

**Mixed-species quantum logic with trapped ions for gate
teleportation, metrology and high-fidelity indirect readout**

by

Stephen Dale Erickson

B.S., Physics and Mathematics, Brigham Young University, 2015

M.S., Physics, University of Colorado, 2018

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics
2021

Committee Members:

Dietrich Leibfried, Chair

Andrew Wilson

Adam Kaufman

Graeme Smith

Gordana Dukovic

Erickson, Stephen Dale (Ph.D., Physics)

Mixed-species quantum logic with trapped ions for gate teleportation, metrology and high-fidelity indirect readout

Thesis directed by Dr. Dietrich Leibfried

Trapped ions are a leading platform for quantum computation, communication, and metrology, in large part due to their promising pathways to scalability and the high fidelity quantum logic gates enabled by laser and microwave control. We use ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions to demonstrate how quantum logic operations between separate species of ion can enhance a broad spectrum of quantum technologies. First, we implement quantum gate teleportation between ${}^9\text{Be}^+$ ions in separate potential wells by means of a split entangled pair of ${}^{25}\text{Mg}^+$ ions, a key primitive for distributing quantum computation across multiple processors or for reducing transport latency in a large segmented trap array. Next, we show how mixed-species multi-ion operations can be used to combine the probe and readout in spectroscopy experiments while providing enhanced sensitivity in the number of spectroscopy ions due to the entangling nature of the operations. Finally, working with two species of ion allows for mid-experiment dissipation to the environment through one species without perturbing the other. This is often used for reinitializing the collective motion to near the ground state, however in the long run the internal ion states will also need to be measured and reinitialized during computation. We therefore propose and demonstrate a protocol for high-fidelity indirect measurement of trapped ion hyperfine qubits that is robust to spontaneous photon scattering and avoids the risk of decohering spectator qubits with stray resonant light.

Dedication

To my wife Brook.

Acknowledgements

Although the experiments described in this dissertation were performed on ions in a vacuum, it was not written in one. There are many people without which this work would not be possible, and to whom I owe a great deal of thanks. I am indebted to my formal advisors throughout this process, first Dave Wineland for offering me a position in the Ion Storage Group and guiding me through the early years of my graduate career, and later Didi Leibfried for taking over that mantle. Both allowed the freedom to pursue independent research interests in the lab, and always had open doors to share their depth and breadth of knowledge when that had me stumped. I am also grateful for the informal guidance given by Andrew Wilson, Daniel Slichter, and Dave Hume.

I would particularly like to thank all of the members of the Quantum II lab that I had the privilege of working with and learning from. Ting Rei Tan was the senior graduate student when I joined the lab and taught me a great deal about ion trapping as I bounced countless dumb ideas off of him. Postdocs Yong Wan and Daniel Kienzler then taught me just how much it takes to run a multi-species ion trapping experiment. As that generation moved on, I was able to work closely alongside a new generation in QII with fellow graduate student Jenny Wu and postdocs Panyu Hou and Dan Cole. More recently we have been joined by Giorgio Zarantonello and Ian Fogarty Florang, and I am excited to see where they take the lab in the future.

In Quantum II we benefited immensely by being embedded in the larger Ion Storage Group, all of whose members I owe a great deal of thanks, both past and present. They

are too many to name all of them individually, but everyone that has passed through the group contributed to the environment of excellence. I am also grateful for others at NIST that helped support this work, including the Time and Frequency Division admin staff who worked tirelessly to shield us from the bureaucracy, and our theory collaborators in the group of Scott Glancy and Manny Knill, including but not limited to Shawn Geller, Alex Kwiatkowski, Karl Mayer, and Hilma Vasconcelos.

Many others helped guide me towards my career in physics. I thank my parents John and Annette for always fostering an environment of learning. My high school physics teacher Scott Panik sparked my interest that got me started down this path. I am thankful to my undergraduate research advisor John Colton for taking a chance on me as an unproven sophomore, showing me how to do research and involving me in every step of the process. Finally, I am endlessly grateful to my wife Brook for supporting me every step along the way and loving me when the ions did not.

Contents

Chapter	
1 Introduction	1
1.1 Three Pillars of Quantum Technology	1
1.2 DiVincenzo Criteria	5
1.3 Thesis Organization	8
2 ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions	11
2.1 Magnetic Field Insensitive Transitions	12
2.2 ${}^2\text{S}_{1/2}$ Transition Frequencies	14
2.3 Excited States	16
2.4 Motion of Ions in the Trap	16
2.5 Optical Pumping and Doppler Cooling	19
2.6 Repumping	20
2.7 Microwave Transitions	22
2.8 State Preparation and Measurement	22
2.9 Stimulated Raman Transitions	24
2.10 Sideband Cooling	26
3 Apparatus	30
3.1 System Overview	30
3.2 Trap	33

3.3	Imaging System	37
3.4	Magnetic Field	39
3.5	Control System	40
3.6	${}^9\text{Be}^+$ lasers	42
3.6.1	${}^9\text{Be}^+$ resonant lasers	45
3.6.2	${}^9\text{Be}^+$ Raman lasers	48
3.7	${}^{25}\text{Mg}^+$ lasers	51
3.7.1	${}^{25}\text{Mg}^+$ resonant lasers	53
3.7.2	${}^{25}\text{Mg}^+$ Raman lasers	55
3.8	Microwaves	58
4	Quantum Logic Gates	60
4.1	State Definitions	60
4.2	Single-Qubit Gates	63
4.3	Multi-Qubit Gates	65
4.4	Mølmer-Sørensen Interaction	68
4.5	Mutually-Controlled Multiflip	72
4.6	Quantum Circuits	74
4.7	Composite Gates	75
4.8	Universal Gate Sets	78
5	Quantum Gate Teleportation	79
5.1	General Algorithm for Quantum Gate Teleportation	79
5.2	Experimental Demonstration	81
5.2.1	Shuttling Sequence	83
5.2.2	Experiment-Specific Circuit	88
5.2.3	Step Durations	90
5.3	Full Process Tomography	90

5.4	Error Budget	96
5.5	Scattered Resonant Light	99
5.6	Consistency Checks	105
5.7	Identifying Experimental Drifts	106
5.8	Validators and Dependencies	111
5.9	Outlook on Quantum Gate Teleportation	113
5.9.1	Utility within the QCCD Architecture	113
5.9.2	QCCD Electrode Structure	115
5.9.3	Roles for Each Species	115
6	Applications of the Mutually-Controlled Multiflip	117
6.1	Multiflip as a Tool for Spectroscopy	121
6.1.1	BM Crystal Multiflip Results	124
6.1.2	BMMB Crystal Multiflip Results	130
6.1.3	Multiflip Use Cases	133
6.1.4	Impact of Logic Ion Properties and Their Mitigation	134
6.1.5	Line-Pulling Effects and Their Cancellation	135
6.1.6	Multiflip Variants	137
6.2	Multiflip in Quantum Information Processing	141
6.2.1	Measure Qubit State Parity with the Multiflip	144
6.2.2	Protocol for the Repetition Quantum Error Correcting Code in a BMBMB Crystal	147
6.2.3	Protocol For Preparation and Readout of Hyperfine Qubits Without Resonant Light	150
7	High-Fidelity Indirect Measurement	158
7.1	Motivation	158
7.2	Repetitive Quantum Nondemolition Measurements	161

7.2.1	Bayesian Analysis	163
7.3	Extending From Optical to Hyperfine Qubits	165
7.4	Minimizing Spontaneous Emission	167
7.4.1	Containing Scattering within a Recoverable Subspace	169
7.4.2	Experimental Limitations	171
7.5	QLS Protocol	171
7.6	Data Filtering	174
7.7	Measurement Results	176
7.7.1	Test Data at 45 GHz Raman Detuning	178
7.7.2	High-Fidelity Measurement	182
7.8	Improvements to the Indirect Measurement Protocol	186
8	Conclusions and Outlook	188
	Bibliography	192

Tables

Table

2.1	$^9\text{Be}^+$ field insensitive transitions	13
2.2	$^{25}\text{Mg}^+$ field insensitive transitions	13
2.3	$^9\text{Be}^+$ transition frequencies	15
2.4	$^{25}\text{Mg}^+$ transition frequencies	15
3.1	$^9\text{Be}^+$ resonant AO switching	48
3.2	$^9\text{Be}^+$ Raman AO switching	51
3.3	$^{25}\text{Mg}^+$ resonant AO switching	55
3.4	$^{25}\text{Mg}^+$ Raman AO switching	58
5.1	Approximate durations of each step in our quantum gate teleportation demonstration	91
5.2	Leading Imperfections for the Remote CNOT Gate	97
5.3	Measured depumping errors	101
7.1	Be-Mg Crystal Normal Modes	168
7.2	Table of High-Fidelity Indirect Measurement Results	183

Figures

Figure

2.1	${}^9\text{Be}^+$ energy levels and control fields	17
2.2	${}^{25}\text{Mg}^+$ energy levels and control fields	18
2.3	Sideband Rabi rate dependence on n	28
3.1	Sketch of vacuum chamber and optical access	31
3.2	Schematic of Ion Trap	34
3.3	Axial micromotion sideband strength	36
3.4	Imagined Objective	38
3.5	Laser power stabilization	43
3.6	${}^9\text{Be}^+$ laser sources	44
3.7	${}^9\text{Be}^+$ resonant beam line	46
3.8	${}^9\text{Be}^+$ Raman beam line	49
3.9	${}^{25}\text{Mg}^+$ laser sources	52
3.10	${}^{25}\text{Mg}^+$ resonant beam line	54
3.11	${}^{25}\text{Mg}^+$ Raman beam line	56
4.1	Bloch Sphere	62
4.2	Mølmer-Sørensen interaction trajectory in phase space	69
4.3	Time evolution of Mølmer-Sørensen interaction on two qubits	71
4.4	Mutually-controlled multiflip on systems of 1-6 qubits	73

4.5	Composite Gates	77
5.1	Procedure for initializing BMMB crystal order	82
5.2	Shuttling sequence for quantum gate teleportation	84
5.3	Calibrating separation shim	86
5.4	Experiment-specific circuit for quantum gate teleportation	87
5.5	Teleported CNOT results in the Pauli transfer matrix representation	94
5.6	Bootstrapped Fidelity Histograms for the Remote CNOT Gate	95
5.7	Circuits for measuring depumping errors	100
5.8	Results of likelihood ratio test	105
5.9	Drifts in ${}^9\text{Be}^+$ copropagating carrier beams	108
5.10	Dead time in performing gates between distant qubits	114
6.1	Traditional Quantum Logic Spectroscopy	119
6.2	Multiflip Spectroscopy Energy Levels	122
6.3	Initializing BM crystal order	125
6.4	Multiflip Results in a BM Crystal	127
6.5	Multiflip Results in a BMMB Crystal	131
6.6	Multiflip resonance shift from loop closing errors	136
6.7	Multiflip error signal to cancel out shifts from loop closing errors	138
6.8	ZX multiflip	140
6.9	Multiflip for Multi-Qubit Parity Measurements	146
6.10	Circuit for Repetition Quantum Error Correction Code With a BMMB Crystal	149
6.11	Encoding Circuit for Repetition Code	151
6.12	Iterative Sideband Pumping Across ${}^9\text{Be}^+ {}^2\text{S}_{1/2}$ Manifold	153
6.13	Multiflip State Preparation Protocol	155
7.1	Fluorescence Detection Counts Histograms	162

7.2	$^{27}\text{Al}^+$ vs $^9\text{Be}^+$	166
7.3	Containing Spontaneous Emission with $^9\text{Be}^+$ Ground State Manifold	170
7.4	Circuit for QLS Protocol	172
7.5	Running Averages of Test QLS	177
7.6	QLS Test Data at 45 GHz Detuning	179
7.7	Histogram of QLS Rounds at 45 GHz Detuning for a 10^4 confidence ratio	181
7.8	Histogram of QLS Rounds at High Detuning	185

Chapter 1

Introduction

Quantum mechanics was first formulated nearly a century ago. Since then, its implications have enabled a number of revolutionary technologies across a wide range of industries. For example, quantum mechanics explains the electronic band structure in solids, which gives rise to the semiconductor transistor at the core of modern computing and electronics. It also led to the development of the laser, which forms the backbone of modern telecommunications through optical fibers, among many other applications.

1.1 Three Pillars of Quantum Technology

More recently, precise manipulation and control of individual quantum systems has become possible, leading to a second generation of quantum technology that directly exploits this control. This fledgling technology can be broken down into three main pillars:

- (1) quantum communication
- (2) quantum metrology
- (3) quantum simulation/computation

These applications directly harness the unique features of quantum mechanics, including superposition, entanglement, and interference.

Analogous to classical bits (i.e. 0's and 1's) in a conventional computer or digital communication protocol, quantum technology typically works on quantum bits (qubits) each

consisting of two fully distinguishable states, which we call $|0\rangle$ and $|1\rangle$. The $|\ \rangle$ surrounding the number serves to remind that we are working with qubit states, not their classical counterparts. What makes a qubit special is that it can be in a state that is made up of components of $|0\rangle$ and $|1\rangle$ *at the same time*. Such an undecided state is called a superposition of $|0\rangle$ and $|1\rangle$.

Two qubits can be in a superposition of states that cannot be described by the two qubits individually. For example, qubit one and two both being in superpositions of $|0\rangle$ and $|1\rangle$ is fundamentally different from them together being in a superposition of both $|0\rangle$ and both $|1\rangle$. In the latter, the first qubit to be measured (independent of which one that is) has a completely undetermined outcome, while a subsequent measurement of the other qubit is certain as soon as the first outcome is determined, and in this example will match it. This means that one cannot separate the measurement outcome on one of the qubits from the other outcome. Such a special state is an example of entangled states, and their existence has profound consequences on the fabric of our reality and for the computational power of qubits.

Finally, interference is analogous to waves in a pool either amplifying constructively or canceling destructively. Similarly, multiple paths between one quantum state and another can either add or subtract to make that transition more or less probable.

The first pillar, quantum communication, involves transferring quantum information between two distant parties and related applications. Naturally, photons form the basis for this technology given their ability to traverse great distances, although other systems may be used as intermediary quantum repeaters to extend their reach. Quantum measurements open up pathways for inherently secure communication, even in the presence of an eavesdropper, in a process called quantum key distribution (QKD) [Bennett and Brassard, 1984, Scarani et al., 2009]. QKD is one of the earliest applications that uses direct control of individual quantum systems, giving both the security of one-time pads and convenience of public key cryptography, and is already commercially available.

The second pillar, quantum metrology, makes precise measurements of physical parameters that influence quantum systems, or uses those systems as references for precise time keeping. Perhaps the most basic example of quantum metrology is atomic clocks that use a stable transition within an atom as a frequency reference. This typically begins by using a laser or microwave source to probe the transition, yielding results that depend on how the probing frequency compares to that of the transition. Such frequency standards have been the basis of the most stable clocks for decades, beginning with microwave transitions in cesium-133 and later with optical transitions [Ludlow et al., 2015]. Besides clock transitions that are selected to be as insensitive to their environment as possible, quantum systems also have other transitions that are exquisitely sensitive to their surroundings and can therefore serve as quantum sensors for many environmental effects like electric and magnetic fields. With trapped ions, the motion of the ion(s) within the trap can also serve as sensitive probe of the trapping environment [McCormick et al., 2019, Burd et al., 2019].

As the size of a quantum system grows, the amount of information needed to fully describe its state grows exponentially with particle number. For example, to fully write down an arbitrary 300 qubit system would require more pieces of information than there are particles in the universe. This makes quantum systems inherently difficult to simulate classically, and available classical computing resources limit simulations of such systems to about 50 qubits. Moreover, given the exponential scaling of this constraint, improvements in classical simulations are rapidly outpaced by adding just a few more qubits. Richard Feynman observed this difficulty and proposed instead using controllable quantum systems themselves to simulate the behavior of other quantum systems, giving birth to quantum simulation/computation as the third pillar [Feynman, 1982]. The general idea is that nature is governed by quantum mechanics, so if we have tunable interactions on a quantum system we can engineer it to emulate a system of interest that we cannot access directly (for example, all individual electron spins in a material like pharmaceuticals or high temperature superconductors). To date, quantum simulations on systems of up to hundreds of particles

have been performed, although typically for relatively simple physical systems [Georgescu et al., 2014].

