experimentally demonstrated the building blocks of an architecture that addresses some of the challenges in scaling trapped-ion quantum computers. Our architecture leverages the vast state space available within a single ion species to encode qubits within different level manifolds, the ground level (G) and the metastable level (M). This enables the implementation of techniques such as mid-circuit measurement, cooling, and reset of the G qubits while preserving the information stored in the M qubits.

Previously, these capabilities have been demonstrated in systems using multiple atomic species, where one species is used for information processing and storage, and the other species is used for cooling and measurement. However, scaling up a system based on multiple different atomic species is challenging. The different masses of the ions in these experiments make implementing many operations difficult, for example, efficient sympathetic cooling, transporting mixed-species crystals across the ion trap chip quickly and with low heating, and implementing high-fidelity mixed-species entangling gates.

Our architecture allows one to implement all of these capabilities by using an additional level manifold within a single type of ion rather than an additional ion species. We demonstrated all fundamental operations required for this architecture using trapped $^{137}Ba^{+}$ ions.

We showed that we can implement clock qubits in the metastable level of $^{137}Ba^{+}$. We experimentally evaluated one of the possible metastable clock qubits and demonstrated that we can achieve a coherence time of 31(5) s, which does not limit the performance of the system and is comparable to the natural lifetime of the metastable level of $^{137}Ba^{+}$. We further demonstrated the isolation of this qubit from the light used to perform mid-circuit dissipative operations. We measured a qubit coherence time in the presence of the laser light used for these operations of 21.9(21) s. This reduction will still not limit the performance of the system and we showed that it could be made lower in future implementations by adjusting the polarisation of the laser beams.

We implemented two-qubit entangling gates between a pair of metastable-level qubits (M-M gates) and between a ground-level qubit and a metastable-level qubit (G-M gates). The M-M gates are required to build a universal unitary gate set and thus be able to implement a generic quantum algorithm. The G-M gates are the equivalent of the mixed-species gates and hence facilitate quantum non-demolition measurements without the need to change the qubit encoding. This makes our architecture completely insensitive to the coherence of the population transfer pulses between the ground and the metastable level. Hence the performance of the system is only dependent on the fidelity of the qubit operations themselves.

Finally, we also showed how the ability to perform mid-circuit measurements can be used to significantly reduce errors in SPAM. We developed a protocol that relies on several mid-circuit measurements during SPAM to detect errors and either remove the corresponding shots using post-selection or convert the errors into erasures, one of the easiest types of error to correct for. Experimentally implementing this protocol in $^{137}Ba^{+}$ has enabled us to demonstrate the lowest ever recorded SPAM errors in all three qubit types – optical (O): $7(4) \times 10^{-6}$, metastable (M): $5(4) \times 10^{-6}$, and ground (G): $8(4) \times 10^{-6}$.

We hope that the techniques we have developed here will be useful in building large-scale trapped-ion quantum computers in the future. In the long term, we expect that ion trap chips will adopt the QCCD architecture: a single trap

chip split up into several trapping zones, each hosting a small number of ions. Connectivity between zones is realised by shuttling the ions across the chip. The architecture we have developed here uses a single ion species, making the fast transport of ion crystals with low heating much easier to implement.

Furthermore, when running long circuits, both the time taken to execute the circuit and the many transport operations required will cause heating of the ions. In our architecture, this effect can be mitigated by cooling the ground-level ions, which in turn sympathetically cools the metastable-level ions. Since the masses of all ions in the crystal are equal, this can be done much more efficiently compared to mixed-species systems.

Additionally, building a large-scale quantum computer will require QEC. One of the essential requirements for QEC is the ability to perform QNDs to identify logical errors. Implementing QNDs requires entangling a computational qubit to an ancilla qubit, followed by measurement of the ancilla qubit that does not destroy the computational qubit. That way, the state of the computational qubit is detected but not destroyed, and the qubit can be used in the rest of the computation. The G-M gate we have demonstrated and the ability to measure ancilla (G) qubits while protecting computational (M) qubits enable performing QNDs in the system without changing the qubit encoding. Thus, the system performance does not depend on the coherence of the pulses used to transfer population between the G and M levels and is only determined by the quality of the qubit operations themselves.

Finally, in large-scale devices, the error with which we can prepare an N-qubit register in the state $|000...\rangle$ scales approximately linearly with the number of qubits in the limit of small single-qubit state preparation error. The SPAM protocol we developed in this thesis has enabled us to achieve world-record low SPAM errors on a single qubit. This protocol can be extended for large registers. The state preparation part of the protocol can be implemented in a repeat-until-success (RUS) way, truncating the number of attempts to give deterministic runtime with minimal truncation error. Any further errors detected during the

mid-circuit measurements of the SPAM protocol can either be removed by post-selecting these shots or by converting the errors into erasures. That way, the SPAM error in large registers is either converted to a reduced data rate, in the case of post-selection, or into an erasure, which is one of the easiest errors to correct.

In summary, we have developed and experimentally verified a set of techniques that overcome some of the roadblocks to scaling trapped-ion systems. These techniques can be used in building useful large-scale quantum computers in the future.

Bibliography

[1] Pengfei Wang et al. “Single ion qubit with estimated coherence time exceeding one hour”. In: *Nature Communications* 12.1 (Jan. 2021), p. 233.

[2] 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.

[3] 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.

[4] Craig R. Clark et al. “High-Fidelity Bell-State Preparation with $^{40}Ca^{+}$ Optical Qubits”. In: *Physical Review Letters* 127.13 (Sept. 2021), p. 130505.

[5] R. Srinivas et al. “High-fidelity laser-free universal control of trapped ion qubits”. In: *Nature* 597.7875 (Sept. 2021), pp. 209–213.

[6] 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.

[7] A. C. Hughes et al. “Benchmarking a High-Fidelity Mixed-Species Entangling Gate”. In: *Physical Review Letters* 125.8 (Aug. 2020), p. 080504.

[8] Richard P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6 (June 1982), pp. 467–488.

[9] Peter W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1484–1509.

[10] Artur K. Ekert. “Quantum cryptography based on Bell’s theorem”. In: *Physical Review Letters* 67.6 (Aug. 1991), pp. 661–663.

[11] Jacob Biamonte et al. “Quantum machine learning”. In: *Nature* 549.7671 (Sept. 2017), pp. 195–202.

[12] Seth Lloyd. “Universal Quantum Simulators”. In: *Science* 273.5278 (Aug. 1996), pp. 1073–1078.

[13] David P. DiVincenzo. “The Physical Implementation of Quantum Computation”. In: *Fortschritte der Physik* 48.9-11 (2000), pp. 771–783.

[14] Serge Haroche and Jean-Michel Raimond. “Quantum Computing: Dream or Nightmare?” In: *Physics Today* 49.8 (Aug. 1996), pp. 51–52.

[15] J. I. Cirac and P. Zoller. “Quantum Computations with Cold Trapped Ions”. In: *Physical Review Letters* 74.20 (May 1995), pp. 4091–4094.

[16] Loïc Henriet et al. “Quantum computing with neutral atoms”. In: *Quantum* 4 (Sept. 2020), p. 327.

[17] M. H. Devoret and R. J. Schoelkopf. "Superconducting Circuits for Quantum Information: An Outlook". In: *Science* 339.6124 (Mar. 2013), pp. 1169–1174.

[18] Lilian Childress and Ronald Hanson. "Diamond NV centers for quantum computing and quantum networks". In: *MRS Bulletin* 38.2 (Feb. 2013), pp. 134–138.

[19] Floris A. Zwanenburg et al. "Silicon quantum electronics". In: *Reviews of Modern Physics* 85.3 (July 2013), pp. 961–1019.

[20] Sonika Johri et al. "Nearest centroid classification on a trapped ion quantum computer". In: *npj Quantum Information* 7.1 (Aug. 2021), pp. 1–11.