Paul Benioff introduced the concept of quantum computation [Benioff, 1980]. Rather than operating on binary bits (0's and 1's) like conventional computing, quantum computing operates on qubits, which can harness the power of superposition, entanglement, and interference. Most physical implementations of qubits are encoded in effective two-level systems within a much richer landscape, for example by selecting two electron energy levels out of many in an atom. Quantum algorithms then proceed by applying a sequence of quantum logic gates, analogous to digital logic gates like AND, OR, and NOT in classical computing (Sec. 4.2-4.3). Quantum gates on qubits allow one to create superpositions, generate entanglement, and harness interference.

The first discovery of a quantum computational task that can be more efficiently performed on a quantum computer than a classical computer came from David Deutsch and Richard Jozsa in 1992 [Deutsch and Jozsa, 1992]. Though of little practical application, it served as a useful proof that such tasks exist, and a number of more useful algorithms soon followed. Interest in quantum computation took off in 1994 when Peter Shor discovered a quantum algorithm for prime factorization that gives an exponential improvement over the best known classical algorithms [Shor, 1994]. This has profound implications for the RSA and discrete logarithm based public key cryptography, which is the basis for most internet security, although the quantum computers that pose a serious threat to existing cryptosystems are probably still decades away, and significant research has since been done on alternative cryptosystems without known quantum lines of attack [Chen et al., 2016]. Regardless, it demonstrated that speedups from quantum computing are not limited to contrived problems like the Deutsch-Jozsa algorithm, but also to ones of practical use. Though only providing a more modest (but still very significant) quadratic improvement, Lov Grover shortly thereafter discovered a quantum algorithm for unsorted database search [Grover, 1996], which can be mapped to many other problems. While these three algorithms are the most famous,

a whole “zoo” of quantum algorithms exists [Jordan, 2020].

1.2 DiVincenzo Criteria

Naturally, theoretical proposals are many steps ahead of experiments executing these algorithms. Certain experimental ingredients are necessary to execute them, which were formalized by David DiVincenzo [DiVincenzo, 2000]:

- (1) A scalable physical system with well characterized qubits
- (2) The ability to initialize the state of the qubits to a simple fiducial state
- (3) Long relevant coherence times, much longer than the gate operation time
- (4) A universal set of quantum gates
- (5) A qubit-specific measurement capability

Many different physical systems have been used to demonstrate some or all of these fundamental building blocks, including nuclear magnetic resonance, quantum dots, photons, superconducting circuits, nitrogen-vacancy centers in diamond, trapped neutral atoms, and trapped ions, among others [Ladd et al., 2010].

Executing quantum algorithms with a reasonable chance of success is no easy task. For simple noise models, the total error for the entire algorithm grows exponentially in the error rate of the constituent operations. It is not obvious that overcoming this hurdle to realize the power of the quantum computing is feasible. In particular, quantum errors can occur on a continuous spectrum, and quantum measurements fundamentally alter the system. Fortunately, with the right choice of multi-qubit states and measurements, the latter can be used to discretize and subsequently correct for the former in a process called quantum error correction. Peter Shor made another foundational contribution to quantum computation by discovering the first quantum error correction scheme that encodes a single

“logical qubit” into nine physical qubits [Shor, 1995]. The key enabling feature is to not measure individual qubits, but instead parities of select subsets of qubits (i.e. whether an odd or even number of qubits are in $|0\rangle$). Then, with a carefully constructed code, the set of parity measurements informs whether certain errors occurred, and if so, how to correct them. For this to be beneficial, it is necessary that the constituent operations have errors below a certain threshold, such that the processes of checking for and correcting errors does not do more harm than good [Nielsen and Chuang, 2002]. Since Shor’s original error correcting code, many others have been invented, with varying degrees of protection, overhead, and qubit connectivity requirements. See Ref. [Devitt et al., 2013] for an excellent introduction to error correction as a whole and various promising codes.

The work presented in this thesis uses trapped ions, which are well equipped to meet each of the DiVincenzo criteria, as will be described below. We select two particular energy levels within the hyperfine ground states of ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions to serve as qubits, the details of which will be discussed in Chapter 2. Ion qubits are inherently simple, with well understood and characterized spectra. Every atom is identical to the next one. This makes adding more qubits to the register a matter of engineering traps that can store more ions while allowing them to controllably interact and without introducing significant noise. The two most prominent approaches to doing this, which ultimately may be combined in the long run, are the quantum charge-coupled device (QCCD) architecture used in this thesis, and modular traps joined through photonic networks [Moehring et al., 2007, Monroe et al., 2014]. Details of the former can be found in Refs. [Wineland et al., 1998] and [Kielpinski et al., 2002], with a brief introduction in Chapter 3.

Fulfilling the second criterion, the ability to initialize the state of the qubits to a simple fiducial state, is readily accomplished with standard atomic physics techniques. Simplified energy level diagrams for the ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions that we use can be seen in Figs. 2.1 and 2.2. Splitting into different magnetic quantum number m_F sublevels by an external magnetic field allows for optical pumping to a particular sublevel by applying a laser with the right

frequency and polarization. The fidelity, or success rate, of this process can be extremely high [Harty et al., 2014].

The third DiVincenzo criterion is for long relevant coherence times, much longer than the gate operation time. In simple terms, the coherence time is how long a qubit state remains unperturbed with the experimenter able to track it. Reduction in coherence is called decoherence. Trapped ions are extremely stable qubits, with decoherence introduced only by the surrounding environment and imperfection in operations used to manipulate them. Typically the environment perturbs qubits through magnetic field fluctuations from the external quantization field (Sec. 3.4) to which the qubit states are referenced. The hyperfine qubits that we use allow for transitions that are first-order insensitive to these fluctuations (Sec. 2.1), resulting in at least seconds long coherence times. Carefully stabilizing the magnetic field allowed for coherence times as long as 50 seconds in $^{43}\text{Ca}^+$ without any dynamical decoupling to refocus drifts. With dynamical decoupling, the coherence time of $^{171}\text{Yb}^+$ ions was extended to one hour [Wang et al., 2021]. In contrast, gate times are typically on the order of tens of microseconds, however time spent transporting and recooling ions in the QCCD architecture or generating remote entanglement in the photonic network architecture must also be considered.

Specific requirements for a universal gate set (criterion four), are outlined in Sec. 4.8. In short, an entangling two-qubit gate and arbitrary single qubit operations are sufficient, the former typically being the more difficult to implement well. The first entangling gate with trapped ions was performed in 1995 [Monroe et al., 1995a], based on the proposal from Ignacio Cirac and Peter Zoller for quantum computation with trapped ions [Cirac and Zoller, 1995], kickstarting interest in the platform. Since then, fidelities have greatly improved, up to around 0.999 (i.e. a 99.9% success rate) [Gaebler et al., 2016, Ballance et al., 2016, Srinivas et al., 2021, Clark et al., 2021].

The fifth and final criterion is for qubit-specific measurement capability. As with the second criterion, trapped ions typically use laser beams with controlled frequency and

polarization. Careful tuning of these parameters allows for driving “cycling” transitions where a certain state will repeatedly scatter photons, while other states are sufficiently far detuned from resonance and/or require the wrong polarization so they scatter few if any photons. Scattered photons are then recorded by a camera or photomultiplier tube, with various techniques to distinguish between qubit states. The simplest just compares the number of observed photons to a preset threshold, and work for a single detection period in a single experiment. Details of how we implement this can be found in Sec. 2.8. Errors for this process as low as $9(1) \times 10^{-5}$ in $^{43}\text{Ca}^+$ and $6(7) \times 10^{-6}$ in $^{171}\text{Yb}^+$ ions have been demonstrated [Myerson et al., 2008, Edmunds et al., 2020].

However, this process comes with risks, namely stray light from the detection laser may scatter off of a spectator qubit and prematurely decohere it, projecting any superposition it may have been in and destroying any entanglement it may have had with other qubits. In Chapter 7 we propose and demonstrate an alternative measurement protocol that avoids this danger while allowing for high fidelity, at the expense of time.

1.3 Thesis Organization

In this thesis, we demonstrate how mixed-species quantum logic, where two different atomic species of ions are used, can enhance various pillars of quantum technology. The next three chapters provide the background needed for understanding the experiments presented in subsequent chapters and how we implement them. We begin with an overview of the $^9\text{Be}^+$ and $^{25}\text{Mg}^+$ ions and basic operations on them in Chapter 2. Next, in Chapter 3 we describe the experimental infrastructure used to trap and control the qubits, including the vacuum system, laser systems, control electronics, etc. Chapter 4 outlines the necessary terminology for describing the experiments presented in remaining chapters.

Our demonstration of quantum gate teleportation [Wan et al., 2019] (Chapter 5) performs a quantum logic gate between spatially separated $^9\text{Be}^+$ qubits by means of an entangled pair of $^{25}\text{Mg}^+$ ions, with applications in both the computing and communication pillars of

quantum technology. In Chapter 6 we enhance the quantum metrology pillar by demonstrating an operation for indirect spectroscopy through the second species that is robust to various experimental imperfections and provides enhanced scaling with ion number through quantum entanglement [Kienzler et al., 2020]. Moreover, when implemented with high fidelity it has powerful applications to quantum computation as is discussed in the latter half of the chapter.

The quantum gate teleportation experiment suffered from unexpected error due to stray resonant laser light scattering off of idle spectator qubits, and hence projecting their superpositions to a single qubit level prematurely. To address this, we develop and demonstrate a protocol for high-fidelity indirect qubit measurement in Chapter 7. Qubits information is mapped to a separate species readout ion, which is then subsequently measured through resonance fluorescence detection. The readout ion resonant wavelength is very different from that of the qubit, leaving spectator qubits unperturbed if they encounter stray light from that fluorescence detection. A judicious choice of subspaces within the qubit ion states renders the indirect readout resilient to spontaneous photon scattering, the typical limiting error in laser-based quantum logic operations. Finally, in Chapter 8 we summarize the findings of this thesis and offer outlook on how quantum technology can be further enhanced through mixed-species quantum logic.

Though not discussed in this thesis, I had the pleasure of participating in several other experiments that strengthen the various pillars of quantum technology. In Ref. [Tan et al., 2017] we perform a chained Bell test that allow for tighter bounds on local realism than is possible in an ideal traditional Bell test, while closing the detection and memory loopholes. In Ref. [Wan et al., 2020] we demonstrate reordering of ions in a two-dimensional trap array while maintaining their coherence, a key primitive for satisfying the first DiVincenzo criterion for quantum computation. We develop frequency-selective tools for sensing electric field noise through ion motion in Ref. [Keller et al., 2020], which can help enhance any of the pillars of quantum technology with trapped ions by identifying noise sources in the trap,

in addition to more direct metrological applications. Finally, in Ref. [Cole et al., 2021] we present a protocol for dissipative generation of entangled W states (Eq. 4.8) that can form a logical qubit out of three physical qubits that is robust to arbitrary global Pauli noise (Eq. 4.9). Although experiments directly involving mixed-species operations for quantum simulation are not presented in this thesis, a second species can be very useful in simulating open quantum systems through controlled dissipation without unintended decoherence, for example in Dicke super- and subradiance [Sutherland, 2019, Reiter et al., 2020].

Chapter 2

${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions

The first DiVincenzo criteria (Sec. 1.2) requires a well characterized physical system in which to encode our qubits. This chapter will hence characterize our choice of qubits encoded within the ground state manifolds of ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions, including the relevant electron energy levels and how they are coupled. These isotopes have nuclear spins of $3/2$ and $5/2$, respectively, giving rise to hyperfine structure within the electronic ground state that enables extremely long coherence times at an appropriate magnetic field strength (Sec. 2.1) [Langer et al., 2005]. Choosing these particular ion species to encode our qubits is further motivated by a number of factors: In their singly-ionized state, they each have a single remaining valence electron, resulting in relatively simple alkali-like electronic structures. They also have light masses, enabling strong spin-motion coupling (see Sec. 2.9) for coupling between qubits via their collective motion in the trap. One drawback is the ultraviolet (UV) wavelengths needed to optically address these species, however recent demonstrations of high-fidelity microwave control could alleviate this concern [Harty et al., 2016, Zarantonello et al., 2019, Srinivas et al., 2021]. Details of how those controls are generated are discussed in Chapter 3, and their use in quantum logic gates that compose quantum algorithms in Chapter 4.

Producing ions from neutral atoms is done through a two-step process. First, one of the valence electrons is excited to the first excited P-state. A second photon then excites that electron to the continuum where it is no longer associated with the atom, leaving behind a positively charged ion. For Be atoms, this is done with two 235 nm photons [Zhou and

Lin, 1995, Wehlitz et al., 2003, Hsiao et al., 2008]. For Mg atoms, this can be done with two 285 nm photons [Rafiq et al., 2007]. However, the cross section from the excited P state in neutral Mg is suboptimal for 285 nm light, and in fact is higher for the ${}^9\text{Be}^+$ ion 313 nm Doppler cooling light (Sec. 2.5), so we also turn it on during Mg photoionization. Rudimentary tests comparing Mg photoionization with and without this 313 nm beam found a factor of roughly two or three improvement in photoionization rate, though the specific benefit will depend on beam intensities and overlap.

The relevant energy level manifolds that we consider include the ${}^2\text{S}_{1/2}$ ground states and ${}^2\text{P}_{1/2}$ and ${}^2\text{P}_{3/2}$ excited states for both ion species. Splittings within each manifold are governed by the combined effects of the hyperfine interaction between the electron and nuclear angular momenta and Zeeman splitting by a static external magnetic field. The Hamiltonian takes the form

$$H = hA \mathbf{I} \cdot \mathbf{J} - \boldsymbol{\mu} \cdot \mathbf{B} \quad (2.1)$$

$$= hA \mathbf{I} \cdot \mathbf{J} + \mu_B (g_J \mathbf{J} + g_I \mathbf{I}) \cdot \mathbf{B} \quad (2.2)$$

where h is Planck's constant, A the hyperfine constant (-625.008837048(10) MHz for ${}^9\text{Be}^+$ [Wineland et al., 1983] and -596.2542487(42) MHz for ${}^{25}\text{Mg}^+$ [Xu et al., 2017] ${}^2\text{S}_{1/2}$ levels), \mathbf{I} the nuclear spin operator, \mathbf{J} the electron angular momentum operator, $\mathbf{B} = B_Z \hat{z}$ a static external magnetic field defining the quantization axis, μ_B the Bohr magneton, g_J the electron g-factor, and g_I the nuclear g-factor. An analytic solution exists for certain combinations of \mathbf{I} and \mathbf{J} [Breit and Rabi, 1931].

2.1 Magnetic Field Insensitive Transitions

This Hamiltonian can be diagonalized to find the energy eigenstates $\{|i\rangle\}$ as a function of magnetic field strength B , with energies $\{E_i(B)\}$. In doing so, an important feature emerges, whereby at certain magnetic fields pairs of states $|i\rangle$ and $|j\rangle$ exhibit equal first order field-sensitivities $\frac{dE_i}{dB} = \frac{dE_j}{dB}$. When operated at these particular values of B , the

Transition	B-field (mT)	Frequency (MHz)	$\Delta \frac{d^2 E}{dB^2}$ (h MHz/mT ²)
$ F = 2, m_F = 0\rangle \leftrightarrow 1, 1\rangle$	11.944629	1207.495843	0.609733
$ 2, 1\rangle \leftrightarrow 1, 0\rangle$	11.964252	1207.352808	0.609764
$ 2, 1\rangle \leftrightarrow 1, 1\rangle$	22.307289	1082.547061	0.725157
$ 1, 1\rangle \leftrightarrow 1, 0\rangle$	160.182070	322.551896	0.00105868
$ 2, 1\rangle \leftrightarrow 2, 0\rangle$	174.719379	324.547849	0.000624669

Table 2.1: The first five magnetic field insensitive transitions in ${}^9\text{Be}^+$ at nonzero field are shown, along with their required magnetic fields, frequencies, and second order magnetic field sensitivities. We typically choose the first of these to encode qubits in ${}^9\text{Be}^+$ ions, though as will be discussed in Chapter 7 different choices may be advantageous for high-fidelity quantum logic readout.