[21] Abhinav Kandala et al. "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets". In: *Nature* 549.7671 (Sept. 2017), pp. 242–246.

[22] Tyson Jones et al. "QuEST and High Performance Simulation of Quantum Computers". In: *Scientific Reports* 9.1 (July 2019), p. 10736.

[23] Feng Pan and Pan Zhang. "Simulation of Quantum Circuits Using the Big-Batch Tensor Network Method". In: *Physical Review Letters* 128.3 (Jan. 2022), p. 030501.

[24] Peter W. Shor. "Scheme for reducing decoherence in quantum computer memory". In: *Physical Review A* 52.4 (Oct. 1995), R2493–R2496.

[25] Daniel Gottesman. *Stabilizer Codes and Quantum Error Correction.* arXiv:quant-ph/9705052. May 1997.

[26] Rajeev Acharya et al. "Suppressing quantum errors by scaling a surface code logical qubit". In: *Nature* 614.7949 (Feb. 2023), pp. 676–681.

[27] John Preskill. "Quantum Computing in the NISQ era and beyond". In: *Quantum* 2 (Aug. 2018), p. 79.

[28] Kishor Bharti et al. "Noisy intermediate-scale quantum algorithms". In: *Reviews of Modern Physics* 94.1 (Feb. 2022), p. 015004.

[29] Justin G. Bohnet et al. "Quantum spin dynamics and entanglement generation with hundreds of trapped ions". In: *Science* 352.6291 (June 2016), pp. 1297–1301.

[30] Michael J. Biercuk et al. "Optimized dynamical decoupling in a model quantum memory". In: *Nature* 458.7241 (Apr. 2009), pp. 996–1000.

[31] Christopher J. Ballance. *High-Fidelity Quantum Logic in* Ca^+. Springer Theses. Springer International Publishing, 2017.

[32] Chi Zhang et al. "Submicrosecond entangling gate between trapped ions via Rydberg interaction". In: *Nature* 580.7803 (Apr. 2020), pp. 345–349.

[33] M. Grassl, Th. Beth, and T. Pellizzari. "Codes for the quantum erasure channel". In: *Physical Review A* 56.1 (July 1997), pp. 33–38.

[34] Charles H. Bennett, David P. DiVincenzo, and John A. Smolin. "Capacities of Quantum Erasure Channels". In: *Physical Review Letters* 78.16 (Apr. 1997), pp. 3217–3220.

[35] Francesco Tacchino et al. "Quantum implementation of an artificial feed-forward neural network". In: *Quantum Science and Technology* 5.4 (Oct. 2020), p. 044010.

[36] M. D. Barrett et al. "Sympathetic cooling of $^{9}Be^{+}$ and $^{24}Mg^{+}$ for quantum logic". In: *Physical Review A* 68.4 (Oct. 2003), p. 042302.

[37] V. Schafer. "Fast gates and mixed-species entanglement with trapped ions". PhD thesis. University of Oxford, 2018.

[38] A. C. Hughes. "Benchmarking memory and logic gates for trapped-ion quantum computing". PhD thesis. University of Oxford, 2021.

[39] 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.

[40] O. Băzăvan et al. "Synthesizing a $\hat{\sigma}_z$ spin-dependent force for optical, metastable, and ground-state trapped-ion qubits". In: *Physical Review A* 107.2 (Feb. 2023), p. 022617.

[41] Takashi Nakajima et al. "Quantum non-demolition measurement of an electron spin qubit". In: *Nature Nanotechnology* 14.6 (June 2019), pp. 555–560.

[42] Xiao Xue et al. "Repetitive Quantum Nondemolition Measurement and Soft Decoding of a Silicon Spin Qubit". In: *Physical Review X* 10.2 (Apr. 2020), p. 021006.

[43] Thomas Harty. *Atomic Physics*. Version 1.0.2. URL: https://github.com/OxfordIonTrapGroup/atomic_physics/.

[44] R Srinivas. "Laser-free trapped-ion quantum logic with a radiofrequency magnetic field gradient". PhD thesis. 2020.

[45] A. C. Wilson et al. "A 750 mW, continuous-wave, solid-state laser source at 313 nm for cooling and manipulating trapped $^{9}Be^{+}$ ions". In: *Applied Physics B* 105.4 (Dec. 2011), pp. 741–748.

[46] Hsiang-Yu Lo et al. "All-solid-state continuous-wave laser systems for ionization, cooling and quantum state manipulation of beryllium ions". In: *Applied Physics B* 114.1 (Jan. 2014), pp. 17–25.

[47] Zhiqiang Zhang et al. "Branching fractions for $P_{3/2}$ decays in Ba^{+}". In: *Physical Review A* 101.6 (June 2020), p. 062515.

[48] Justin Christensen. "High-fidelity operation of a radioactive trapped-ion qubit, $^{133}Ba^{+}$". PhD thesis. UCLA, 2020.

[49] Brendan M. White et al. "Isotope-selective laser ablation ion-trap loading of $^{137}Ba^{+}$ using a $BaCl_2$ target". In: *Physical Review A* 105.3 (Mar. 2022), p. 033102.

[50] X. Shi et al. "Ablation loading of barium ions into a surface-electrode trap". In: *Applied Physics Letters* 122.26 (June 2023), p. 264002.

[51] R. Blatt and G. Werth. "Precision determination of the ground-state hyperfine splitting in $^{137}Ba^{+}$ using the ion-storage technique". In: *Physical Review A* 25.3 (Mar. 1982), pp. 1476–1482.

[52] Roger E. Silverans et al. "Hyperfine structure of the $5d^2D_J$ states in the alkaline-earth Ba ion by fast-ion-beam laser-rf spectroscopy". In: *Physical Review A* 33.3 (Mar. 1986), pp. 2117–2120.

[53] D. Szwer. "High fidelity readout and protection of a $^{43}Ca^+$ trapped ion qubit". PhD thesis. Oxford University, UK, 2009.

[54] 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.

[55] 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.

[56] D. J. Wineland and Wayne M. Itano. "Laser cooling of atoms". In: *Physical Review A* 20.4 (Oct. 1979), pp. 1521–1540.

[57] C. J. Foot. *Atomic Physics*. Oxford Master Series in Physics. Oxford University Press, Nov. 2004.

[58] H. G. Dehmelt. "Mono-Ion Oscillator for Ultimate Resolution Laser Spectroscopy". In: *Laser Spectroscopy V*. Ed. by A. Robert W. McKellar, Takeshi Oka, and Boris P. Stoicheff. Berlin, Heidelberg: Springer, 1981, pp. 353–359.

[59] 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 (May 1998), p. 259.

[60] D. Leibfried et al. "Quantum dynamics of single trapped ions". In: *Reviews of Modern Physics* 75.1 (Mar. 2003), pp. 281–324.

[61] Laird Nicholas Egan. "Scaling Quantum Computers with Long Chains of Trapped Ions". Ph.D. United States – Maryland: University of Maryland, College Park, 2021.

[62] D. Kielpinski, C. Monroe, and D. J. Wineland. "Architecture for a large-scale ion-trap quantum computer". In: *Nature* 417.6890 (June 2002), pp. 709–711.

[63] L. Stephenson. "Entanglement between nodes of a quantum network". PhD thesis. University of Oxford, 2019.

[64] I. V. Inlek et al. "Multispecies Trapped-Ion Node for Quantum Networking". In: *Physical Review Letters* 118.25 (June 2017), p. 250502.

[65] C. Monroe et al. "Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects". In: *Physical Review A* 89.2 (Feb. 2014), p. 022317.

[66] D. P. Nadlinger et al. "Experimental quantum key distribution certified by Bell's theorem". In: *Nature* 607.7920 (July 2022), pp. 682–686.

[67] P. Drmota et al. "Verifiable Blind Quantum Computing with Trapped Ions and Single Photons". In: *Physical Review Letters* 132.15 (Apr. 2024), p. 150604.