Transition	B-field (mT)	Frequency (MHz)	$\Delta \frac{d^2 E}{dB^2}$ (h MHz/mT ²)
$ F = 3, m_F = 0\rangle \leftrightarrow 2, 1\rangle$	10.946394	1763.030898	0.433018
$ 3, 1\rangle \leftrightarrow 2, 0\rangle$	10.958430	1762.973812	0.433020
$ 3, 1\rangle \leftrightarrow 2, 1\rangle$	21.278445	1686.462050	0.465593
$ 3, 1\rangle \leftrightarrow 2, 2\rangle$	33.156830	1539.101739	0.497859
$ 3, 2\rangle \leftrightarrow 2, 1\rangle$	33.167299	1538.928892	0.497872
$ 3, 2\rangle \leftrightarrow 2, 2\rangle$	42.556890	1333.265317	0.855934

Table 2.2: The first six magnetic field insensitive transitions in ${}^{25}\text{Mg}^+$ at nonzero field are shown, along with their required magnetic fields, frequencies, and second order magnetic field sensitivities. We choose to set our magnetic field to a value where a certain transition in ${}^9\text{Be}^+$ is field-insensitive (Table 2.1), resulting in residual first order field sensitivity for ${}^{25}\text{Mg}^+$ qubits.

corresponding pairs of states make an excellent choice for encoding a qubit, because magnetic field fluctuations shift both levels equally to first order. This suppresses what is otherwise the dominant source of decoherence in trapped ions and allowing for seconds-long coherence times [Langer et al., 2005, Langer, 2006]. This fulfills the third DiVincenzo criteria (Sec. 1.2). Dynamical decoupling has been demonstrated to extend such times to up to one hour in trapped ion qubits [Wang et al., 2021]. First order magnetic field-insensitive transitions, associated magnetic field strengths, and residual second order field sensitivities $\Delta \frac{d^2 E}{dB^2} = \left| \frac{d^2 E_i}{dB^2} - \frac{d^2 E_j}{dB^2} \right|$ are listed in Table 2.1 for ${}^9\text{Be}^+$ and Table 2.2 for ${}^{25}\text{Mg}^+ {}^2\text{S}_{1/2}$ energy levels. Energy eigenstates are labeled by the total angular momentum $\mathbf{F} = \mathbf{I} + \mathbf{J}$ and its projection $m_F = m_I + m_J$ along the quantization axis of their largest component, though they are not pure eigenstates of \mathbf{F} at nonzero magnetic field strength. Several first order field-insensitive transitions can be found for each species, however there is no non-zero magnetic field for which one can be found for both.

2.2 ${}^2\text{S}_{1/2}$ Transition Frequencies

Weighing the desire for moderately low magnetic field, residual first-order field sensitivity of the second species, and ease of mapping into the qubit from an optically pumped stretch state (Sec. 2.8), we choose a magnetic field of approximately 11.9446 mT. At this field the ${}^9\text{Be}^+ |F = 2, m_F = 0\rangle \leftrightarrow |1, 1\rangle$ transition is first order field-insensitive with a second order sensitivity of $0.6097 h \text{ MHz/mT}^2$. Naturally, we then define our ${}^9\text{Be}^+$ qubit as $|2, 0\rangle \equiv |\downarrow\rangle$ and $|1, 1\rangle \equiv |\uparrow\rangle$. We encode our ${}^{25}\text{Mg}^+$ qubit as $|3, 1\rangle \equiv |\downarrow\rangle$ and $|2, 0\rangle \equiv |\uparrow\rangle$. The magnetic field for which this ${}^{25}\text{Mg}^+$ qubit transition has vanishing first order magnetic field sensitivity is nearby at 10.9584 mT, resulting in a relatively low residual first order field-sensitivity of $0.4271 h \text{ MHz/mT}$ at our chosen field. The resulting transition frequencies between all pairs of ${}^2\text{S}_{1/2}$ levels are shown in Table 2.3 for ${}^9\text{Be}^+$ and Table 2.4 for ${}^{25}\text{Mg}^+$ (with the exception of $|3, -3\rangle$ and $|3, -2\rangle$).

	$ 2, -2\rangle$	$ 2, -1\rangle$	$ 2, 0\rangle$	$ 2, 1\rangle$	$ 2, 2\rangle$	$ 1, 1\rangle$	$ 1, 0\rangle$	$ 1, -1\rangle$
$ 2, -2\rangle$	0	69.07	145.46	232.15	334.95	1352.96	1439.51	1515.75
$ 2, -1\rangle$	69.07	0	76.39	163.08	265.88	1283.89	1370.43	1446.68
$ 2, 0\rangle$	145.46	76.39	0	86.69	189.49	1207.50	1294.04	1370.29
$ 2, 1\rangle$	232.15	163.08	86.69	0	102.80	1120.81	1207.35	1283.60
$ 2, 2\rangle$	334.95	265.88	189.49	102.80	0	1018.01	1104.55	1180.80
$ 1, 1\rangle$	1352.96	1283.89	1207.50	1120.81	1018.01	0	86.55	162.79
$ 1, 0\rangle$	1439.51	1370.43	1294.04	1207.35	1104.55	86.55	0	76.25
$ 1, -1\rangle$	1515.75	1446.68	1370.29	1283.60	1180.80	162.79	76.25	0

Table 2.3: ${}^9\text{Be}^+ {}^2\text{S}_{1/2}$ frequency differences for all pairs of hyperfine states at $B = 11.9446$ mT, given in MHz and rounded to the nearest two decimal places. Note that transitions are forbidden by selection rules for many of these pairs.

	$ 3, -1\rangle$	$ 3, 0\rangle$	$ 3, 1\rangle$	$ 3, 2\rangle$	$ 3, 3\rangle$	$ 2, 2\rangle$	$ 2, 1\rangle$	$ 2, 0\rangle$	$ 2, -1\rangle$	$ 2, -2\rangle$
$ 3, -1\rangle$	0	53.30	109.93	170.58	236.27	1755.96	1816.55	1873.11	1926.35	1976.80
$ 3, 0\rangle$	53.30	0	56.62	117.28	182.97	1702.65	1763.25	1819.81	1873.05	1923.50
$ 3, 1\rangle$	109.93	56.62	0	60.65	126.35	1646.03	1706.62	1763.18	1816.43	1866.87
$ 3, 2\rangle$	170.58	117.28	60.65	0	65.69	1585.38	1645.97	1702.53	1755.77	1806.22
$ 3, 3\rangle$	236.27	182.97	126.35	65.69	0	1519.69	1580.28	1636.84	1690.08	1740.53
$ 2, 2\rangle$	1755.96	1702.65	1646.03	1585.38	1519.69	0	60.59	117.15	170.40	220.84
$ 2, 1\rangle$	1816.55	1763.25	1706.62	1645.97	1580.28	60.59	0	56.56	109.80	160.25
$ 2, 0\rangle$	1873.11	1819.81	1763.18	1702.53	1636.84	117.15	56.56	0	53.24	103.69
$ 2, -1\rangle$	1926.35	1873.05	1816.43	1755.77	1690.08	170.40	109.80	53.24	0	50.45
$ 2, -2\rangle$	1976.80	1923.50	1866.87	1806.22	1740.53	220.84	160.25	103.69	50.45	0

Table 2.4: ${}^{25}\text{Mg}^+ {}^2\text{S}_{1/2}$ frequency differences for all pairs of hyperfine states at $B = 11.9446$ mT, except those involving $|3, -3\rangle$ and $|3, -2\rangle$, given in MHz and rounded to the nearest two decimal places. Note that transitions are forbidden by selection rules for many of these pairs.

2.3 Excited States

Our qubits are encoded in $^2S_{1/2}$ ground states, however excited states in the $^2P_{1/2}$ and $^2P_{3/2}$ manifolds are used for other tasks like optical pumping, Doppler cooling (Sec. 2.5), detection (Sec. 2.8), and repumping during sideband cooling (Sec. 2.6 and 2.10). In $^9\text{Be}^+$, the $^2S_{1/2}$ states can be coupled to the $^2P_{1/2}$ states with ~ 313.197 nm laser light, and the $^2P_{3/2}$ states with ~ 313.133 nm laser light. Both excited states have natural linewidths of approximately $2\pi \times 19.4$ MHz [Poulsen et al., 1975]. Each of these manifolds is further split by the hyperfine interaction and Zeeman splittings, however we only resonantly couple to the $^2P_{1/2} |2, 2\rangle$ and $^2P_{3/2} |3, 3\rangle$ stretch states in $^9\text{Be}^+$. The other excited-state levels contribute to stimulated Raman transitions (Sec. 2.9) and AC Stark shifts. Similarly, in $^{25}\text{Mg}^+$ the state $^2P_{1/2} |3, 3\rangle$ is resonantly coupled to with ~ 280.353 nm laser light, while $^2P_{3/2} |4, 4\rangle$ and $|3, 3\rangle$ are coupled to with ~ 279.635 nm light. The additional excited state coupled to in $^{25}\text{Mg}^+$ arises from differences in how we implement our repumping (Sec. 2.6). These excited states have natural linewidths of approximately $2\pi \times 41.3$ MHz [Clos et al., 2014]. Figs. 2.1 and 2.2 show the relevant $^9\text{Be}^+$ and $^{25}\text{Mg}^+$ ion states along with laser and microwave control fields, which are explained below. Details of how the various control fields are generated can be found in Chapter 3.

2.4 Motion of Ions in the Trap

The trapping fields (Sec. 3.2) provide harmonic confinement of the ions in all three spatial dimensions. If sufficiently low in energy, the motion of the ions takes the form of quantum harmonic oscillators, governed by the Hamiltonian

$$H = \hbar\omega (\hat{n} + 1/2) \quad (2.3)$$

where \hbar is the reduced Planck constant, ω the frequency of oscillation and \hat{n} the number operator. Eigenstates $\{|n\rangle\}$, typically referred to as “Fock states” or “number states,” have

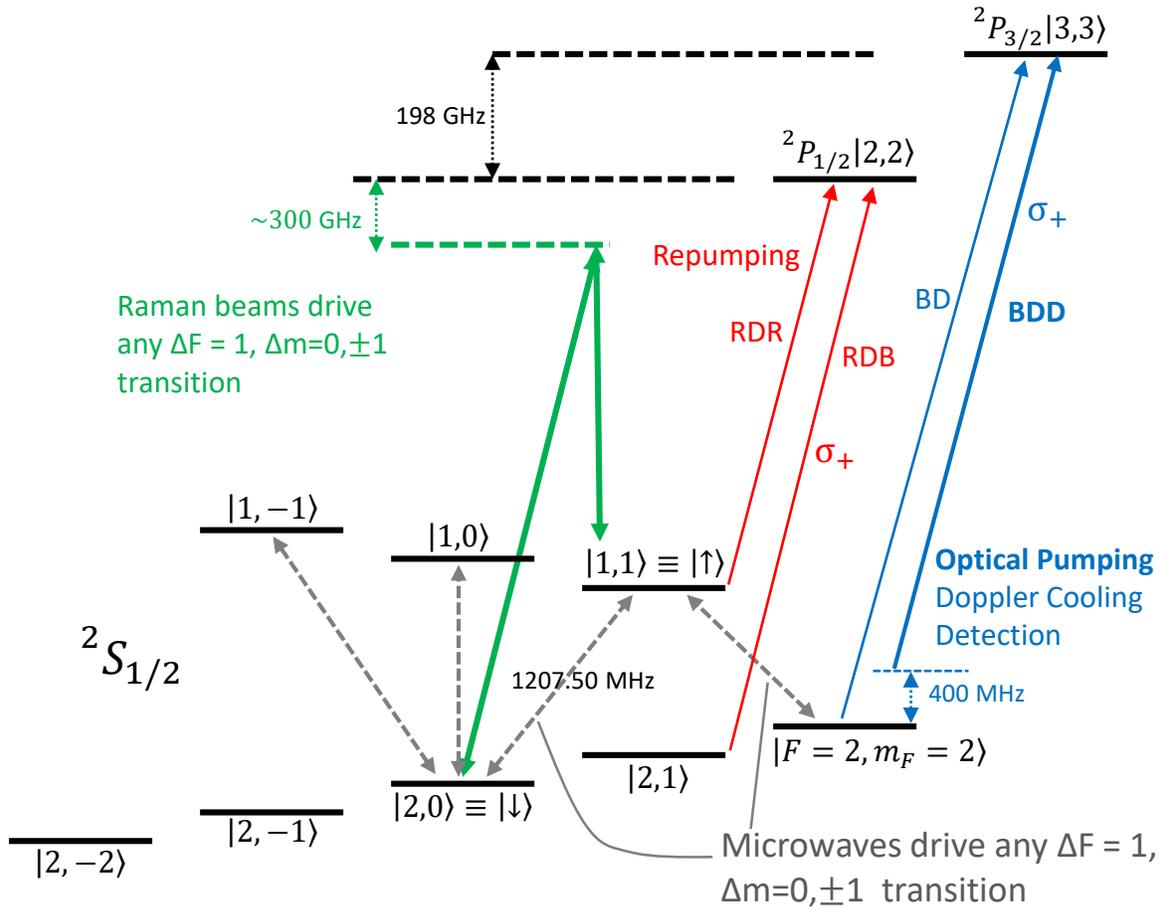


Figure 2.1: Important ${}^9\text{Be}^+$ ion energy levels at 11.9446 mT (not to scale) and control fields used to address them. All laser beams have a wavelength of approximately 313 nm. Beams coupling to the ${}^2\text{P}_{3/2}$ manifold are shown in blue. These include the “blue Doppler” (BD) beam near-resonant with the ${}^2\text{S}_{1/2} |2, 2\rangle \rightarrow {}^2\text{P}_{3/2} |3, 3\rangle$ cycling transition that is used for both detection (on resonance) and Doppler cooling (red-detuned by $\Gamma/2$) and the “blue Doppler detuned” (BDD) beam that is 400 MHz red-detuned from the cycling transition for optical pumping and cooling of highly energetic states of the motion. The BDD is drawn thicker to emphasize the high intensity of this beam relative to other resonant beams. The “red Doppler red” (RDR) and “red Doppler blue” (RDB) repumper beams are shown in red, coupling ${}^2\text{S}_{1/2} |1, 1\rangle$ and $|2, 1\rangle$, respectively, to ${}^2\text{P}_{1/2} |2, 2\rangle$, which decays at a total rate of $\Gamma = 2\pi \times 19.4$ MHz back to those two states and ${}^2\text{S}_{1/2} |2, 2\rangle$. Raman beams are shown in green coupling two particular states, though they can be used to drive any $\Delta F = \pm 1$, $\Delta m_F = 0, \pm 1$ transition in the ${}^2\text{S}_{1/2}$ manifold with appropriate polarizations. Likewise, microwaves can be used to drive any transition with the same selection rules, of which several key transitions are shown by grey dashed lines. Details on how the laser beams are generated can be found in Sec. 3.6 and microwaves in Sec. 3.8.