[68] H. Hahn et al. "Multilayer ion trap with three-dimensional microwave circuitry for scalable quantum logic applications". In: *Applied Physics B* 125.8 (July 2019), p. 154.

[69] James E. Tarlton, Richard C. Thompson, and David M. Lucas. "Surface-electrode ion trap design for near-field microwave quantum gates". In: *Applied Physics B* 129.6 (May 2023), p. 89.

[70] Kaushal Choonee, Guido Wilpers, and Alastair G. Sinclair. "Silicon microfabricated linear segmented ion traps for quantum technologies". In: *2017 19th International Conference on Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS)*. June 2017, pp. 615–618.

[71] Patrick See et al. "Fabrication of a Monolithic Array of Three Dimensional Si-based Ion Traps". In: *Journal of Microelectromechanical Systems* 22.5 (Oct. 2013), pp. 1180–1189.

[72] Guido Wilpers et al. "A monolithic array of three-dimensional ion traps fabricated with conventional semiconductor technology". In: *Nature Nanotechnology* 7.9 (Sept. 2012), pp. 572–576.

[73] M Brownnutt et al. "Monolithic microfabricated ion trap chip design for scaleable quantum processors". In: *New Journal of Physics* 8.10 (Oct. 2006), pp. 232–232.

[74] Guido Wilpers et al. "A compact UHV package for microfabricated ion-trap arrays with direct electronic air-side access". In: *Applied Physics B* 111.1 (Apr. 2013), pp. 21–28.

[75] A. Burrell. "High fidelity readout of trapped ion qubits". PhD thesis. Oxford University, UK, 2010.

[76] Stephan Gulde. "Experimental Realization of Quantum Gates and the Deutsch-Jozsa Algorithm with Trapped $^{40}Ca^{+}$ Ions". PhD thesis. University of Innsbruck, 2003.

[77] H.C. Nägerl et al. "Ion strings for quantum gates". In: *Applied Physics B* 66.5 (May 1998), pp. 603–608.

[78] A. D. Leu et al. *Fast, high-fidelity addressed single-qubit gates using efficient composite pulse sequences*. arXiv:2305.06725 [quant-ph]. May 2023.

[79] R. Srinivas et al. "Coherent Control of Trapped-Ion Qubits with Localized Electric Fields". In: *Physical Review Letters* 131.2 (July 2023), p. 020601.

[80] D. Leibfried. "Individual addressing and state readout of trapped ions utilizing rf micromotion". In: *Physical Review A* 60.5 (Nov. 1999), R3335–R3338.

[81] U. Warring et al. "Individual-Ion Addressing with Microwave Field Gradients". In: *Physical Review Letters* 110.17 (Apr. 2013), p. 173002.

[82] Shannon X. Wang et al. "Individual addressing of ions using magnetic field gradients in a surface-electrode ion trap". In: *Applied Physics Letters* 94.9 (Mar. 2009), p. 094103.

[83] Christopher M. Seck et al. "Single-ion addressing via trap potential modulation in global optical fields". In: *New Journal of Physics* 22.5 (May 2020), p. 053024.

[84] S. Crain et al. "Individual addressing of trapped $^{171}Yb^{+}$ ion qubits using a microelectromechanical systems-based beam steering system". In: *Applied Physics Letters* 105.18 (Nov. 2014), p. 181115.

[85] Chung-You 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), pp. 1–8.

[86] Ye Wang et al. “High-Fidelity Two-Qubit Gates Using a Microelectromechanical-System-Based Beam Steering System for Individual Qubit Addressing”. In: *Physical Review Letters* 125.15 (Oct. 2020), p. 150505.

[87] I. Pogorelov et al. “Compact Ion-Trap Quantum Computing Demonstrator”. In: *PRX Quantum* 2.2 (June 2021), p. 020343.

[88] Ali Binai-Motlagh et al. *A guided light system for agile individual addressing of* Ba^+ *qubits with* 10^{-4} *level intensity crosstalk*. arXiv:2302.14711 [physics, physics:quant-ph]. Feb. 2023.

[89] Flavia Timpu et al. “Laser-Written Waveguide Array Optimized for Individual Control of Trapped Ion Qubits in a Chain”. In: *2022 European Conference on Optical Communication (ECOC)*. Sept. 2022, pp. 1–4.

[90] Bangshan Sun et al. “On-chip beam rotators, adiabatic mode converters, and waveplates through low-loss waveguides with variable cross-sections”. In: *Light: Science & Applications* 11.1 (July 2022), p. 214.

[91] M. Cetina et al. “Control of Transverse Motion for Quantum Gates on Individually Addressed Atomic Qubits”. In: *PRX Quantum* 3.1 (Mar. 2022), p. 010334.

[92] Adam D. West et al. “Tunable transverse spin–motion coupling for quantum information processing”. In: *Quantum Science and Technology* 6.2 (Jan. 2021), p. 024003.

[93] Shelby Kimmel, Guang Hao Low, and Theodore J. Yoder. “Robust calibration of a universal single-qubit gate set via robust phase estimation”. In: *Physical Review A* 92.6 (Dec. 2015), p. 062315.

[94] Kenneth Rudinger et al. “Experimental Demonstration of a Cheap and Accurate Phase Estimation”. In: *Physical Review Letters* 118.19 (May 2017), p. 190502.

[95] E. L. Hahn. “Spin Echoes”. In: *Physical Review* 80.4 (Nov. 1950), pp. 580–594.

[96] F. Schmidt-Kaler et al. “How to realize a universal quantum gate with trapped ions”. In: *Applied Physics B* 77.8 (Dec. 2003), pp. 789–796.

[97] Lorenza Viola, Emanuel Knill, and Seth Lloyd. “Dynamical Decoupling of Open Quantum Systems”. In: *Physical Review Letters* 82.12 (Mar. 1999), pp. 2417–2421.

[98] Kenneth R. Brown, Aram W. Harrow, and Isaac L. Chuang. “Arbitrarily accurate composite pulse sequences”. In: *Physical Review A* 70.5 (Nov. 2004), p. 052318.

[99] Alexandre M. Souza, Gonzalo A. Álvarez, and Dieter Suter. “Robust Dynamical Decoupling for Quantum Computing and Quantum Memory”. In: *Physical Review Letters* 106.24 (June 2011), p. 240501.

[100] Nathan Wiebe and Chris Granade. “Efficient Bayesian Phase Estimation”. In: *Physical Review Letters* 117.1 (June 2016), p. 010503.

[101] D. W. Berry et al. “How to perform the most accurate possible phase measurements”. In: *Physical Review A* 80.5 (Nov. 2009), p. 052114.

[102] Cristian Bonato and Dominic W. Berry. “Adaptive tracking of a time-varying field with a quantum sensor”. In: *Physical Review A* 95.5 (May 2017), p. 052348.

[103] F. Martínez-García, D. Vodola, and M. Müller. “Adaptive Bayesian phase estimation for quantum error correcting codes”. In: *New Journal of Physics* 21.12 (Dec. 2019), p. 123027.

[104] Jeremy Flannery et al. “Optical Crosstalk Mitigation for Individual Addressing in a Cryogenic Ion Trap”. In: *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. Sept. 2022, pp. 816–817.

[105] A. H. Myerson et al. “High-Fidelity Readout of Trapped-Ion Qubits”. In: *Physical Review Letters* 100.20 (May 2008), p. 200502.

[106] S. Olmschenk et al. “Manipulation and detection of a trapped Yb^+ hyperfine qubit”. In: *Physical Review A* 76.5 (Nov. 2007), p. 052314.

[107] Anthony Ransford et al. “Weak dissipation for high-fidelity qubit-state preparation and measurement”. In: *Physical Review A* 104.6 (Dec. 2021), p. L060402.

[108] Fangzhao Alex 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.

[109] Yue 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.

[110] Aleksander Kubica et al. *Erasure qubits: Overcoming the T_1 limit in superconducting circuits*. arXiv:2208.05461 [quant-ph]. Aug. 2022.

[111] Mingyu Kang, Wesley C. Campbell, and Kenneth R. Brown. “Quantum Error Correction with Metastable States of Trapped Ions Using Erasure Conversion”. In: *PRX Quantum* 4.2 (June 2023), p. 020358.

[112] Boris Mihailov Varbanov et al. “Leakage detection for a transmon-based surface code”. In: *npj Quantum Information* 6.1 (Dec. 2020), pp. 1–13.

[113] J. E. Johnson et al. “Heralded State Preparation in a Superconducting Qubit”. In: *Physical Review Letters* 109.5 (Aug. 2012), p. 050506.

[114] Samuel R. Vizvary et al. *Eliminating qubit type cross-talk in the omg protocol*. arXiv:2310.10905 [physics, physics:quant-ph]. Oct. 2023.

[115] F Schmidt-Kaler et al. “The coherence of qubits based on single Ca^+ ions”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 36.3 (Feb. 2003), pp. 623–636.

[116] Iris Cong et al. “Hardware-Efficient, Fault-Tolerant Quantum Computation with Rydberg Atoms”. In: *Physical Review X* 12.2 (June 2022), p. 021049.

[117] S. Schneider and G. J. Milburn. “Decoherence in ion traps due to laser intensity and phase fluctuations”. In: *Physical Review A* 57.5 (May 1998). arXiv:quant-ph/9710044, pp. 3748–3752.

[118] Ludwig Krinner et al. *A low phase noise cavity transmission self-injection locked laser system for atomic physics experiments*. arXiv:2311.03461 [physics, physics:quant-ph]. Nov. 2023.

[119] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Dec. 2010.

[120] Peter Drmota. "Blind Quantum Computing with Trapped Ions and Single Photons". PhD thesis. University of Oxford, 2023.

[121] I. D. Moore et al. "Photon scattering errors during stimulated Raman transitions in trapped-ion qubits". In: *Physical Review A* 107.3 (Mar. 2023). arXiv:2211.00744 [physics, physics:quant-ph], p. 032413.

[122] Matthew J. Boguslawski et al. "Errors in stimulated-Raman-induced logic gates in $^{133}Ba^{+}$". In: *Physical Review Letters* 131.6 (Aug. 2023). arXiv:2212.02608 [physics, physics:quant-ph], p. 063001.

[123] Daniel A Steck. *Quantum and Atom Optics*.

[124] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. "Measurement of ion motional heating rates over a range of trap frequencies and temperatures". In: *Physical Review A* 91.4 (Apr. 2015), p. 041402.

[125] Nitzan Akerman et al. "Universal gate-set for trapped-ion qubits using a narrow linewidth diode laser". In: *New Journal of Physics* 17.11 (Nov. 2015), p. 113060.

[126] Lukas Gerster et al. "Experimental Bayesian Calibration of Trapped-Ion Entangling Operations". In: *PRX Quantum* 3.2 (June 2022), p. 020350.

Appendices

A

Barium clock qubits

Contents

A.1 Metastable level clock qubits 169
A.2 Ground level clock qubits 171

A.1 Metastable level clock qubits

In this section, we provide a list of the metastable level clock qubits in $^{137}Ba^{+}$. There are 24 states in the metastable level and thus many possibilities for clock qubits at various magnetic fields. Here, we only consider clock qubits that can be used for quantum information processing. Therefore, when searching for clock qubits we apply the following criteria:

- The qubit transition must be a dipole-allowed transition, i.e. $|\Delta m_F| \leq 1$.
- The qubit frequency must be much larger than typical secular frequencies encountered in ion traps. Thus, here we set the cut-off at 20 MHz.
- The field-insensitive point must be sufficiently large such that fluctuations in the external magnetic field are much smaller than the quantisation field itself. Here, we set the cut-off at 1 G.
- The field must be sufficiently small so it can be generated using existing techniques for trapped-ion set-ups. We set this cut-off at 1 T.

Under these constraints, a list of the metastable level clock qubits is provided in Table A.1.

Qubit states	Field-insensitive point (G)	Qubit frequency (MHz)	Second-order sensitivity (kHz/G^2)
$\|1,-1\rangle \leftrightarrow \|2,-1\rangle$	15.410	59.674	104.327
$\|1,-1\rangle \leftrightarrow \|3,-2\rangle$	14.670	120.511	150.089
$\|1,-1\rangle \leftrightarrow \|3,-1\rangle$	19.914	101.614	184.608
$\|2,-2\rangle \leftrightarrow \|3,-3\rangle$	10.403	60.113	61.754
$\|2,-2\rangle \leftrightarrow \|3,-2\rangle$	20.048	42.285	107.587
$\|2,-2\rangle \leftrightarrow \|3,-1\rangle$	25.325	19.220	137.805
$\|2,-2\rangle \leftrightarrow \|4,-1\rangle$	8.721	59.696	92.626
$\|2,-1\rangle \leftrightarrow \|3,-2\rangle$	12.704	60.733	36.504
$\|2,-1\rangle \leftrightarrow \|3,-1\rangle$	23.531	40.185	121.740
$\|2,-1\rangle \leftrightarrow \|4,0\rangle$	8.588	61.956	23.727
$\|2,0\rangle \leftrightarrow \|3,-1\rangle$	1.693	63.054	27.113
$\|2,0\rangle \leftrightarrow \|3,0\rangle$	18.483	50.214	88.117
$\|2,+1\rangle \leftrightarrow \|3,+1\rangle$	12.287	58.133	68.479
$\|2,+2\rangle \leftrightarrow \|3,+2\rangle$	3.333	62.696	64.311
$\|3,-3\rangle \leftrightarrow \|3,-2\rangle$	55.077	27.016	2.227
$\|3,-2\rangle \leftrightarrow \|3,-1\rangle$	39.098	27.506	19.598
$\|3,-1\rangle \leftrightarrow \|4,0\rangle$	30.556	26.111	53.697
$\|3,0\rangle \leftrightarrow \|4,+1\rangle$	48.375	27.578	6.320

Table A.1: A list of the metastable level clock qubits in $^{137}Ba^{+}$. The states are labelled as $|F, m_F\rangle$.

A.2 Ground level clock qubits

A list of the available ground-level clock qubits is provided in Table A.2. We use the same criteria as in Section A.1.

Qubit states	Field-insensitive point (G)	Qubit frequency (MHz)	Second-order sensitivity (kHz/G^2)
$\|1,-1\rangle \leftrightarrow \|2,-1\rangle$	1433.662	6960.888	1.129
$\|1,-1\rangle \leftrightarrow \|2,0\rangle$	768.797	7763.496	0.949
$\|1,0\rangle \leftrightarrow \|2,-1\rangle$	767.793	7764.228	0.949

Table A.2: A list of the ground level clock qubits in $^{137}Ba^{+}$. The states are labelled as $|F, m_F\rangle$.