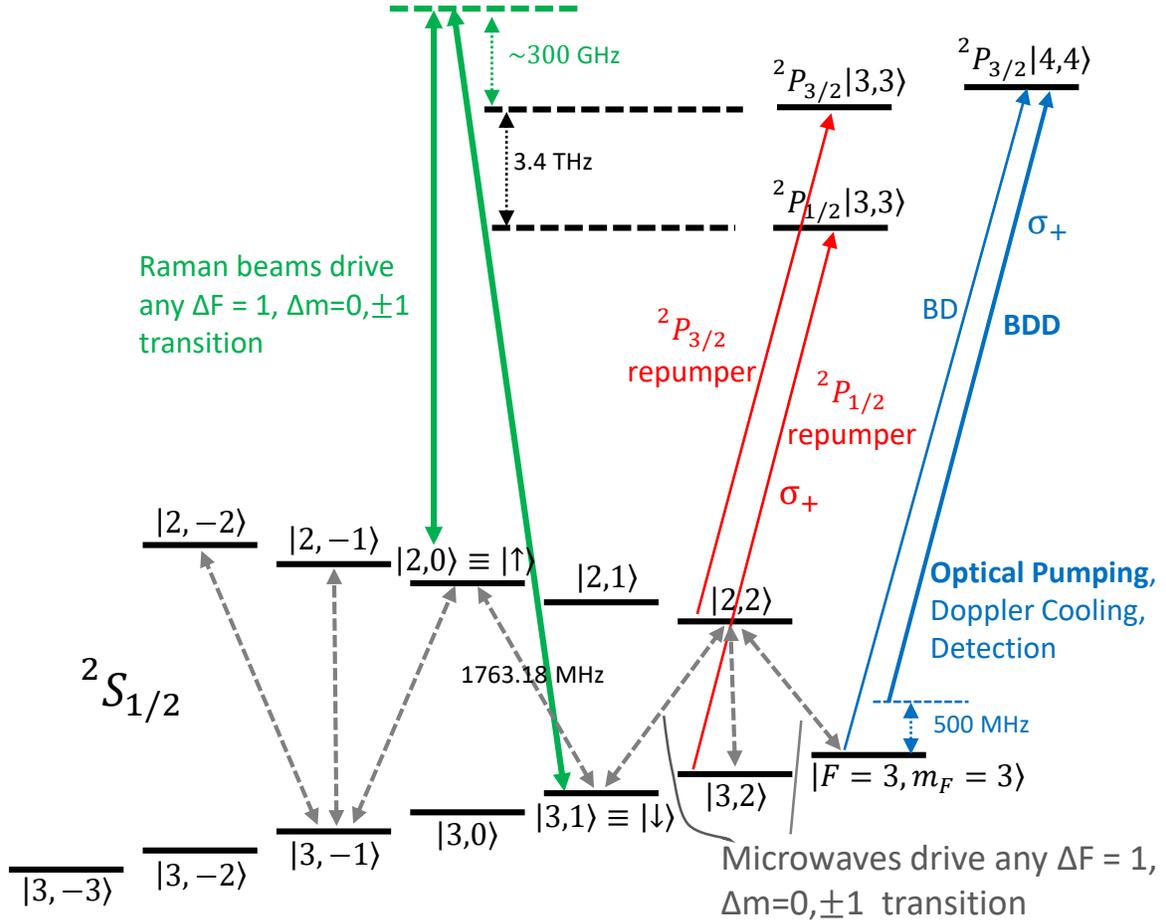


Figure 2.2: Important $^{25}\text{Mg}^+$ ion energy levels at 11.9446 mT (not to scale) and control fields used to address them. All laser beams have a wavelength of approximately 280 nm. Beams coupling to the $^2\text{P}_{3/2}$ manifold are shown in blue. These include the “blue Doppler” (BD) beam near-resonant with the $^2\text{S}_{1/2} |3, 3\rangle \rightarrow ^2\text{P}_{3/2} |4, 4\rangle$ cycling transition that is used for both detection (on resonance) and Doppler cooling (red-detuned by $\Gamma/2$) and the “blue Doppler detuned” (BDD) beam that is 500 MHz red-detuned from the cycling transition for optical pumping and cooling of highly energetic states of the motion. The BDD is drawn thicker to emphasize the high intensity of this beam relative to other resonant beams. The $^2\text{P}_{1/2}$ repumper couples $^2\text{S}_{1/2} |3, 2\rangle \rightarrow ^2\text{P}_{1/2} |3, 3\rangle$ and the $^2\text{P}_{3/2}$ repumper couples $^2\text{S}_{1/2} |2, 2\rangle \rightarrow ^2\text{P}_{3/2} |3, 3\rangle$, each decaying at a total rate of $\Gamma = 2\pi \times 41.3$ MHz back to those two states and $^2\text{S}_{1/2} |3, 3\rangle$. Raman beams are shown in green coupling two particular states, though they can be used to drive any $\Delta F = \pm 1$, $\Delta m_F = 0, \pm 1$ transition in the $^2\text{S}_{1/2}$ manifold with appropriate polarizations. Likewise, microwaves can be used to drive any transition with the same selection rules, of which several key transitions are shown by grey dashed lines. Details on how the laser beams are generated can be found in Sec. 3.7 and microwaves in Sec. 3.8.

corresponding equally-spaced energies $E_n = \hbar\omega (n + 1/2)$ where the motional quantum number $n \geq 0$ is an integer. The number operator can be rewritten in terms of the creation and annihilation operators $\hat{n} = \hat{a}^\dagger \hat{a}$. Here \hat{a}^\dagger is the creation operator, which adds one quantum of motion to a number state by effecting $\hat{a}^\dagger |n\rangle = \sqrt{n+1} |n+1\rangle$. On the other hand, \hat{a} is the annihilation operator, which removes one quantum of motion from a number state, assuming the oscillator is not in the ground state, by enacting $\hat{a} |n\rangle = \sqrt{n} |n-1\rangle$.

Single ions have three normal modes, one along each axis. A system of m ions in a linear string will have $3m$ modes, with different participation factors depending on the mode and ion. This gets more complicated when the ions have different masses, as is the case with ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$, where certain modes will have higher participation for one species than the other. We refer to modes along the axis connecting the ions as axial modes and the modes orthogonal to that as radial modes. As will be discussed below, these modes serve as the information bus through which different qubits are coupled, so this imbalance has important implications for quantum logic gates.

2.5 Optical Pumping and Doppler Cooling

Experiments typically begin with optical pumping, employing the “blue Doppler detuned” (BDD) beam for each species. The ${}^9\text{Be}^+$ BDD is 400 MHz red-detuned from the ${}^2\text{S}_{1/2} |2, 2\rangle \rightarrow {}^2\text{P}_{3/2} |3, 3\rangle$ cycling transition, and the ${}^{25}\text{Mg}^+$ BDD 500 MHz from the ${}^2\text{S}_{1/2} |3, 3\rangle \rightarrow {}^2\text{P}_{3/2} |4, 4\rangle$ cycling transition (Figs. 2.1 and 2.2). The “blue” part of the name refers to its relation with the “red Doppler” beams discussed in Sec. 2.6 for historical reasons, and should not be confused with it being blue-detuned from the cycling transition. Each of these beams is high intensity, on the order of 1 mW with a beam waist of 25 μm , to sufficiently power broaden such that it can excite any ${}^2\text{S}_{1/2}$ state to the ${}^2\text{P}_{3/2}$ manifold. They are σ_+ polarized, so each excitation to the ${}^2\text{P}_{3/2}$ manifold involves $\Delta m_J = +1$, while spontaneous decay back down to the ${}^2\text{S}_{1/2}$ levels brings $\Delta m_J = 0, \pm 1$. The combined effect gives $\Delta m_J = 0, +1, +2$, hence population is pumped towards ${}^2\text{S}_{1/2} |2, 2\rangle$ for ${}^9\text{Be}^+$ ($|3, 3\rangle$ for

$^{25}\text{Mg}^+$). These beams also serve to Doppler cool ions that are in highly excited motional states directly after loading or due to background gas collisions, where the “blue Doppler” (BD) cooling beam is Doppler shifted out of resonance for most of the ion’s trajectory and hence unable to cool efficiently.

After optical pumping with the BDD, the BD (red-detuned by $\Gamma/2$) is used for cooling to the Doppler limit, giving a final average motional quantum number \bar{n} . Assuming the weak binding regime for a motional mode of frequency ω ($\omega \ll \Gamma$), the final temperature after Doppler cooling can be calculated as

$$T_{Doppler} = \frac{\hbar\Gamma}{8k_B} (1 + \xi) \left[(1 + s) \frac{\Gamma}{2\delta} + \frac{2\delta}{\Gamma} \right] \quad (2.4)$$

where \hbar is the reduced Planck constant, k_B Boltzmann’s constant, $\xi > 0$ a geometric factor that considers beam geometry relative to the mode axis and the dipole emission pattern, s the laser saturation parameter, and δ the laser detuning from the cooling transition [Leibfried et al., 2003a, Wineland and Itano, 1979, Itano and Wineland, 1982]. The average motional quantum number can then be calculated from Bose-Einstein statistics

$$\bar{n} = \frac{1}{e^{\hbar\omega_z/k_B T_{Doppler}} - 1}. \quad (2.5)$$

For our typical trapping potentials, beam geometries, and mode frequencies, this gives \bar{n} of several quanta. The occupation of each Fock state is then given by the thermal distribution

$$P_n = \frac{\bar{n}^n}{(\bar{n} + 1)^{n+1}}. \quad (2.6)$$

2.6 Repumping

The power broadening of the BDD beam allows it to optically pump any hyperfine ground state towards a stretch state (i.e. a state with maximum $|m_F|$), however that also means that any polarization impurities in the beam cause population to leak out of the stretch state to nearby states. To combat this, we use additional σ_+ -polarized resonant repumping beams to pump leaked population back to the stretch state after optical pumping with the

BDD. These beams are low power to avoid broadening, so any polarization impurity is off resonance from being able to leak population out of the target stretch state. In addition to state preparation, the repumping beams also serve as the dissipative step during sideband cooling (Sec. 2.10).

For ${}^9\text{Be}^+$, these beams are referred to as “red Doppler red” (RDR) and “red Doppler blue” (RDB) for historical reasons. Together, they are referred to as the “red Doppler” (RD) beams or repumpers. The “red” at the beginning references it being red of the “blue Doppler” beams, and the “red” and “blue” at the end reference their relationship to each other. The RDR resonantly drives ${}^2\text{S}_{1/2} |1, 1\rangle \rightarrow {}^2\text{P}_{1/2} |2, 2\rangle$ and the RDB drives ${}^2\text{S}_{1/2} |2, 1\rangle \rightarrow {}^2\text{P}_{1/2} |2, 2\rangle$ (Fig. 2.1). To avoid trapping a coherent superposition of $|1, 1\rangle$ and $|2, 1\rangle$ through electromagnetically-induced transparency (EIT) [Marangos, 1998], one of these beams is detuned by roughly $\Gamma/2$.

For ${}^{25}\text{Mg}^+$, the two repumping beams couple to different excited states, so no EIT is possible. The ${}^2\text{P}_{1/2}$ repumper couples ${}^2\text{S}_{1/2} |3, 2\rangle \rightarrow {}^2\text{P}_{1/2} |3, 3\rangle$ and the ${}^2\text{P}_{3/2}$ repumper couples ${}^2\text{S}_{1/2} |2, 2\rangle \rightarrow {}^2\text{P}_{3/2} |3, 3\rangle$. It is important to repump ${}^2\text{S}_{1/2} |3, 2\rangle$ to the ${}^2\text{P}_{1/2}$ manifold because if coupled to ${}^2\text{P}_{3/2} |3, 3\rangle$ it would still be close enough in frequency to off-resonantly drive the ${}^2\text{S}_{1/2} |3, 3\rangle \rightarrow {}^2\text{P}_{3/2} |4, 4\rangle$ cycling transition. This would scatter a significant number photons and heat the ions through photon recoil. Some of the experiments discussed in this thesis (Chapters 5 and 6) were performed before the ${}^2\text{P}_{1/2}$ repumper was installed, so the repumping step during state preparation was performed by iterating ${}^2\text{P}_{3/2}$ repumper pulses with microwave π -pulses on the $|2, 2\rangle \leftrightarrow |3, 2\rangle$ transition. Efficient sideband cooling on ${}^{25}\text{Mg}^+$ was not possible at the time. With the addition of the ${}^2\text{P}_{1/2}$ repumper, we found that the final state preparation reached the highest fidelity by first pulsing both repumpers together, followed by a shorter sequence of iterated ${}^2\text{P}_{3/2}$ repumping and microwave pulses. We believe this to be caused by polarization impurities in the ${}^2\text{P}_{1/2}$ laser due to differential dispersion in shared optics elements and its small but nonzero wavelength difference from the ${}^2\text{P}_{3/2}$ laser. The beams are both coupled into the same optical fiber (Sec. 3.7) and are co-

propagating after emerging from it, allowing for polarization differences to accumulate when passing through waveplates, reflecting off of a dichroic mirror, and propagating through a window in the vacuum chamber that exhibits a small amount of stress-induced birefringence. Polarization is optimized for the ${}^2\text{P}_{3/2}$ laser to minimize state preparation and measurement (SPAM) error by being ideally purely σ^+ when interacting with the ions. This form of state preparation was used for the experiments discussed in Chapter 7.

2.7 Microwave Transitions

Magnetic dipole transitions between any ${}^2\text{S}_{1/2}$ levels that satisfy the selection rules $\Delta F = \pm 1, \Delta m_F = 0, \pm 1$ can be driven with microwaves for either species, with the Hamiltonian

$$H_{\text{microwave}} = \hbar\Omega |F + 1, m_F \pm 1, 0\rangle \langle F, m_F| e^{-i(\delta t + \phi)} + h.c. \quad (2.7)$$

where Ω is the transition Rabi rate, δ the detuning from resonance, ϕ the phase of the microwave drive, and $h.c.$ the Hermitian conjugate. Ideally δ is calibrated to be zero, however fluctuating magnetic fields can shift the energy levels on field sensitive transitions resulting in shot-to-shot variations in this detuning. This can be mitigated through composite pulse sequences (see Sec. 2.8). The exclusion of $\Delta F = 0, \Delta m_F = \pm 1$ in the above selection rules is an artificial limitation imposed by the filters used in our microwave control lines, not a fundamental limit (details in Sec. 3.8).

2.8 State Preparation and Measurement

The stretch state can be prepared as described in Sec. 2.5 and is a useful deterministic starting point for experiments, fulfilling the second DiVincenzo criteria (Sec. 1.2), however it is not part of a magnetic-field insensitive transition for either species at any magnetic field. The transitions needed to map into the qubit levels are field sensitive, which can lead to incomplete population transfer when the transition frequency is off-resonant with the driving

field due to magnetic field fluctuations. Composite pulse sequences are used to mitigate this effect, where an odd-numbered sequence of π -pulses with select phases result in one effective π -pulse that is less sensitive to such errors [Levitt, 1982]. We typically choose five-pulse sequences for each step, balancing time with the increased robustness from the number of pulses. For ${}^9\text{Be}^+$, we use a single composite pulse sequence to map $|2, 2\rangle \rightarrow |1, 1\rangle \equiv |\uparrow\rangle$. For ${}^{25}\text{Mg}^+$, we need to make two composite pulse transitions $|3, 3\rangle \rightarrow |2, 2\rangle \rightarrow |3, 1\rangle \equiv |\downarrow\rangle$.

State detection involves reversing the map in process, as well as shelving the other qubit state to be maximally distinguishable from the state resonant with the detection laser [Langer et al., 2005]. For ${}^9\text{Be}^+$, we shelve $|\downarrow\rangle \equiv |2, 0\rangle \rightarrow |1, -1\rangle$ with a single microwave π -pulse. Magnetic field fluctuations lead to pulse errors, which one could again mitigate through a composite pulse sequence. However, for state detection we do not require coherent population transfer to a single dark state, so instead we simply include a second shelving pulse from $|2, 0\rangle$ to a darker state $|1, 0\rangle$. Likewise, for ${}^{25}\text{Mg}^+$ we first shelve $|\uparrow\rangle \equiv |2, 0\rangle \rightarrow |3, -1\rangle \rightarrow |2, -2\rangle$ and then $|3, -1\rangle \rightarrow |2, -1\rangle$.

After shelving, we apply the BD detection laser, shifted onto resonance, for $330 \mu\text{s}$ for ${}^9\text{Be}^+$ and $200 \mu\text{s}$ for ${}^{25}\text{Mg}^+$. These durations were chosen based on the competing effect of collecting more photons to better resolve the resulting bright/dark Poisson distributions and detection errors from either the dark states off resonantly pumping to the bright state or the bright state depumping to a dark states due to polarization impurities. Each of these detections yield on average 30 counts per detection window, with the laser power controlled to stabilize this average against long term drifts. This provides a clear state-dependent measurement signal, fulfilling the fifth DiVincenzo criteria (Sec. 1.2). Results can be analyzed in many different ways. If shot-by-shot detection results are required, setting a counts threshold to distinguish between the two states can be done as in Ref. [Tan et al., 2017] or the gate teleportation experiment described in Chapter 5 [Wan et al., 2019]. If only the average over many runs of the same experiment is needed, results can be analyzed through techniques like histogram fitting or maximum likelihood estimation [Keith et al.,

2018].

2.9 Stimulated Raman Transitions

Transitions that are described by a Hamiltonian of the same form as that for microwaves (Eq. 2.7) can be driven through stimulated Raman transitions. These are two-photon processes that transition from one state to another through a virtual excited state by tuning the difference frequency of two laser beams $\Delta\omega = \omega_1 - \omega_2$ to the desired transition frequency ω_0 [Wineland et al., 1998]. In our experiment, we drive Raman transitions with two sets of laser beams for each species, one set where the beams are copropagating and the other where they are orthogonal to each other. Our Raman laser beam wavelengths are chosen to typically be a few hundred GHz detuned from the closest 2P manifold. Ideally we would detune as far as possible to minimize Raman scattering errors [Ozeri et al., 2007], but finite laser power and the desire to keep transitions reasonably fast introduce a trade-off. If transitions are too slow/narrow, drifts in the experiment or motional decoherence if coupling to the motion can negate further benefit from increased detuning. For $^9\text{Be}^+$, we red-detune from the $^2P_{1/2}$ level, while for $^{25}\text{Mg}^+$ we blue-detune from the $^2P_{3/2}$ level.