List of Figures

2.1 Fundamental concept of our architecture [6]. We encode qubits within two level manifolds of the ion: the ground level (G) and the metastable level (M). The laser light used to perform cooling, readout, and qubit reset (dissipative light) does not couple to the M level. Thus, we use the M level qubits as our computational qubits and the G level qubits as our ancilla qubits. 11

3.1 Gross level structure of most ions used for ion trap quantum computing. The ground level is $S_{1/2}$. There are two short-lived levels, $P_{1/2}$ and $P_{3/2}$, and two long-lived levels, $S_{1/2}$ and $D_{5/2}$ (lifetime of ~ 1 s). The D-levels do not exist in beryllium and magnesium. In ytterbium, they are very short-lived, and the $F_{7/2}$ level is used in cases where a long-lived level is necessary in addition to the ground level. 15

3.2 Gross level structure of a barium ion. All wavelengths we need to work with are in the visible range, with the exception of the 1762 nm laser, which is in the near-IR range. 17

3.3 Two-step photoionisation of neutral barium atoms. The two valence electrons are initially in the state $6s^2$. A resonant 553 nm laser excites one of the electrons to the $6p$ state. This is followed by a 405 nm laser that brings the p-electron to the continuum. 19

3.4 Sidebands used during operations using the ion's fluorescence cycle [51, 52]. The 493 nm laser carrier is tuned on resonance with the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. We apply two RF tones to an EOM, at 8.037 GHz and 6.56 GHz driving the $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ and the $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 1\rangle$ transitions, respectively. Due to sideband mixing in the EOM, these tones also generate a sideband at 1.48 GHz, which drives the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. The 650 nm laser carrier is set halfway in between transitions from $|D_{3/2}, F = 0\rangle$ and $|D_{3/2}, F = 1\rangle$ to $|P_{1/2}, F = 1\rangle$. We apply a single RF tone to an EOM at 430 MHz, which generates sidebands driving the $|D_{3/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 1\rangle$ and $|D_{3/2}, F = 3\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. We use ~ 1000 saturation intensities in the 650 nm beam such that we power-broaden the transition to address all states within $D_{3/2}$ 21

3.5 Preparation into a well-known state within $S_{1/2}$ using optical pumping. We turn on the 493 nm laser, turning on only the sidebands resonant with the $|S_{1/2}, F = 2\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ and $|S_{1/2}, F = 1\rangle \leftrightarrow |P_{1/2}, F = 2\rangle$ transitions. Since the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ transition is forbidden, this process results in the majority of the population being pumped in the $|S_{1/2}, F = 2, m_F = 0\rangle$ state. 23

3.6 A rate equations simulation of the state preparation error in the $|S_{1/2}, F = 2, m_F = 0\rangle$ state as a function of the applied magnetic field. In this simulation, we use $1I_{\text{sat}}$ of 493 nm light and $1000I_{\text{sat}}$ of 650 nm light. We assume pure π-polarisation of the 493 nm beam. The 493 nm sidebands used are as shown in Fig. 3.5, and the 650 nm sidebands used are as shown in Fig. 3.4. We assume we have equal powers in each sideband (for both the 493 nm and the 650 nm beams). The state preparation duration used is 100 ms: this is much longer than necessary in practice; we used this duration for the simulation only to ensure that we have reached a steady state. As the magnetic field increases, the matrix element of the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 2, m_F = 0\rangle$ increases as well, and we enter the intermediate-field regime [57]. Hence, F is no longer a good quantum number, resulting in changes in the transition selection rules. The increase of this matrix element with magnetic field poses a fundamental limit on the fidelity of the state preparation scheme described in Section 3.3.3. The minimum state preparation error of $\sim 0.1\%$ is due to off-resonant excitation of the $|S_{1/2}, F = 2, m_F = 0\rangle \leftrightarrow |P_{1/2}, F = 1, m_F = 0\rangle$ transition. 24

3.7 Hyperfine splitting of the metastable level of $^{137}Ba^+$ in the absence of an external magnetic field [52]. 26

3.8 Evolution of the metastable level state energies as a function of the external magnetic field. The frequencies provided are relative to the $D_{5/2}$ state when only fine structure is considered. 27

3.9 Evolution of the frequency splitting between the states $|F = 2, m_F = +2\rangle$ and $|F = 3, m_F = +2\rangle$ in the $D_{5/2}$ level of $^{137}Ba^+$. The frequency splitting reaches its minimum at $B \approx 3.333\,\text{G}$. As can be seen from the blue curve, this point also corresponds to $df/dB = 0$, i.e. at this magnetic field, these states form a clock qubit. 28

3.10 Hyperfine splitting of the ground $S_{1/2}$ level of $^{137}Ba^+$ in the absence of an external magnetic field [51]. 28

3.11 Evolution of the energies of the states within the ground $S_{1/2}$ level in $^{137}Ba^+$ as the applied quantisation field is increased. The frequencies provided are relative to the $S_{1/2}$ state when only fine structure is considered. 29

3.12 Evolution of the frequency splitting and the first-order magnetic field sensitivity of the $|F = 2, m_F = 0\rangle \leftrightarrow |F = 1, m_F = 0\rangle$ in the $S_{1/2}$ level. This pair of states forms a clock qubit at zero quantisation field. As the field is increased, this is no longer a perfect clock qubit and its first-order magnetic field sensitivity becomes non-zero. However, at the fields we want to work with (constrained by the metastable level), this is the ground level transition with the lowest field sensitivity. 29

3.13 Raman transitions used to drive the two qubit types in $^{137}Ba^+$ . . 30

4.1 The trap vacuum chamber and pumping assembly mounted on the optics table. 36

4.2 An overview of the trap layout. (a) A photo of the trap chip. (b) A cross-section of the trap layout in a plane perpendicular to the ion chain. The distance between the trap electrodes is 350 µm resulting in an ion-electrode distance of ≈ 247.5 µm. (c) Trap electrode map [same view as in (a)]. The ion chain is at the centre of the slot, parallel to the RF electrode. Zones 1 and 11 (brown) are end electrodes that could be used as additional degrees of freedom in generating electric field compensation voltagesets. Zone 10 (green) is the loading zone. Zones $2-8$ (red) are the trap operation zones. Zone 9 (purple) is the 'transfer' zone used for ion shuttling between the loading zones and the trap operation zones. Unless otherwise specified, during experiments the ions are confined in Zone 5. Zone 10 is only used during loading, and zones 3, 4, and 6 are used as overflow zones for ion merging, as described in Section 4.1.4. . . . 37

4.3 A photo of the PCB used to deliver signals to the trap chip. . . . 38

4.4 RC filter used to filter the input DC voltage to each electrode before delivery to the trap. The component values were chosen to minimise high-frequency noise, thus minimising the trap heating rate while facilitating shuttling operations. 38

4.5 LC resonator circuit used to convert the input low-voltage high-input current RF signal into a high-voltage low-current signal before delivery to the trap. The resonance frequency of this circuit, when connected to the trap chip, is 17.685 MHz. A small amount of the signal delivered to the trap is picked off using the 820 pF capacitor and is used to monitor the applied voltage. The measured ratio between the RF amplitude at the trap and the pickoff signal is ≈ 540. The discrepancy between the capacitor values shown in the circuit above and the measured pickoff fraction is most likely due to the capacitance of the pickoff cable to ground. 39

4.6 Isotope-selective loading sequence. (i) We start with a trapping potential at Zone 5 of the trap that contains a certain number of ions (one in the example above). (ii) When more ions need to be loaded, the ions that are currently trapped are transported to Zone 3 for storage, and a trapping potential is created in Zone 10 (the loading zone). (iii) After a chosen loading duration (~ 10 ms in our experiment), the trap in Zone 10 is transported to Zone 5, where it is cooled using the 493 nm and the 650 nm beams at their idle cooling parameters, as defined in Section 3.3.2. The ion is then imaged using a camera, and the resulting fluorescence counts are recorded. If the fluorescence counts are below a certain threshold (corresponding to either the absence of an ion or an ion of the wrong isotope), the ion is transported back to Zone 10 and then kicked out of the trap by applying DC voltages on the electrodes in Zones 9 and 10, resulting in an anticonfining potential in these zones. (iv) If the fluorescence counts are above this threshold, indicating that the ion is a $^{137}Ba^{+}$ ion, the ions stored in Zone 3 are transported to Zone 4, and the newly trapped ion is transported from Zone 5 to Zone 6. (v) Finally, a merging waveform is used to merge the ions in Zones 4 and 6 into Zone 5. This sequence is repeated as many times as necessary until the required number of ions is reached. 41

4.7 (a) Ion imaging system. Light from the ions is collected using an $NA = 0.5$ commercial microscope objective. The light then passes through a narrow line filter and is focused on an sCMOS camera using an $f = 35$ mm lens. The measured magnification of the imaging system is ≈ 6.5. (b) An image of two trapped $^{137}Ba^{+}$ ions. To measure the ion fluorescence, we sum the counts from a 3×3 region around the centre of the ion image. 43

4.8 Raman laser system. The source 1064 nm laser is amplified and split up into two arms. One arm (the carrier) is amplified and frequency-doubled to produce ~ 10 W of 532 nm light. The second arm (the offset) is frequency-modulated using an EOM driven at 4 GHz. It then passes through a cavity and the laser is frequency-locked to the $+1^{\text{st}}$ sideband of the cavity. The output of the cavity is then amplified again and frequency-doubled to generate ~ 10 W of 532 nm light detuned by 8 GHz relative to the carrier. 50

4.9 AOM network used to tune the frequencies of the Raman lasers on resonance with the ground and the metastable level qubits. The carrier and the offset lasers are denoted in Fig. 4.8. The frequencies shown in the diagram above are the centre AOM frequencies generating roughly the right frequency splitting. The double-pass AOMs in the ground and the metastable level beam paths are used for fine-tuning the beam frequencies on resonance with the desired transition. The ground beam, together with either the co- or counter-propagating beam, is used to drive ground-level transitions. The metastable beam, together with either the co- or counter-propagating beam, is used to drive metastable-level transitions. The ground, metastable, and co-propagating beams are tightly focused to ~ 1 µm spot size onto the ions using our individual addressing system (see Chapter 5) and therefore require little power. Hence, we use TeO_2 AOMs for these beam paths. The counter-propagating beam is global for the qubit register and, therefore, has a much larger spot size (~ 10 µm). Hence, it requires much higher power, as we will see in Chapter 8. Therefore, we use a set of quartz AOMs to modulate its frequency. 51

4.10 Generic trap beam path. The polarisation of the beam is cleaned up using a PBS, which also turns any polarisation fluctuations into power fluctuations. A small amount of the beam power is sent to a photodetector (PD) used in a feedback loop for power stabilisation. A subsequent set of waveplates and lenses is used to set the beam polarisation and the spot size at the ion position. 52

4.11 Directions of incidence of the individual laser beams onto the ions. The ground, metastable, and co-propagating Raman beams overlap in practice but are drawn separately here for clarity. 53

5.1 Requirements for individual optical addressing in long chains of trapped ions. The spacing between the individual laser beams needs to match the ion separaion, typically $4 - 10$ µm. The spot size needs to be small enough to minimise the intensity at the neighbouring ions, but larger than the Fourier limit for the laser wavelength in use to avoid cross-talk due to aberrations in the optical system. 56

5.2 An outline of the fibre network used for individual addressing. The free-space 532 nm laser source is coupled to a series of fibre splitters. Each of the resulting channels is then connected to a fibre AOM that enables individual channel switching as well as individual phase and amplitude control. The outputs of the fibre AOMs are connected to a fibre VGA whose output is an array of fibre cores with an MFD of 3.6 µm and a pitch of 127 µm. To map between the output of the VGA and the ion pattern, we use an integrated photonic chip whose input matches the VGA output pattern and whose output matches the ion pattern (up to an overall scaling factor). 60

5.3 (a) Setup used to couple light from the VGA into the photonic chip. The VGA is placed on a 6-axis positioning stage (left) while the photonic chip is held stationary (platform on the right in front of the objective). The objective on the right is used to image the photonic chip output onto a camera (not shown) to evaluate the coupling quality. (b) A photo of the VGA glued to the photonic chip after the coupling is optimised. (c) A photo of the VGA and photonic chip assembly glued onto a stainless steel plate before integrating with the optical system in the ion trap setup. 62

5.4 (a) A photo of the photonic chip output after glueing to the VGA taken using a CMOS camera, and a profile taken by summing the camera counts along the vertical axis. The slight difference in the vertical position of the channels comes from non-uniform coupling into each of the channels, resulting in multi-mode propagation. (b) A zoomed-in view of one of the channels (channel 6 left-to-right). We observe no measurable cross-talk compared to the baseline in this measurement. 63

5.5 Telecentricity requirement for the optical system used to image the photonic chip output onto the ion chain. The ion chain is parallel to the x axis. The motion of the ions in the chain can be decomposed into three sets of normal modes along the x, y, and z axes. To ensure all 532 nm beams couple to only one set of motional modes, their $\vec{k}$-vectors should be parallel to each other and to the motional mode of interest (here along the z axis). . . . 64

5.6 Trapped-ion setup with the individual optical addressing system. Light from the ions at 493 nm is collimated using a Mitutoyo NA0.5 objective with an effective focal length (EFL) of 4 mm and refocused onto an sCMOS camera for fast spatially-resolved readout. The 532 nm light from the photonic chip is first collimated using a Mitutoyo objective with an EFL of 20 mm to minimise aberrations. The focal lengths and positions of the subsequent $f = 250$ mm and $f = 100$ mm lenses are chosen to achieve the required 2:1 scaling and satisfy the telecentricity requirement of the optical system. A dichroic mirror (DM) overlaps the 532 nm beam path with the 493 nm fluorescence beam path. This allows us to focus the 532 nm light onto the ions using the same NA0.5 objective we use for imaging them. A 3.5 mm N-BK7 glass plate is used to fine-tune the beam positions at the trap with nanometre precision. A pair of mirrors (not shown) placed after the glass plate provide the necessary degrees of freedom for larger adjustments. 65

5.7 Displacement of the beam centre at the trap as a function of the angular displacement of the glass plate in Fig. 5.6. The distribution of datapoints along the vertical axis is dominated by positional drifts in the optical system between consecutive measurements. . . 68

5.8 Robust phase estimation (RPE) sequence for estimating the phase introduced on the $|\rightarrow\rangle$ state as a result of the AC Stark shift introduced by a single 532 nm beam. The ion is prepared in the state $|0\rangle$ and the subsequent 1762 nm $\pi_x/2$ pulse prepares the state $|\rightarrow\rangle$. Then the 532 nm beam is turned on for a duration $\tau_{532} = 2^k\tau_0$ where τ_0 is the *base* duration defining the maximum observable AC stark shift of $\delta_{\text{ac, max}} = 1/\tau_0$. The final 1762 nm $\pi_\phi/2$ pulse defines the measurement basis. The measurement is performed for $\phi = 0, \pi/2, \pi, 3\pi/2$ to minimise the effect of polarising noise in one direction. 70

5.9 (a) A Knill composite pulse robust to amplitude and detuning errors. (b) A Knill dynamical decoupling (KDD) sequence block consisting of two Knill pulses offset in phase by $\pi/2$, with delays of length τ_d inserted between consecutive π pulses. Without loss of generality here we have set the initial phase $\phi = 0$. Depending on the choice of delay time and the total required probe duration, this block is repeated N_{KDD} times. 73

5.10 Sequence used to measure the AC Stark shift introcuced by a single 532 nm beam onto the $|0\rangle \leftrightarrow |1\rangle$ transition, combining a KDD sequence to extend the system coherence time and and RPE protocol for accurate phase estimation. 74

5.11 A measurement of the displacement of a single trapped ion along the x (axial) and y (radial) directions as denoted in Fig. 4.11 and linear fits in both directions. We measure an axial gradient of 69.7 nm/V/m. The small radial gradient (-1.1 nm/V/m) is caused by tilts in the imaging system and is not a displacement of the ion relative to the trap. 75

5.12 Dependence of the Rabi frequency (top) and the detuning (bottom) of the $|0\rangle \leftrightarrow |1\rangle$ transition as function of the ion's axial position. . 76

5.13 Intensity measurement of the 532 nm beams at the ion position obtained by shuttling a single ion across the trap axis in steps of 69.7 nm and measuring the AC Stark shift introduced on one of the quadrupole transitions. The data from each of the channels is normalised relative to the maximum intensity of that channel. (a) A measurement of all 8 channels. The two outliers in channel 7 (left-to-right) are caused by experiment instabilities, most likely jumps in the frequency of the 1762 nm laser. (b) A zoomed-in view of the third channel [left-to-right in (a)]. This is the channel we typically use when working with a single ion in the trap. In most cases, the errorbars are smaller than the marker size. 78