The first form of Raman transitions we introduce are called carrier transitions, where only the internal state of the ion is affected. In these transitions, the relative propagation direction of the two laser beams does not matter (ignoring errors from motional imperfections like Debye-Waller type couplings [Wineland et al., 1998]), so either the copropagating or orthogonal Raman beams can be used. The Hamiltonian for carrier rotations in the interaction picture is given by

$$H_{carrier} = \hbar\Omega_0\hat{\sigma}_+e^{-i(\delta t+\Delta\phi)} + h.c. \quad (2.8)$$

where Ω_0 is the transition Rabi rate, $\delta = \omega_1 - \omega_2 - \omega_0$ is the detuning of the difference frequency from resonance, and $\Delta\phi = \phi_1 - \phi_2$ the phase difference between the two Raman lasers. Here we introduce the $\hat{\sigma}_+$ operator, which typically refers to $|\uparrow\rangle\langle\downarrow|$, but can apply

to any pair of states that obey the selection rules $\Delta F = \pm 1, \Delta m_F = \pm 1, 0$. Its Hermitian conjugate is denoted $\hat{\sigma}_-$. As with microwave transitions, our exclusion of $\Delta F = 0$ is technical, rather than fundamental (see Sec. 3.6 and 3.7).

Raman transitions also give access to the collective motion of the ions in the trap (Sec. 2.4). By tuning $\Delta\omega = \omega_0 \pm n\omega_k$, ω_k being the frequency of the k th motional mode and n an integer, we can drive transitions that both flip the spin of the qubit and add or subtract n quanta of motion, $|\downarrow, n_0\rangle \leftrightarrow |\uparrow, n_0 \pm n\rangle$. We call these transitions “ n th blue/red sidebands,” or often just red/blue sidebands when $n = \pm 1$, where red sidebands remove quanta of motion when going from $|\downarrow\rangle$ to $|\uparrow\rangle$ and blue sidebands add quanta of motion. Here $|\downarrow\rangle$ and $|\uparrow\rangle$ can again refer to any pair of $^2S_{1/2}$ states, with lower and higher energies, respectively. Higher order sideband transitions are typically only used in initial stages of sideband cooling to bypass zero points in the first order sideband Rabi rate (Fig. 2.3) [Wineland et al., 1998]. Conservation of momentum imposes that these can only be driven with Raman beams whose difference in wave vector $\Delta\mathbf{k} = \mathbf{k}_1 - \mathbf{k}_2$ has nonzero projection onto the motional mode axis. Therefore copropagating Raman beams, having negligible $\Delta\mathbf{k}$, cannot drive these transitions. Our orthogonal Raman beams, detailed in Sec. 3.6 and 3.7, are oriented such that $\Delta\mathbf{k}$ aligns approximately with the trap axis, allowing them to couple predominantly to the axial motional modes.

Taking the rotating wave and Lamb-Dicke approximations, the ($n = 1$) red sideband (RSB) Hamiltonian for a single ion in the interaction picture is

$$H_{RSB} = \hbar\Omega_{RSB}\hat{\sigma}_+\hat{a}e^{-i(\delta t + \Delta\phi)} + h.c. \quad (2.9)$$

where Ω_{RSB} is the Rabi rate and \hat{a} the annihilation operator. For transitions between the ground and first excited motional states, $\Omega_{RSB} \approx \eta\Omega_0$ where Ω_0 is the carrier Rabi rate and η the Lamb-Dicke parameter defined as

$$\eta = |\Delta\mathbf{k}_\parallel|z_0. \quad (2.10)$$

Here $\Delta\mathbf{k}_\parallel$ is the projection of the Raman beam differential wave vector onto the motional

mode axis (in our system they are approximately parallel for axial modes and approximately orthogonal for radial modes) and $z_0 = \sqrt{\hbar/2m\omega}$ is the ground state wavepacket extent for an ion of mass m .

Similarly, blue sideband (BSB) transitions take the form:

$$H_{BSB} = \hbar\Omega_{BSB}\hat{\sigma}_+\hat{a}^\dagger e^{-i(\delta t + \Delta\phi)} + h.c. \quad (2.11)$$

where \hat{a}^\dagger is the creation operator. The Rabi rate for sideband transitions involving arbitrary number states is defined in Eq. 2.12

Finally, to implement the Mølmer-Sørensen interaction for entangling gates on multiple qubits we apply detuned red and blue sidebands simultaneously, discussed further in Sec. 4.4.

2.10 Sideband Cooling

The motional state of the ions must also be initialized to a pure quantum state at the start of each experiment that uses the motion to couple qubits to each other. In many instances, in particular when performing entangling gates or driving sidebands, fidelity is significantly improved by cooling the motion to the ground state. Several different techniques exist for cooling below the Doppler limit (Sec. 2.5), including polarization gradient cooling [Wineland et al., 1992, Ejtemaee and Haljan, 2017], electromagnetically induced transparency cooling [Morigi et al., 2000, Roos et al., 2000, Lin et al., 2013b, Feng et al., 2020, Qiao et al., 2021], and resolved sideband cooling [Monroe et al., 1995b]. Each method has its advantages and disadvantages. Though comparatively slow, we choose sideband cooling because it can get to very low \bar{n} and has no additional requirements beyond sideband transitions (Sec. 2.9) and repumping (Sec. 2.6) which are also needed for entangling gates and state initialization. This is done in one of two ways, depending on the experiment.

Pulsed sideband cooling interleaves two steps: First, a red sideband π -pulse on the $|2, 2\rangle \leftrightarrow |1, 1\rangle$ ($|3, 3\rangle \leftrightarrow |2, 2\rangle$) transition for ${}^9\text{Be}^+$ (${}^{25}\text{Mg}^+$) decreases the motional quantum

number while flipping the internal state. Second, the ions are repumped back to $|2, 2\rangle$ ($|3, 3\rangle$), resetting the internal states of the ions and providing the crucial dissipative path to the environment. The number of cycles is determined by the starting temperature and desired motional ground state fidelity. If multiple motional modes need to be cooled, pulses are interleaved between the different modes.

Continuous sideband cooling is similar, except that the red sideband and repumping are applied simultaneously as a single long pulse rather than as interleaved timed pulses. Efficient cooling requires that the repump rate and sideband Rabi rate are roughly equal. If multiple modes need to be cooled, the cooling can be done sequentially for each mode in its entirety, or broken up into several shorter periods and interleaved to mitigate the effects of heating on a previously cooled mode.

During sideband cooling, it is important to consider the Rabi rate dependence on the motional state, given by

$$\Omega_{n,n'} = \Omega e^{-\eta^2/2} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta^{|n'-n|} L_{n_{<}}^{|n'-n|}(\eta^2) \quad (2.12)$$

where Ω is the carrier Rabi rate, η the Lamb-Dicke parameter, $n_{<}$ the lesser of n' and n , $n_{>}$ the greater of n' and n , and $L_n^\alpha(x)$ the generalized Laguerre polynomial [Wineland et al., 1998]. Notably, the generalized Laguerre polynomial has zeros at which the Rabi rate vanishes (Fig. 2.3). If the initial thermal distribution has significant population beyond the first zero point, higher order sidebands must be used to cool that population.

The achievable final \bar{n} for each mode is ultimately limited by heating mechanisms that compete with the cooling. Technical noise at the mode frequency that reaches the electrodes can excite that motion. Moreover, ion motion suffers from anomalous heating of the ion motion, believed to be due to surface effects on the trap electrodes [Hite et al., 2013]. Modes with a larger displacement of their center of mass will couple more strongly to noise fields, for example the lowest energy axial mode where all ions oscillate in phase (COM mode), compared to modes with less center-of-mass motion where ions primarily oscillate

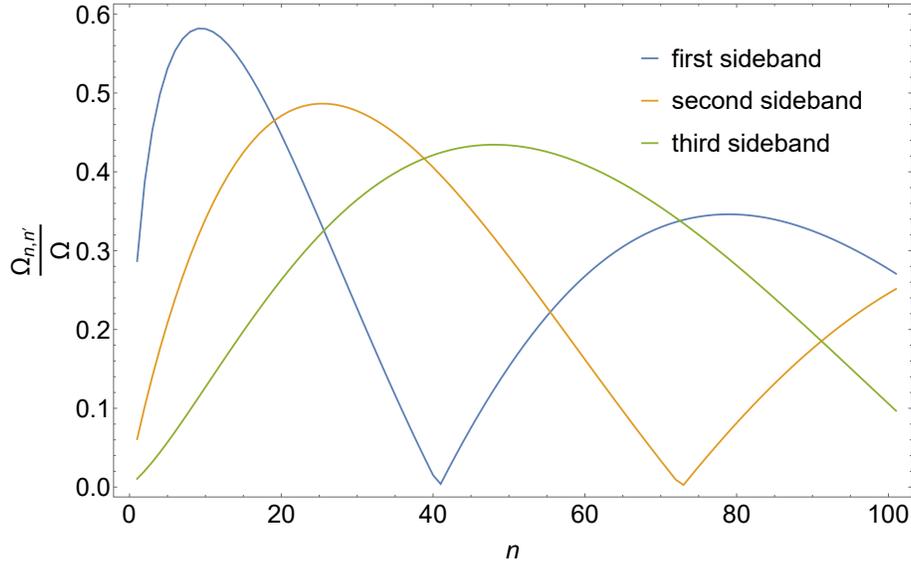


Figure 2.3: Dependence of the sideband Rabi rate $\Omega_{n,n'}$ (Eq. 2.12) on the initial and final motional quantum numbers n and n' of a $|n - n'|$ th order sideband, as a fraction of the carrier Rabi rate Ω for $\eta = 0.3$. First, second, and third sidebands are shown where $n' = n + 1, n + 2, n + 3$, respectively.

out of phase. In addition to heating due to external noise fields, photon recoil during the repumping step of sideband cooling can also excite the motion. In most experiments, each axial mode can typically be cooled to $\bar{n} < 0.1$, though for larger crystals this becomes more difficult for the modes with higher heating rates.

From the same initial state, continuous sideband cooling can bring the motion to near the ground state faster than pulsed cooling, however typically to a higher \bar{n} . Therefore, when using continuous sideband cooling we follow it with a short sequence of pulsed sideband cooling, usually around five pulses per axial mode.

Mixed-species ion crystals bring additional complications to cooling. As mentioned in Sec. 2.4, certain modes will have imbalanced participation between the two species, which impacts the coupling strength and thus how fast they can be cooled. At the beginning of an experiment the modes can be divided between the two species based on participation for

faster cooling. However, that is not possible if recooling is needed in the middle of an experiment where one of the species still stores quantum information that would be destroyed by cooling that species. Ongoing experiments in our lab are developing ways to couple motional modes to each other by modulating appropriate terms of the trapping potential at their difference frequency, extending the single ion demonstration in Ref. [Gorman et al., 2014] to multi-ion crystals. This coupling exchanges motional states, thereby allowing efficient cooling of all modes despite some having poor cooling-ion participation. For example, a mode with high participation can be sideband cooled, then exchanged with a low participation mode, and then sideband cooled again.

Chapter 3

Apparatus

In Chapter 2 we outlined the ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions and the laser and microwaves used to control them. In this chapter, we describe the details of how this control is realized, including the vacuum system, trap, magnetic field, lasers, and microwaves.

3.1 System Overview

Here we provide only a brief overview of the trap and vacuum system, since details are available in Ref. [Blakestad, 2010, Blakestad et al., 2011] and the trap has remained under vacuum since then, so performance should be similar. Our trap is enclosed in a room temperature vacuum system consisting primarily of standard stainless steel parts, with the exception of a protruding custom fused-silica glass envelope (Fig. 3.1). A copper quarter-wave coaxial resonator for the trap radio frequency (RF) drive of 82.529 MHz [Jefferts et al., 1995] extends into this glass envelope. An additional half-wave filter is located outside the vacuum system to reduce technical noise on the RF electrodes that can heat the ion motion. The trap itself sits inside the outer conductor of the RF resonator at the intersection of the optical access paths. The axial direction \hat{z} in the experiment zone aligns with the axis of the quarter-wave resonator.

External to the vacuum system, a set of Helmholtz coils produce the magnetic field defining the quantization axis (45 degrees relative to the trap axis and in the plane of the laser beams) and with the necessary strength for the chosen ${}^9\text{Be}^+$ qubit to be first order

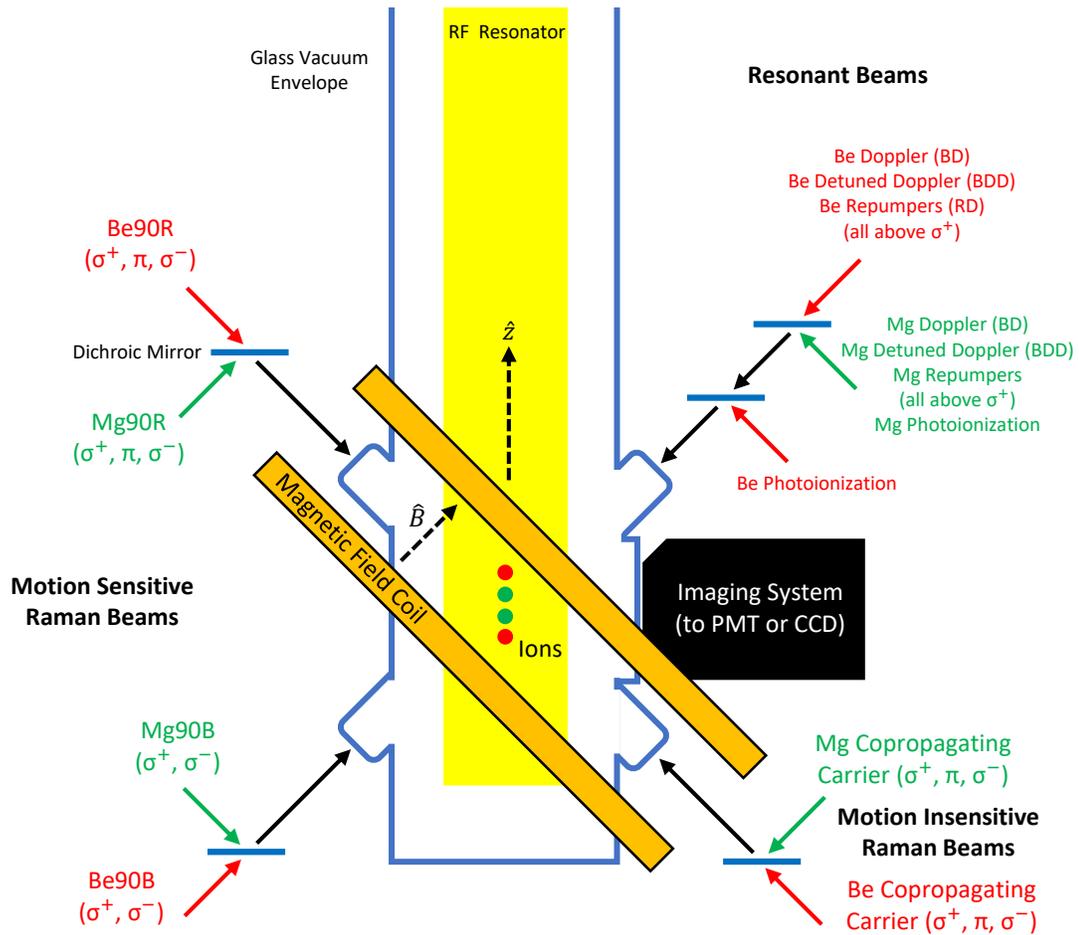


Figure 3.1: Sketch of the vacuum chamber and laser beam access. A glass vacuum envelope (cross section shown in blue) extends out from the vacuum system containing a coaxial RF resonator (yellow) for the trap drive. The trap sits vertically (out of the page) within that resonator with the trap axis in the experiment zone aligned with the axis of the resonator. External coils produce a magnetic field at a 45 degree angle to the trap axis. A primarily reflective imaging system (Sec. 3.3) collects ion fluorescence and directs it to either a PMT or CCD. Four windows allow laser access, two parallel and two orthogonal to the magnetic field. The top right window admits resonant and photoionization beams, the bottom right the copropagating Raman beams, and the two windows on the left admit the motion-sensitive Raman beams for each species (Sec. 3.6 and 3.7). All mirrors shown are dichroic and used to combine beams of significantly different wavelengths. Photoionization beams go to a separate loading zone, and resonant beams for both species are split between there and the experiment zone (see Fig. 3.2).

insensitive to magnetic field fluctuations (Sec. 2.1). The primary coils are wound out of hollow conductors with a rectangular cross section and have distilled water circulated through them to maintain their temperature and provide greater stability in the magnetic field they produce. A chiller stabilizes the temperature of the distilled water such that the resulting temperature of the coils roughly matches the surrounding environment, minimizing air temperature gradients near the trap. Approximately 62.7 A of current are used to target a magnetic field of 11.9446 mT at the ions' location, with fine tuning of the magnetic field done through a second smaller pair of Helmholtz coils. Gradients in the magnetic field throughout the QCCD array can be problematic in experiments where ions are widely separated, as in the gate teleportation experiment discussed in Chapter 5, so a set of anti-Helmholtz shim coils are used to reduce the magnetic field gradient.

Five flat windows provide clean optical access to the trap, four providing orthogonal and counter-propagating laser access (Sec. 3.6 and 3.7), with the window orthogonal to the propagation direction of the beams, and the fifth for collecting fluorescence during detection at high numerical aperture greater than 0.5 (Sec. 3.3). Our resonant and photoionization beams all enter through the same window. Beams of similar wavelength are combined through acousto-optic modulators (AOMs) or beam splitters, typically before going through an optical fiber, and beams of significantly different wavelength (e.g. 313, 280, 235 nm) are combined with dichroic mirrors. Two of the remaining windows are reserved for the motion sensitive Raman beams that propagate at 90 degrees relative to each other and are used for ground state cooling (Sec. 2.10) and gate operations (Sec. 4.2-4.3). The final laser window is used for copropagating Raman beams that drive high fidelity single qubit carrier rotations.

All laser beams except the 235 nm ${}^9\text{Be}^+$ photoionization beam pass through UV optical fibers on their way to the trap. They are photonic crystal fibers with a 10 μm core that are loaded with hydrogen in a high pressure environment before being cured with UV light. This helps prevent formation of color centers that are opaque to our UV wavelengths that would otherwise result in intolerable losses. These fibers provide a stable launching point

near the trap, reducing beam pointing fluctuations on the ions, and filter the spatial mode of UV doubler output. Details of the fibers can be found in Ref. [Colombe et al., 2014].

3.2 Trap

The ion trap used in this thesis is a segmented Paul trap that allows for separation, shuttling, and recombination of ions throughout the trap [Paul, 1990, Blakestad, 2010, Blakestad et al., 2011]. It also includes an X-shaped junction through which ions can be reordered with high fidelity and little loss to coherence of the qubit states [Wan et al., 2020]. These elements serve as a basic proof of principle implementation of ion transport, separation, and recombination as envisioned for the quantum charged-coupled-device (QCCD) architecture for scaling trapped ion quantum information processing to a large number of qubits [Wineland et al., 1998, Kielpinski et al., 2002]. In this architecture, large arrays of segmented traps are joined by junctions to provide numerous trapping zones and the ability to arbitrarily reconfigure ion positions, giving all-to-all connectivity. The trap is formed from two wafers separated by $250\ \mu\text{m}$, with a schematic shown in Fig. 3.2, reproduced from Ref. [Blakestad, 2010]. More details of the trap can be found therein, along with Refs. [Blakestad et al., 2009, Blakestad et al., 2011].

Our RF electrodes (blue electrodes in Fig. 3.2) are driven at $\omega_{RF} = 2\pi \times 82.529$ MHz to provide radial confinement (along \hat{x} and \hat{y} in Fig. 3.2 when ions are confined near the region marked \mathcal{E}) and allow for stable trapping of both ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ ions. Static potentials are applied to the DC electrodes (orange electrodes in Fig. 3.2) to provide axial confinement (along \hat{z} in Fig. 3.2 when ions are confined near \mathcal{E}). The DC electrodes can also have time-varying potentials applied to separate, reorder, and recombine ion crystals using the arbitrary waveform generators described briefly in Sec. 3.5 and in detail in Ref. [Bowler et al., 2013]. Ions are loaded at the load zone L positioned at DC electrodes 3 and 4 (Fig. 3.2) and shuttled to the experiment zone S at electrodes 19 and 20 where all of our laser beams are focused, with the exception of loading beams. Typical axial trap frequencies

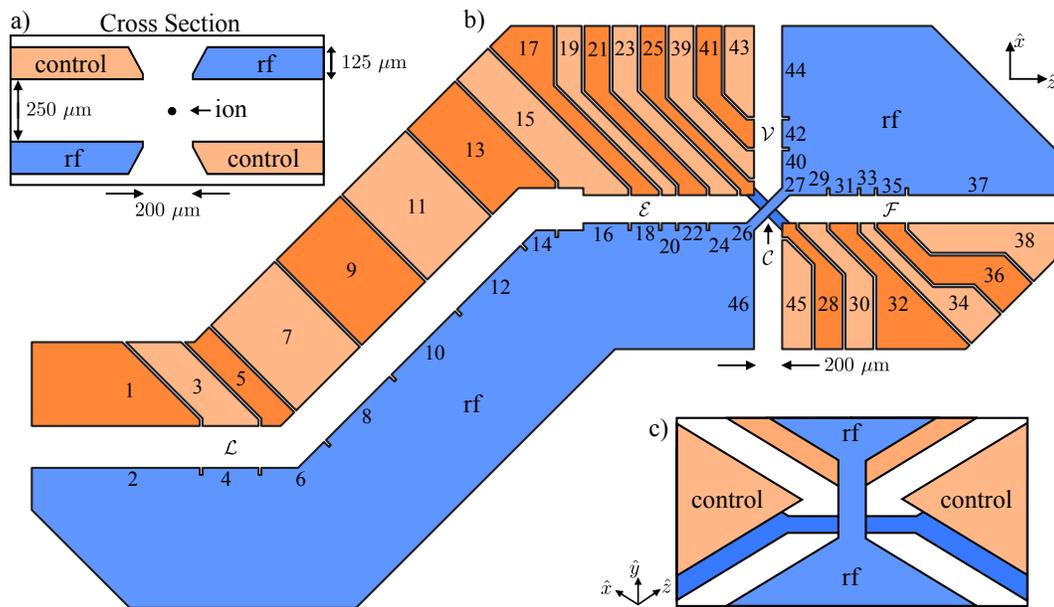


Figure 3.2: Schematic of the Ion Trap used in this thesis, reproduced from Ref. [Blakestad, 2010]. RF electrodes are shown in blue and numbered DC control electrodes in orange. This trap includes a junction for reordering of ions, though this feature is not used for the experiments detailed in this thesis. Regions important for the work described here include the load zone at the position of electrodes 3 and 4 and the experiment zone at electrodes 19 and 20. Zones nearby the load zone from electrodes 15-24 are also used to temporarily store ions for the experiment in Chapter 5. Part (a) shows the view in the plane of the page where both wafers are visible with the typical quadrupolar geometry, (b) shows the top view where the bottom wafer is obscured by the top wafer, and (c) shows an angled view near the junction.

are a few MHz while radial trap frequencies are 10-12 MHz. The experiment described in Chapter 5 involves separation and transport, while all other experiments described in this thesis took place in static potential wells.

All Paul traps exhibit “micromotion” at the RF drive frequency, arising from the fact that the RF potential is ponderomotive. In the presence of stray electric fields, ions are pushed away from their equilibrium points, resulting in excess driven motion, which can be minimized by superimposing a compensating field to ensure that the ions sit at the appropriate minimal RF potential. In most traps, this minimum coincides with a point where the RF electric field approximately vanishes. In our trap, a residual component in the axial direction remains even at the potential minimum over the radial plane for a given axial position. This effect is primarily caused by the nearby junction [Blakestad, 2010]. We therefore only minimize micromotion over the radial plane while holding the ion at a certain axial position of our choice with suitable potentials on the DC electrodes. When confining ions near the center of electrodes 19 and 20, to apply compensation fields that span the radial plane we use a differential shim on those two electrodes and a shim on a separate bias electrode which sits on a third wafer below the two main wafers shown in Fig. 3.2. Calibrations are done by tickling the radial normal modes of a single ion at its resonance frequency while scanning the appropriate shim. The tickle drive is applied to the RF electrodes by mixing the radial mode frequency into the RF trap drive. We tickle on the RF rather than DC electrodes to avoid the low-pass filters that suppress technical noise reaching the DC electrodes. When off the RF null, resonant tickling leads to strong heating of the mode such that ion fluorescence drops, giving a clear signal. The shims are not completely orthogonal, so iterating between them is necessary.

In an infinitely long linear Paul trap, this kind of motion is absent in the axial direction. For finite length there are small residual electric fields along the axis, except in a point of mirror symmetry along the axis, if such a point exists. However, our trap does not have these properties and exhibits significant axial micromotion at the experiment zone. The

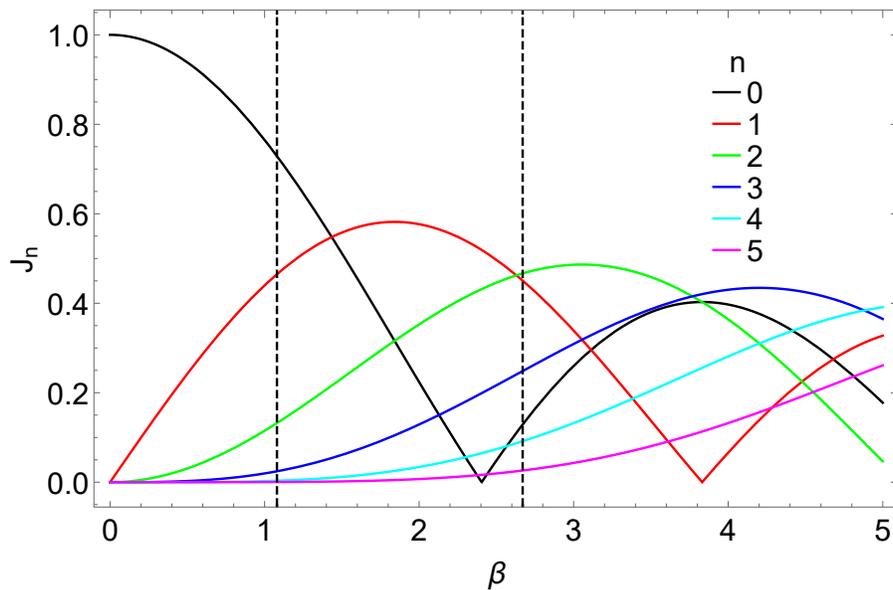


Figure 3.3: Modulus of Bessel functions of the first kind $J_n(\beta)$ for $n = \{0,1,2,3,4,5\}$ and β the axial micromotion modulation index. Motion sensitive Raman transition Rabi rates are proportional to the Bessel function evaluated at these values. Index $n \neq 0$ corresponds to driving a micromotion sideband, where the Raman laser difference frequency is shifted by $\pm n \cdot \omega_{RF}$. Vertical dashed lines mark the modulation indices observed in the experiment zone, 2.67 for ${}^9\text{Be}^+$ ions and 1.08 for ${}^{25}\text{Mg}^+$ ions.

RF junction (\mathcal{C}), angled section between the load (\mathcal{L}) and experiment zones, electrode gaps, finite size, and fabrication impurities can all contribute to this. For ${}^9\text{Be}^+$, this results in a modulation index β of 2.67 at electrodes 19 and 20, significantly impacting our motion sensitive Raman transition Rabi frequencies. Note that this is different from the value reported in Ref. [Tan, 2016] where gates were performed in a different trapping zone at electrodes 17 and 18, and possibly at a different RF power. We can instead drive transitions on micromotion sidebands, where the Raman beam difference frequency is tuned plus or minus an integer times ω_{RF} . The Rabi rate when driving on the n th micromotion sideband scales as the Bessel function $J_n(\beta)$ (Fig. 3.3). This gives relative scaling of approximately $\{0.13, 0.45, 0.47, 0.25, 0.09, 0.03\}$ for $n = \{0, 1, 2, 3, 4, 5\}$. In some experiments, we compensate for this effect with an electro-optic modulator (EOM) as discussed in Sec. 3.6. In others, the axial micromotion can be utilized to extend the range of frequencies reachable by our AO system (Sec. 3.6) or to use our Mølmer-Sørensen infrastructure (Sec. 4.4) to simultaneously drive two widely separated transitions for dissipative entanglement generation (not covered in this thesis).

${}^{25}\text{Mg}^+$ ions only see a modulation index of approximately 1.08 due to their larger mass, which leaves the micromotion carrier the strongest, with relative scaling of approximately $\{0.73, 0.47, 0.13, 0.024, 0.003, 0.0003\}$ for $n = \{0, 1, 2, 3, 4, 5\}$. We therefore drive all of our ${}^{25}\text{Mg}^+$ transitions on the micromotion carrier.

3.3 Imaging System

Simultaneously imaging both species requires careful design of the imaging system for proper focusing of both 280 nm and 313 nm light. Refractive optics are particularly prone to chromatic aberrations, so they should be avoided where possible. Therefore we primarily use reflective optics, coated for both 280 nm and 313 nm. Although these reflective optics avoid chromatic aberrations, the vacuum system window is unavoidable, so some chromatic aberration is inevitable in a system like ours where light exiting the chamber is not collimated.

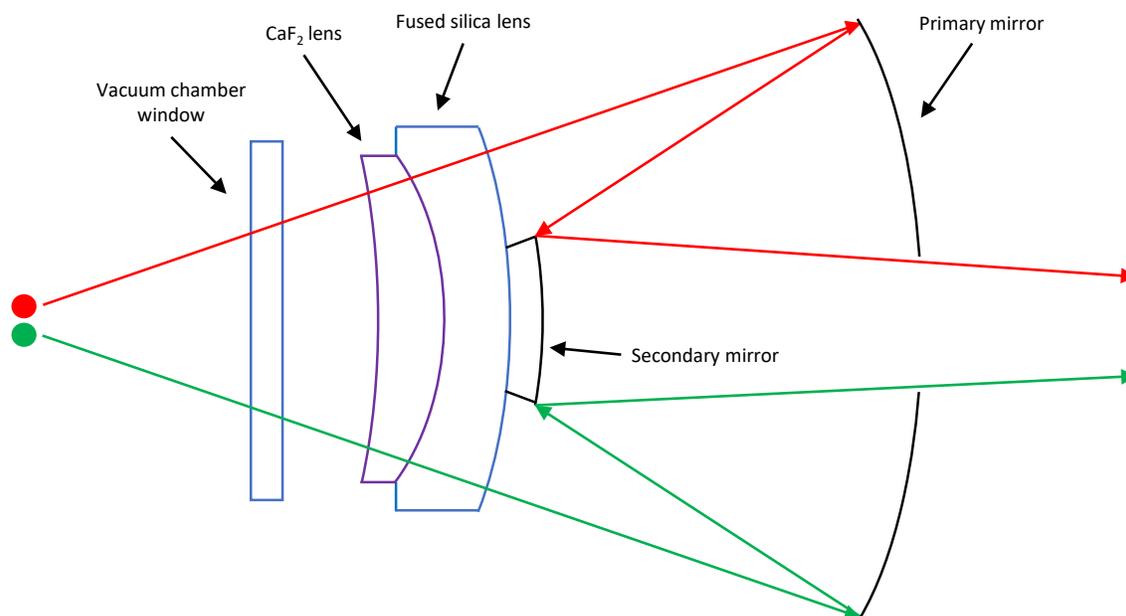


Figure 3.4: Sketch of the objective used to collect ion fluorescence for both species. Ideally, only reflective optics would be used to minimize chromatic aberrations that arise from working with two separate species of ion. However, the vacuum chamber window produces aberrations, which are corrected by a CaF_2 and fused silica pair of lenses with no net focusing power. After that, fluorescence is collected by a curved primary mirror then reflected off a curved secondary mirror before being directed to the rest of the imaging system. A second reflective magnification stage is used before directing the light to either a PMT or CCD with a flipper mirror.

The design of our objective is described in Ref. [Huang and Leibfried, 2004] and sketched in Fig. 3.4. After passing through the fused-silica vacuum chamber window, the objective begins with a pair of refractive elements to correct for chromatic aberrations from the window. A combination of a positive focal length calcium fluoride lens and a negative focal length fused-silica lens serves to correct for window aberrations while having no net focusing power itself. Light then reflects off of a curved primary mirror with a hole in the center, then off a secondary mirror attached to the center of the correcting lenses, which then reflects through the hole in the primary mirror. A second reflective magnification stage is then used before directing the light to a flipper mirror which selects between either a CCD camera or a photo-multiplier tube (PMT). All together, a magnification of 50 is achieved while allowing simultaneous imaging of both 280 nm and 313 nm light. We even found that the 235 nm photoionization light for ${}^9\text{Be}^+$ can be observed on the camera simultaneously in focus with light at 280 nm and 313 nm, which is useful for aligning the ${}^9\text{Be}^+$ photoionization beam in the load zone. Details of how we perform our resonant fluorescence detection and state discrimination can be found in Sec. 2.8.