5.14 Estimated single-qubit error due to intensity cross-talk in the device accounting for contributions from all nearest-neighbour channels. The errorbars are smaller than the marker size. 79

6.1 SPAM protocol. (a) Minimal requirements on the quantum system for the implementation of the protocol. (b) The complete sequence of the SPAM protocol. The camera icons represent the *population detection* steps with corresponding results R_1-R_4. The steps shaded in light blue are not necessary if the qubit is encoded exclusively in B. The step U_q corresponds to the implementation of the desired quantum algorithm. (c) A truth table showing whether a shot is accepted or rejected based on the results of the population detection steps $R_1 - R_4$ in (b). 85

6.2 The pairs of states used as O, M, and G qubits in the demonstration of the SPAM protocol. The states in black are the qubit states. The states in grey are used during any additional population transfer pulses necessary for the implementation of the protocol. Population transfer between states within $S_{1/2}$ and states within $D_{5/2}$ is realised using coherent 1762 nm pulses, tuned on resonance with the desired transition. 88

6.3 Choice of camera ROI for readout. (a) A $n \times n$ square region centred around the ion. We considered $n = 1, 3, 5$ and found that we achieve the highest signal-to-noise ratio for $n = 3$. (b) A 'cross' region. This is the only size of the cross region we considered. A smaller size would be equivalent to $n = 1$ in (a). A larger size (extending by a pixel horizontally and vertically) would be worse compared to $n = 3$ in (a) since the corner pixels in the 3×3 square region will contain more information as the ion image is equivalent to a $2D$ Gaussian distribution. (c) Using pixel weights to try and improve the signal-to-noise ratio. In the example presented here, we sum the counts from the centre cross region and then add the counts from the corners divided by a factor of two. We found that changing the pixel weights had only a small ($< 10\%$) effect on the optical detection error. Hence, for the experiments in this chapter, we chose to work with the ROI shown in (a). 92

6.4 Optical detection error for discriminating between a bright and a dark state. We observe an error on the bright state of $3.2(28) \times 10^{-6}$ and an error on the dark state of $0.0(11) \times 10^{-6}$. The bright and the dark data contain $1,850,000$ shots each (before post-selection). An uncalibrated nominal camera dark count offset of 100 counts per pixel has been subtracted from each raw count value in this and following histograms, leading to occasional negative count readings. 95

6.5 Probability of decay of the metastable level as a function of the delay time between preparation in the metastable level and measurement. The solid line is an exponential fit to the data, $\epsilon_{\text{decay}} = 1 - e^{-t_d/\tau}$. The fit corresponds to a decay time of $\tau = 27.2(2)$ s. 98

6.6 SPAM of the metastable level qubit in $^{137}Ba^{+}$. The histograms show the number of counts from the detection pulse R_3 in Fig. 6.1(b) observed after preparation of the $|0\rangle$ (pink) state and the $|1\rangle$ (blue) state after implementation of the full protocol. The data for both the $|0\rangle$ and the $|1\rangle$ states contain $1,000,000$ shots (before post-selection). The error due to the decay of the metastable level during the detection duration is present only in the measurement of the dark ($|1\rangle$, blue) state. The measured SPAM error for the bright ($|0\rangle$, pink) state is $0.0(20) \times 10^{-6}$, and the measured SPAM error for the dark state is $1.1(7) \times 10^{-5}$, giving an average error of $5(4) \times 10^{-6}$. The uncertainties represent one Wilson interval. . . . 101

6.7 State preparation and measurement of the metastable qubit in $^{137}Ba^{+}$. (a) Reduction in SPAM error after each post-selection step. (b) Reduction in the SPAM data rate as a function of the post-selection step. If we perform no post-selection, the average SPAM error is 2.807(23)%. The post-selection steps R_0 and R_6 are used at the very start and very end of the protocol to ensure that the ion remained in the trap. We expect reasonably infrequent ion loss, thus they have a negligible effect on the data. The step R_1 checks for successful state preparation and the step R_2 checks for successful mapping of the qubit states to the metastable level and for decay of the metastable level. Finally, the steps R_3 and R_4 are used to check that at least one of the two qubit states result in a bright measurement and are hence plotted together. Thus, they ensure that the qubit states were successfully transferred to $S_{1/2}$ for measurement and that there was no leakage outside of the qubit manifold. We see that, while as a result of the protocol, the data rate is reduced by $< 10\%$, the SPAM error is reduced by four orders of magnitude. The differences between the two states arise from different fidelities of the 1762 nm pulses used to transfer the qubit states to and from $S_{1/2}$ (see Section 6.4.2). 102

6.8 Ground level qubit SPAM data. The SPAM error is $8(4) \times 10^{-6}$. (a) The number of counts from the detection pulse R_3. The data contain $1,000,000$ shots for each state. (b) Reduction in the SPAM error after each post-selection step. (c). Reduction in the SPAM data rate after each post-selection step. 104

6.9 Optical qubit SPAM data. The SPAM error is $7(4) \times 10^{-6}$. (a) The number of counts from the detection pulse R_3. The data contain $1,000,000$ shots for each state. (b) Reduction in the SPAM error after each post-selection step. (c). Reduction in the SPAM data rate after each post-selection step. 106

6.10 Measured SPAM errors for all three qubit types in $^{137}Ba^{+}$ — O: $7(4) \times 10^{-6}$, M: $5(4) \times 10^{-6}$, and G: $8(4) \times 10^{-6}$. The dashed line and the shaded region represent the expected error and its uncertainty based on the independent measurements of the optical detection error and metastable level decay probability in Section 6.3.3. The errors for all three qubit types agree within the error bars and also agree with the expected error from Section 6.3.3. . . 110

7.1 Evolution of the frequency splitting between the states $|F = 2, m_F = +2\rangle$ and $|F = 3, m_F = +2\rangle$ in the $D_{5/2}$ level of $^{137}Ba^{+}$. The field-insensitive point, at which $df/dB = 0$ is at $B_c \approx 3.333\,\mathrm{G}$. The corresponding clock qubit frequency is $f_q \approx 62.696$ MHz, and the second-order magnetic field sensitivity is $d^2f/dB^2 \approx 64.312\,\mathrm{kHz/G^2}$. 113

7.2 Magnetic field at the ion position as a function of the current applied to the coil. A linear fit to the data results in a gradient of $0.87418(25)\,\mathrm{G/A}$. The fitted magnetic field offset is $-0.4302(10)\,\mathrm{G}$. We expect that this offset is due to stray magnetic fields in the lab in the direction of the field produced by the 'main' coil. The errorbars are smaller than the marker size. 115

7.3 A detuning scan of the qubit transition obtained by applying a π-pulse to the qubit at varying RF frequencies. This scan was taken at a magnetic field of $2.0951(5)\,\mathrm{G}$. The qubit frequency is obtained by fitting a sinc^2 function to the data. The measured qubit frequency in this scan is 62.74703(11) MHz. 116

7.4 Metastable level qubit frequency as a function of the quantisation field. The blue line is obtained using numerical simulations of the level structure [43]. Using a quadratic fit to the data we infer a measured field-insensitive point of 3.33(5) G and a corresponding qubit frequency of 62.6970(16) MHz. These values agree with the simulated values within the errorbars. 117

7.5 (a) A simple Ramsey sequence used for measuring the qubit coherence time. To obtain the Ramsey contrast we perform two measurements with analysis phases of $\phi = 0.0$ and $\phi = \pi$, and the contrast is calculated by taking the difference between the two points. (b) Simulation of the observed Ramsey fringes for $\tau = 1\,\mathrm{s}$ in the following scenarios: a perfectly calibrated experiment in the absence of noise (red), magnetic field noise resulting in $T_2 = 1\,\mathrm{s}$ (blue), coherent qubit frequency offset of 0.2 Hz (yellow), and $\pi/2$ pulse duration error of 20% (green). During a two-point measurement, we only probe at $\phi = 0$ and $\phi = \pi$, therefore, all of these effects would result in a reduction of the measured contrast. 118