3.4 Magnetic Field

We apply an external magnetic field that defines the quantization axis and makes our ${}^9\text{Be}^+$ qubit transition insensitive to magnetic field fluctuations to first order (Sec. 2.1). This field is primarily generated through a large pair of Helmholtz coils external to the vacuum system (Fig. 3.2) carrying a current of roughly 62.7 A to target a magnetic field of 11.9446 T. This value was arrived at by numerically solving Eq. 2.1 for the magnetic field at which the ${}^9\text{Be}^+ |2, 0\rangle \leftrightarrow |1, 1\rangle$ transition is first-order insensitive to magnetic field fluctuations. Counter-intuitively, we run the power supply driving this current in constant voltage mode, rather than constant current. We do this so that we can feedback on the voltage sensor port of the power supply, tricking it into thinking the voltage is incorrect and adjusting its output to compensate. The error signal fed into this sensor port is generated by measuring

the output current with a flux-gate current sensor, dropping that current across a stable resistor, and comparing the voltage across that resistor against a stable voltage set point [Merkel et al., 2019]. In addition to this current stabilization, we feed forward on AC line noise by adding experimentally determined 60 Hz and 180 Hz signals to the error signal [Wan et al., 2019].

Such a high current leads to significant Joule heating, so the coil is wound with hollow wires of rectangular cross section with chilled water circulating through it. The temperature of the chilled water is chosen to make the resulting steady state coil temperature roughly equal to the ambient temperature to minimize temperature gradients near the trap.

To finely tune the magnetic field, we use an additional pair of Helmholtz shim coils. In experiments where ions are separated over large distances, or are particularly sensitive to differential qubit frequency shifts, the magnetic field gradient over all positions that any ions visit must also be as small as possible. Therefore, we use an additional set of anti-Helmholtz shim coils to minimize the gradient at our experiment zone.

3.5 Control System

Experiments discussed in this thesis were controlled by one of two different, but similar, control systems. The first we refer to as the “hfgui system” in reference to the home made C-based software to control it (detailed in Ref. [Langer, 2006]), and the second we refer to as the “ARTIQ system” after the python-based programming language used to control it [Bourdeauducq et al., 2017]. Partway through my graduate studies we replaced the hfgui system with the ARTIQ control system. Each has a central master field-programmable gate array (FPGA) that precisely times transistor-transistor logic (TTL) and direct digital synthesizer (DDS) signals.

TTL signals are primarily used to control RF switches in AOM/microwave drive lines to turn laser beams or microwaves on/off and to trigger the laser pulse shaping and ion transport waveforms discussed below. One TTL channel is also used as an input to count

clicks from the PMT during fluorescence detection (Sec. 2.8 and 3.3). The hfgui system had 16 fast TTL channels while the ARTIQ system has 40 (and supports extension through modularization). To get around the limited number of TTL channels in the hfgui system, two additional measures were taken. First, we used four of the fast TTL channels as inputs to an auxiliary FPGA. This gave 16 possible output configurations, primarily used for pulsing resonant beams where timing requirements are not as strict as for Raman beams driving gates. Second, we used a hacked DDS card that utilized the DDS bus signal as TTL channels. Different TTL output configurations were chosen by setting the DDS “frequency” and used to control which transport or pulse shaping waveform was executed by the AWGs. When using this hacked TTL, appreciable wait time was needed to avoid locking up the DDS bus.

DDS signals locked to a common 10 MHz clock referenced to the NIST maser ensemble are used to drive microwave antennas and many of our AOMs by outputting a digitized sine wave with precise frequency, phase, amplitude, and timing control. AOMs that are used to control laser frequency and phase are connected to FPGA controlled DDS channels. Others that do not need real-time frequency/phase tuning, for example resonant beam switches, are driven by standalone DDS channels that lack the precise timing control. The hfgui system had 16 DDS channels while the ARTIQ system has 24 (which can also be extended through modularization).

Waveforms for ion transport and laser pulse shaping are produced by arbitrary waveform generators (AWGs) described in Refs. [Bowler et al., 2013, Bowler, 2015]. In the hfgui system, these were programmed via USB. The hacked DDS card described above was used as TTL to select which waveform to output, and a fast TTL was used to trigger the waveforms. A maximum of eight separate waveforms could be stored, limited by the number of TTL configurations for selecting the waveform. In the ARTIQ system, similar AWGs were used except with updated firmware to only use one of the three output channels, thereby providing all available memory to each individual channel instead of dividing it. This updated firmware also allows for 32 distinct waveforms. They can be programmed, selected, and

triggered via the SPI protocol, for which the ARTIQ system can output four channels. We demultiplex one of these SPI channels to program AWGs with voltage waveforms to control our DC electrodes. Though triggering via SPI is possible, we choose to trigger the AWGs via an external TTL signal to ensure precise timing. The AWG takes the logical OR of the two triggering signals if both types of triggering are used.

AWG waveforms are also used to shape Raman laser pulses (Fig. 3.5), minimizing the bandwidth of the pulses to reduce content far off the center frequency that could drive unwanted spectator transitions. This is done by sending the output waveform to a digital servo [Leibrandt and Heidecker, 2015] that is compared against the measured laser power through a pickoff photodiode after the UV fiber but before the trap. Having the pickoff after the UV fiber converts beam pointing fluctuations before the fiber into correctable laser intensity noise. The same TTL that triggers the AWG waveform is used to gate the digital servo. The output of the digital servo then feeds back to the DC port of an RF multiplier that controls the RF drive strength of a final switching AOM right before the UV fiber. In addition to shaping the pulse, this also serves as an intensity noise eater for upstream sources, including the UV doubler, beam pointing into the fiber, and prior AOMs.

Finally, additional slow TTLs are used for switching photoionization beams and the loading ovens on/off. Timing accuracy below a second is not needed for these tasks, so this control is outsourced to separate computers and communicated with via TCPIP from the host control computer. The core FPGA serves no role in these functions.

3.6 ${}^9\text{Be}^+$ lasers

All of our ${}^9\text{Be}^+$ laser beams, with the exception of photoionization, are generated by two infrared (IR) sources [Wilson et al., 2011]. First, sum frequency generation with a temperature stabilized periodically poled lithium niobate (PPLN) crystal sums two IR fiber lasers near 1050 and 1550 nm to a 626 nm beam. The BD, BDD, and RD resonant beams (for Doppler cooling, optical pumping, and repumping, respectively, as described in Sec. 2.5

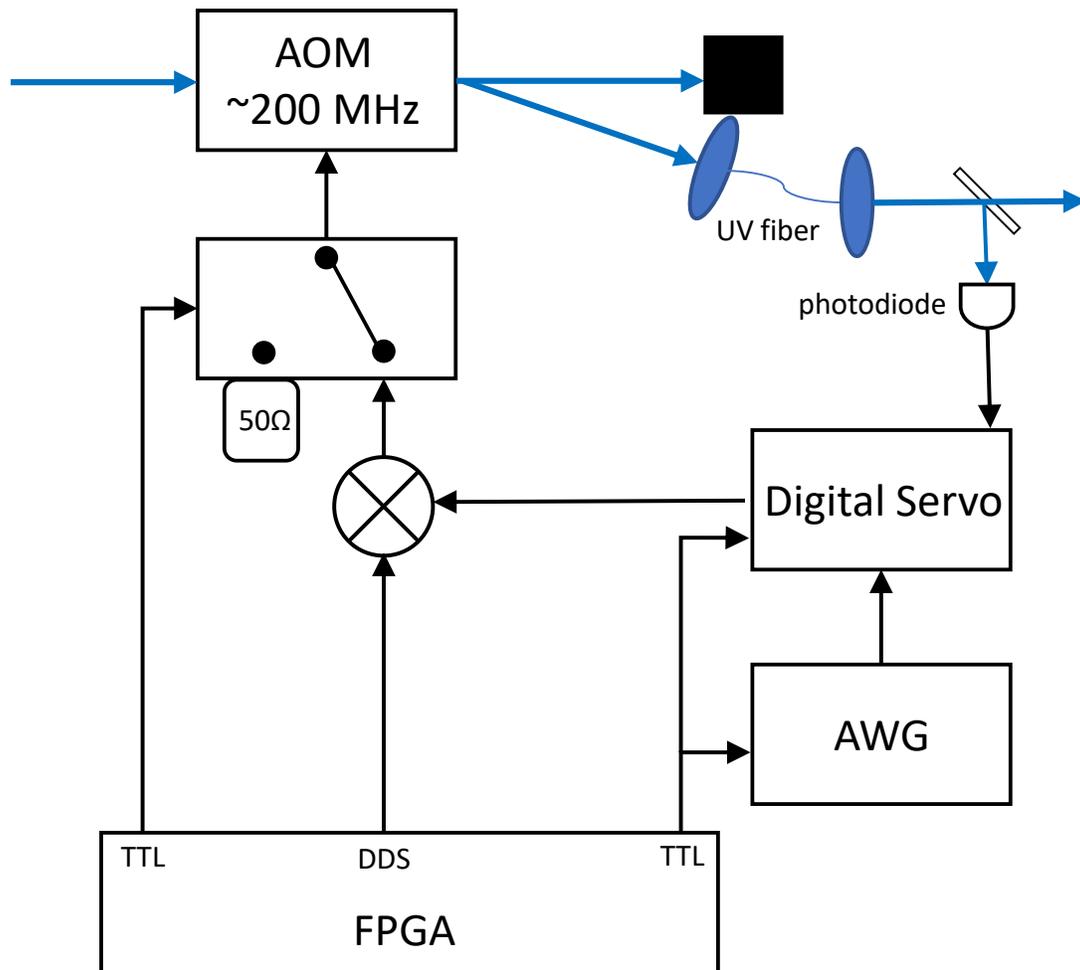


Figure 3.5: Block diagram for laser power stabilization. A desired waveform is output from an AWG to a digital servo which compares that waveform to the measured laser intensity on a pickoff photodiode. The output of the digital servo feeds back to a mixer for controlling the amplitude of the RF drive of a switch AOM before the UV optical fiber.

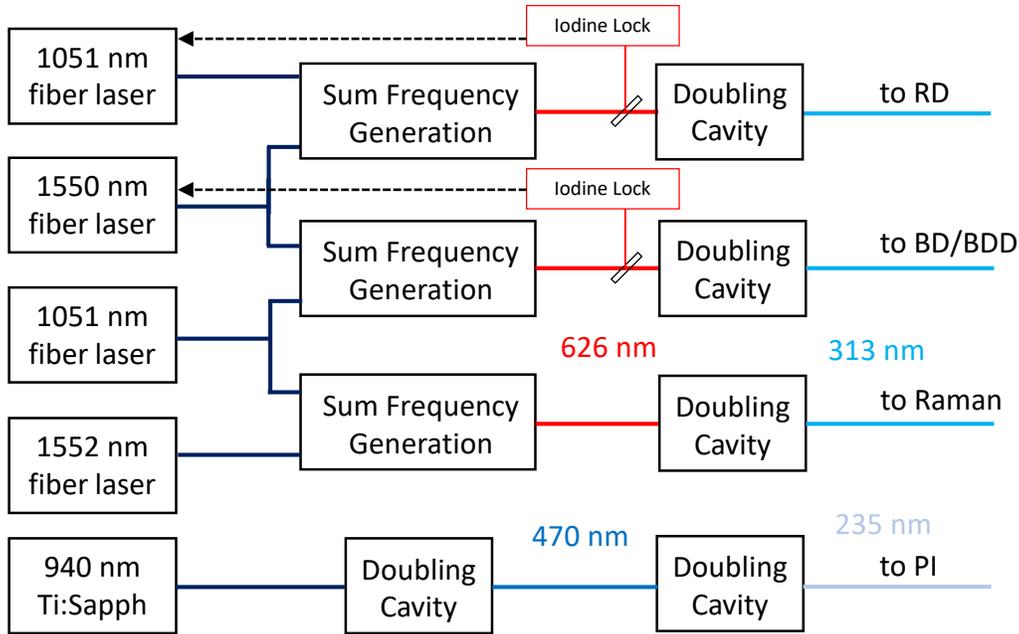


Figure 3.6: Diagram of ${}^9\text{Be}^+$ laser sources. All laser beams except the photoionization laser are generated from sum-frequency generation of two infrared fiber lasers to visible ~ 626 nm light, followed by doubling with a BBO crystal in a doubling cavity to ~ 313 nm. In resonant beamlines we sample the 626 nm light to send to a Doppler-free iodine lock for frequency stabilization. Finally, the 235 nm photoionization light is generated by doubling a Ti:Sapph laser twice.

and 2.6), are frequency stabilized by sampling the 626 nm light and locking to a fringe of the Doppler-free saturated absorption spectrum of iodine. The three frequency groupings that we need to drive, the BD/BDD, RD, and Raman beams, are generated by only four fiber laser sources, with two of them shared (Fig. 3.6). We have also tested a vertical external-cavity surface-emitting laser (VECSEL) that instead generates ~ 1252 nm light that gets frequency doubled twice to 313 nm [Burd et al., 2020].

Early on in my graduate studies, we generated our 235 nm ${}^9\text{Be}^+$ photoionization light by free space doubling 705 nm light from a mode-locked Ti:Sapph, then summing the doubled 352.5 nm light with residual 705 nm. This source was used, for example, early on in my graduate studies in the chained-Bell test experiment described in Ref. [Tan et al., 2017]. In

all experiments since then, we used a continuous-wave (CW) laser source generated by a 940 nm CW Ti:Sapphire laser that is frequency doubled twice with resonant doubling cavities.

3.6.1 $^9\text{Be}^+$ resonant lasers

The BD and RD lasers get split and recombined into multiple different beams with AOMs and beamsplitters, the details of which are shown in Fig. 3.7. Both single pass and double-pass AOM configurations are used to shift frequencies and switch beams on and off. Moreover, AOMs are used to recombine beams with low loss in the following manner: If two beams A and B come in at an angle, such that the zeroth diffraction order of beam A overlaps with the first diffraction order of beam B after the AOM, the two can be made to copropagate. Beam B is switched by this AOM, while A must be switched elsewhere. Both beams can be turned on simultaneously, at the expense of reduced power in A due to AOM diffraction.

The output of the BD UV doubler, after passing through a quadrant detector position lock, is split into the BD and BDD beams with the BD Split AOM driven at 200 MHz, where the zeroth order becomes the BDD and the plus first order the BD. The BD is then sent through the ~ 200 MHz double pass BD Switch AOM where the zeroth order is blocked and the plus first order retroreflected back through. This serves to both shift the frequency of the BD between cooling and detection configurations (Sec. 2.5 and 2.8) and individually switch the beam. Afterwards, the BD and BDD are recombined at the BDD Switch +200 MHz single pass AOM as described above, with the BD zeroth order overlapping the BDD first order. To ensure complete distinction of the beams and avoid decoherence from leaked resonant light, these beams then pass through a pair of single pass shared switch AOMs, one using the plus first and the other the minus first order to provide no net frequency shift.