7.6 Coherence time measurement of the metastable level qubit using the Ramsey experiment shown in Fig. 7.5. The red curve shows the decay in contrast of the metastable level qubit measured by directly driving the qubit transition using an RF antenna placed near the trap. We use a decay model of the contrast as a function of Ramsey delay time of $C = e^{-\tau/T_2}$. The fit corresponds to a coherence time of $T_2 = 31(5)$ s. The green curve shows the decay in contrast of the metastable level qubit measured using the Raman lasers. The fit corresponds to a coherence time of $T_2 = 9.1(10)$ s. The shaded regions correspond to a 1σ uncertainty interval. . . . 122

7.7 Coherence time measurement of the metastable level qubit in the presence of the 493 nm and the 650 nm beams set to the parameters used during readout (blue data). The coherence time we measure is $T_2 = 21.9(21)$ s. The pink curve shows the measured qubit coherence time (31(5) s) in the absence of the dissipative beams – this is the same measurement as in Fig. 7.6 and is plotted here again for easier comparison. We see a small reduction in the coherence time due to the dissipative beams. We expect that this is due to intensity noise in these beams. The measured coherence time corresponds to an intensity noise level of $\approx 6\%$. The shaded regions represent a 1σ uncertainty interval. 124

8.1 Blue (BSB) and red (RSB) sideband tones applied for driving an MS gate. The qubit frequency is ω_0. The motional mode frequency is ω_m. The two tones are detuned by δ from either of the motional modes, that is, we apply two beatnotes with frequencies $\omega_0 + \omega_\mathrm{m} - \delta$ and $\omega_0 - \omega_\mathrm{m} + \delta$. 130

8.2 Beam path configuration of the counter-propagating beam used for implementing two-qubit gates (see also Section 4.3.4). We apply a static RF tone at 125 MHz to the first quartz AOM. The RF tone passes through a pulse shaper board to smooth the rising/falling edges of the pulse, effectively shaping the Rabi frequency experienced by the ions. This minimises off-resonant coupling to nearby transitions during the gate. The BSB and RSB tones are generated separately, combined using an RF power splitter and then applied to the second quartz AOM. The frequencies of the BSB and RSB tones are 125 MHz $\pm\,\omega_m \mp \delta$, respectively, where ω_m is the frequency of the mode we wish to couple to and δ is the MS gate detuning. 132
8.3 Population transfer within the $F = 2$ manifold of the $S_{1/2}$ level due to the intense counter-propagating Raman beam. The lines are numerical simulations of the 5-level system formed by the states in $F = 2$ in $S_{1/2}$. The parameters of the numerical simulation (frequency splitting and field Rabi frequency) are extracted by fitting the data for the $|F = 2, m_F = 0\rangle$ states, and are used to subsequently numerically simulate the evolution of the remaining states. 135
8.4 Simulation of the time evolution of each of the states in a five-level system acted upon with a DC electric field. The frequency splitting between states in these simulations, in arbitrary units, was chosen to be $\omega_0 = 2$. We see that as the coupling strength Λ becomes comparable to the frequency splitting, we are able to coherently transfer population between the $F = 2$ states in $S_{1/2}$ 136
8.5 G-G gate results. The measured contrast of the pariy scan is 0.864(19) and the populations at the gate duration are 0.951(12) resulting in a gate fidelity of 90.8(11)%. 145
8.6 M-M gate results. The measured contrast of the parity scan is 0.831(26) and the populations at the gate duration are 0.927(14) resulting in a gate fidelity of 87.9(15)%. 147
8.7 G-M gate implementation. We see that the gate duration scan is comparable to the M-M gate duration scan. This is expected since the duration of this gate is determined by the M-M gate duration, and, as we will see in Section 8.9, the gate quality is limited by the qubit coherence, common to both the M-M and the G-M gates. Thus, we expect the fidelity of this gate to be comparable to the M-M gate fidelity. 151
8.8 Coherence time of the metastable level qubit measured using the RF antenna in the presence of the counter-propagating Raman beam. The measured coherence time is 6.1(11) ms. 152
8.9 Coherence time measurement of the ground level qubit using the co-propagating Raman beams. The measured coherence time is 3.55(18) ms. 154

List of Tables

3.1 A table containing the most commonly used atom species for trapped-ion quantum information processing [6, 44–47]. The wavelengths listed correspond to the shortest wavelengths necessary for driving transitions between levels used for quantum information processing in the ions. They typically correspond to the $S_{1/2} \leftrightarrow P_{3/2}$ transitions. The two species with the longest lifetimes of the metastable level are barium and ytterbium. While the ytterbium metastable level lifetime is longer compared to barium, it requires work with UV lasers and several population transfer pulses to reach the metastable level from the ground level. Therefore, in this work we choose to work with barium ions. . . . 17

3.2 barium isotopes [48]. All of the isotopes listed in this table are naturally abundant except $^{133}Ba^{+}$, which is normally created by neutron bombardment of $^{132}Ba^{+}$. 19

6.1 A breakdown of the errors during each of the steps used in the SPAM sequences for all three qubit types. The first row gives the error of preparing into $|S_{1/2}, F = 2, m_F = 0\rangle$ using optical pumping. This measurement is performed by first performing optical pumping, then doing four consecutive π pulses from $|S_{1/2}, F = 2, m_F = 0\rangle$ to $|D_{1/2}, F = 1, m_F = -1\rangle$, $|D_{5/2}, F = 1, m_F = +1\rangle$, $|D_{5/2}, F = 2, m_F = -1\rangle$, and $|D_{5/2}, F = 2, m_F = +1\rangle$, then performing a detection step and measuring the probability of a bright outcome. The multiple 1762 nm pulses ensure that the error due to mapping to $D_{5/2}$ is $\sim \epsilon_{1762}^4$ where ϵ_{1762} is the error of a single pulse. The remaining rows give the errors on the π-pulses used to map between states in $S_{1/2}$ and $D_{5/2}$. The first state is the state in $S_{1/2}$ and the second state is the state in $D_{5/2}$. The measurements of the π-pulse errors are performed by first preparing the metastable level state (using post-selection) and then performing a mapping pulse from that state to $S_{1/2}$ and measuring the probability of a dark outcome. . . 108

6.2 A comparison of the measured and expected total fraction of rejected shots during the SPAM protocol. We see a good agreement between the two. Any discrepancies are attributed to drifts in system performance between measurements. 109

A.1 A list of the metastable level clock qubits in $^{137}Ba^{+}$. The states are labelled as $|F, m_F\rangle$. 170

A.2 A list of the ground level clock qubits in $^{137}Ba^{+}$. The states are labelled as $|F, m_F\rangle$. 171

List of Abbreviations

AOM . . . . . Acousto-optic modulator
DD . . . . . Dynamical decoupling
DDS . . . . . Direct digital synthesizer
EFL . . . . . Effective focal length
EOM . . . . . Electro-optic modulator
FTQC . . . . . Fault-tolerant quantum computing
KDD . . . . . Knill Dynamical Decoupling
NA . . . . . Numerical aperture
NISQ . . . . . Noisy intermediate-scale quantum (computers)
PCB . . . . . Printed circuit board
QCCD . . . . . Quantum charge-coupled device
QEC . . . . . Quantum error correction
QND . . . . . Quantum non-demolition (measurement)
ROI . . . . . Region of interest
RUS . . . . . Repeat until success
SPAM . . . . . State preparation and measurement
SPIM-WG . . Spherical phase-induced multiscan waveguide
STD . . . . . Standard deviation
VGA . . . . . V-groove array