Finally, before coupling into an optical fiber to the trap, the beams are modulated by a UV electro-optic modulator (EOM) driven at ω_{RF} to compensate for the effects of axial micromotion. This requires that the EOM drive exactly matches ω_{RF} and has a stable

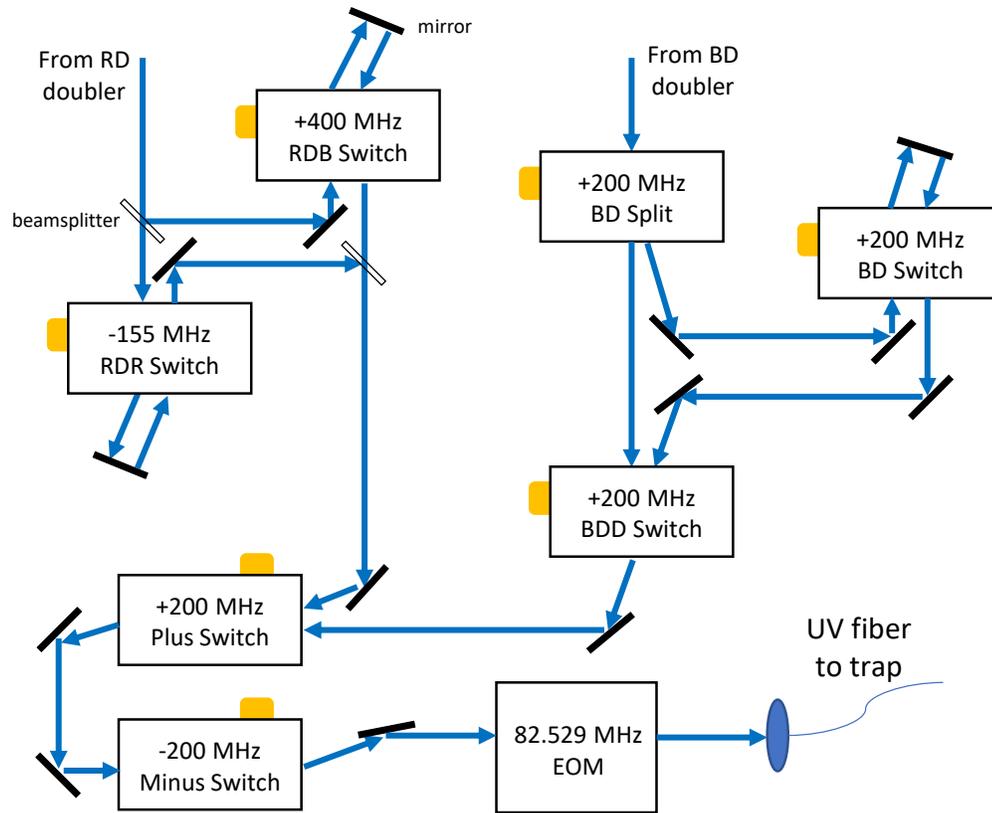


Figure 3.7: AOM/EOM setup used to generate our ${}^9\text{Be}^+$ resonant beams. The BD and RD beamlines are generated by summing separate lasers near 1051 nm and a shared 1550 nm laser, then doubling in a resonant cavity (Fig. 3.6). The BD doubler output first gets split into the BD and BDD beamlines. The BD gets frequency shifted and switched with a double pass AOM, then combined again with the BDD by overlapping their zeroth and first orders, respectively, at the BDD switch AOM. The combined beams then pass through two common switches to ensure complete extinction when turned off, through an EOM to compensate for axial micromotion (Sec. 3.2), then through an optical fiber to the trap. After the RD doubler, the beam is divided into the RDR and RDB which are each sent through their own frequency shifting and switching AOMs then recombined on a beamsplitter. The RD beams are then combined with the BD beams on the first shared switch by overlapping their zeroth order with the BD first order.

relative phase, so we source it by directly picking off from the trap drive. Because of this, any frequency or phase noise in the trap drive, and resultant axial micromotion, will be common mode with the EOM, thus preserving its cancelling effect. If the the phase of the drive relative to the trap is incorrect, it is possible to actually amplify the axial micromotion effect, rather than cancel it out, so care was taken to tune the phase and amplitude of the EOM drive. Provided that nothing changes in either the trap or EOM drive lines, this configuration is stable long term due to the long wavelength.

To calibrate the EOM drive, we first ensure that the EOM resonance matches the trap drive by tuning its temperature while monitoring the backreflected RF power. Next we load a single ${}^9\text{Be}^+$ ion, scan a radial shim on electrodes 19 and 20 (Fig. 3.2), and detect. This shim will induce radial micromotion when pushed off the RF null. Importantly, the ${}^9\text{Be}^+$ detection beam (BD) comes in at a 45 degree angle to the trap axis (Fig. 3.1). In the presence of axial micromotion, there is a shim value for which the net micromotion direction, radial plus axial, is orthogonal to the detection beam and thus has no effect on it. This results in a fluorescence peak at this condition. As an aside, prior to installing this EOM, such a shim was used during detection to maximize fluorescence [Tan, 2016]. If this peak occurs at a nonzero shim (relative to the shims necessary to compensate radial micromotion, see Sec. 3.2), either the EOM is not modulating with the correct strength or the phase is incorrect. We first iterate this scan with changes in phase, which are coarsely done through cable length (since 82.529 MHz corresponds to a 3.63 m wavelength). Once the fluorescence peak is at zero shim, the strength should be tuned to maximize the height of the peak.

The output of the RD doubler also goes through a position lock before being split into the RDR and RDB paths with a beamsplitter, as shown in Fig. 3.7. Each then passes through its own double pass switching AOM, with the RDB shifted positively and the RDR shifted negatively to provide the ~ 1 GHz splitting between the two beams. They are then recombined with each other on a beamsplitter, then with the BD beams on the first common switch such that their zeroth order overlaps with the BD diffracted first order.

Beam	RDR Switch	RDB Switch	BD Split	BD Switch	BDD Switch	+ Switch	- Switch
BD	off	off	on	on	off	on	on
BDD	off	off	off	off	on	on	on
RDR	on	off	off	off	off	off	on
RDB	off	on	off	off	off	off	on

Table 3.1: ${}^9\text{Be}^+$ resonant AOM switching configurations to turn on the various laser beams, as defined in Fig. 3.7. To turn on multiple beams simultaneously, take the logical OR of the desired rows within each column.

There is a pickoff that sends light to a logarithmic photodiode after the fiber to monitor the power in all the beams and to stabilize the intensity during detection. Unlike the Raman beam intensity stabilization (Fig. 3.5), precise shaping of the pulses is not necessary, so the photodiode signal is simply compared to a static voltage offset in the servo.

All together, each beam can be switched on independently. Any combination can also be turned on, though sometimes at the expense of power. For example, if all beams are turned on simultaneously, the BD will lose power at the BDD switch that it would not have if switched on alone (in which case that AOM would be off). Likewise, the RD beams will lose power at the Plus Switch if turned on with either of the BD beams. Table 3.1 lists which AOMs need to be switched on to turn on each of the beams. To turn on multiple, take the logical OR of the relevant rows within each column.

3.6.2 ${}^9\text{Be}^+$ Raman lasers

To mitigate errors from spontaneous photon scattering, our ${}^9\text{Be}^+$ Raman beams are red-detuned by several hundred GHz from the ${}^2\text{P}_{1/2}$ manifold [Ozeri et al., 2007]. As with the resonant beamlines, the ${}^9\text{Be}^+$ Raman laser gets divided into several beams after the UV doubler. After passing through a position lock, the laser gets divided into the Be90B and Be90R beamlines (Fig. 3.8), forming the two arms of our motion sensitive Raman beams (Sec. 2.9). The Be90B goes through a single pass +200 MHz switch before coupling into an optical fiber to the trap.

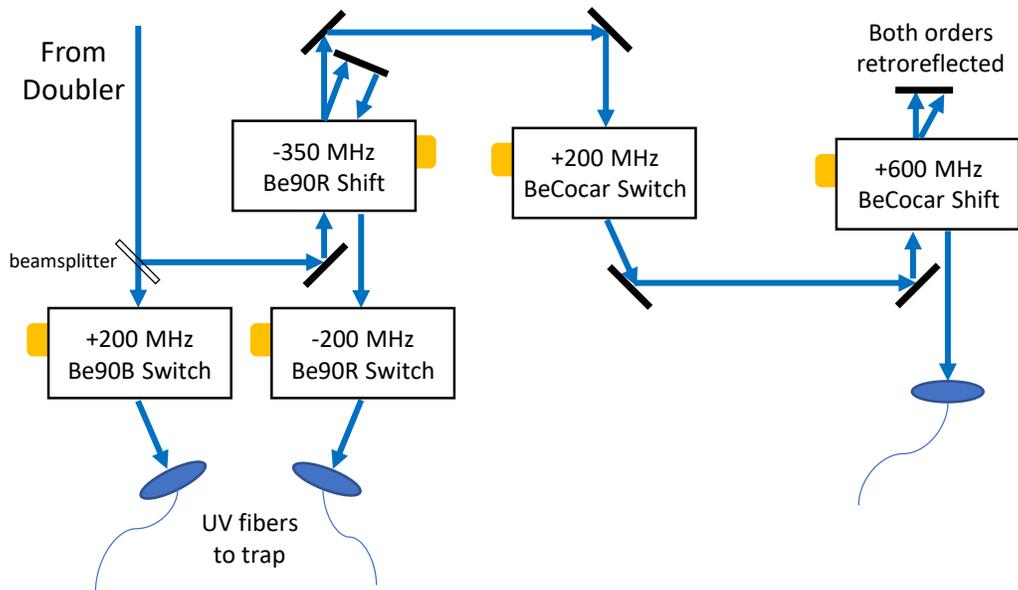


Figure 3.8: AOM setup used to generate our ${}^9\text{Be}^+$ Raman beams. The output of the UV doubler is split into the Be90B and Be90R beamlines with a beamsplitter. The Be90B beam goes through a single pass +200 MHz switch AOM before coupling into a fiber to the trap. The Be90R beam first passes through a double pass AOM to shift its frequency on the minus first order before going through its own -200 MHz switch. The BeCocar beam is split off from the Be90R beam as the zeroth order of the Be90R Shift AOM. It then passes first through a switch, and then through a 600 MHz double pass AOM where both the zeroth and plus first order are retroreflected and coupled into an optical fiber.

The Be90R first passes through the Be90R Shift AOM twice to shift its frequency by $2 \times 300 - 400$ MHz, providing most of the hyperfine splitting needed to couple the different F levels and the tunability to drive multiple transitions. A negative frequency shift is used in this beam line to go the opposite direction of the shift from the Be90B switch. The zeroth order after the first pass is dumped, and the minus first order retroflected to be shifted again during the second pass. The minus first order of that second pass through the AOM is sent to a -200 MHz switch AOM and then couples into an optical fiber to the trap. Our key entangling operation, the Mølmer-Sørensen interaction (Sec. 4.4), requires simultaneously driving detuned red and blue sidebands. To achieve this we inject two tones into the Be90R Switch AOM, separated by $\sim 5-10$ MHz depending on the motional mode, both near 200 MHz.

The Be90B and Be90R take separate paths to the ion to come in at a 90 degree angle, providing a large $\Delta \mathbf{k}$ along the trap axis. This allows these beams to drive both carrier transitions (Sec. 4.2) and spin-motion coupling of axial modes for entangling gates (Sec. 4.4). Since these two beams take drastically different paths to the ions, air turbulence, table vibrations, and temperature drifts lead to a relative phase fluctuation between the two beams. The resulting fluctuating phase difference translates to phase noise on the operations they drive. A method to work around this is discussed in Sec. 4.3.

Finally, the BeCocar copropagating carrier beam line branches off from the Be90R as the zeroth order of the Be90RShift AOM. This pair of overlapped beams can only be used to drive carrier rotations since it has negligible $\Delta \mathbf{k}$ and therefore cannot couple to the motion (Sec. 2.9). It first passes through a +200 MHz switch, then through a 600 MHz double pass AOM where both the zeroth and plus first order are retroreflected and then coupled into an optical fiber to the trap. This requires careful lens positioning relative to the double pass AOM to ensure that the two tones are copropagating and thus able to simultaneously couple into the fiber. The BeCocar switch is placed first to minimize the distance between the BeCocar Shift AOM and the optical fiber, making it easier to couple

Beam	Be90R Shift	Be90R Switch	Be90B Switch	BeCocar Switch	BeCocar Shift
Be90R	on	on	off	off	off
Be90B	off	off	on	off	off
BeCocar	off	off	off	on	on

Table 3.2: ${}^9\text{Be}^+$ Raman AOM switching configurations to turn on the various laser beams, as defined in Fig. 3.8. To turn on multiple beams simultaneously, take the logical OR of the relevant rows within each column.

in both beams. To minimize phase fluctuations between the two tones, which translates to phase errors on the resulting operations, the distance during which they take separate paths in the retroreflector should be as small as reasonably possible. All together, we can address the entire range of shelving transitions with these beams, from $|2, 2\rangle \leftrightarrow |1, 1\rangle$ around 1018 MHz to $|2, 0\rangle \leftrightarrow |1, -1\rangle$ around 1370 MHz.

After their fibers, each of these three beam lines is sampled to provide feedback for pulse shaping and laser power stabilization (Fig. 3.5). They also each pass through a Glan-Thompson polarizer and waveplates to provide stable polarization control. The AOM configurations for switching these beams on, as sketched in Fig. 3.8, are listed in Table 3.2. Like with the ${}^9\text{Be}^+$ resonant beams, multiple beams can be turned on by taking the logical OR of the relevant rows within each column, noting that the BeCocar will have reduced power when the Be90R is also on.

3.7 ${}^{25}\text{Mg}^+$ lasers

Like our ${}^9\text{Be}^+$ lasers, all of our ${}^{25}\text{Mg}^+$ lasers begin in the IR. However, rather than summing and then doubling two different IR sources, we take single IR sources and double them twice. These take the form of a fiber laser for coupling the ground states to the ${}^2\text{P}_{3/2}$ manifold, VECSELs for our ${}^2\text{P}_{1/2}$ and photoionization lasers [Burd et al., 2016], and a diode laser for our Raman beams (Fig. 3.9). The photoionization VECSEL outputs 1140 nm light that is then divided between several different experiments, with our portion getting

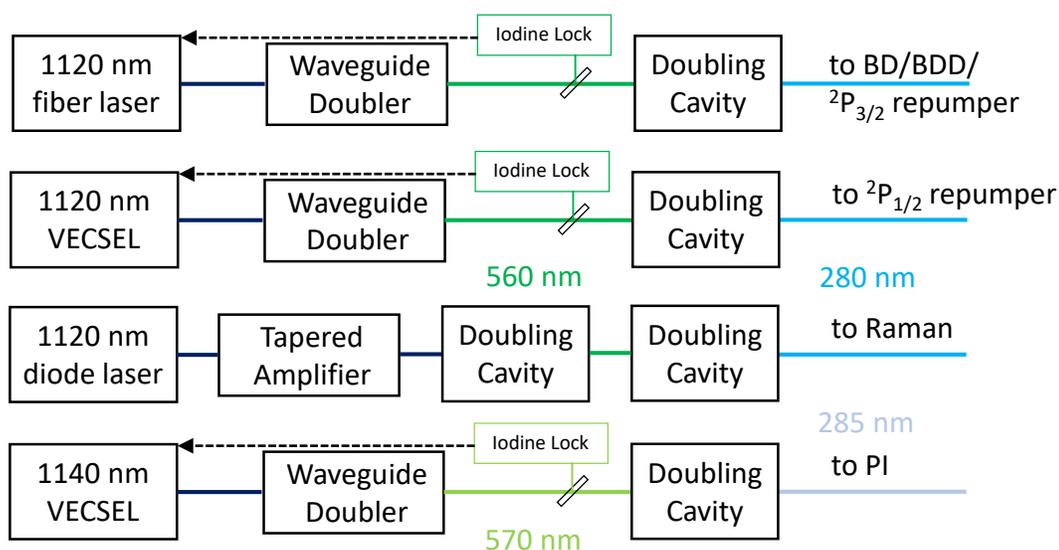


Figure 3.9: Diagram of $^{25}\text{Mg}^+$ laser sources. All beams are generated by frequency doubling an IR source twice. Resonant and photoionization beams are first doubled with waveguide doublers, then with resonant doubling cavities. The Raman beams are doubled with two successive resonant doubling cavities to enable higher power. Visible light from the resonant and photoionization beams is used to lock the frequency to iodine reference cells. The resonant beam locks are in Doppler-free configurations, while the photoionization beam is a Doppler-broadened setup because of looser requirements.

shifted by an AOM before coupling into a fiber to a waveguide doubler. The 570 nm output of the waveguide doubler is again doubled to 285 nm by a resonant doubling cavity before coupling into a UV optical fiber to the trap.

3.7.1 $^{25}\text{Mg}^+$ resonant lasers

The $^2\text{P}_{3/2}$ laser beams, including the BD/BDD and $^2\text{P}_{3/2}$ repumper (See Sec. 2.5 and Sec. 2.6) are derived from the same IR fiber laser source near 1120 nm. A waveguide doubler then uses the light from this source to produce a 560 nm beam. These waveguide doublers are very straight forward and convenient, requiring only fiber coupling into the module and controlling its temperature. However, they are limited by a damage threshold of only a few hundred mW of 560 nm power. This makes them ill-suited for generating Raman beams since this limits the UV power and thus how far we can detune (to suppress spontaneous emission errors) while maintaining reasonable gate times. The 560 nm beam then splits off a sample to a Doppler-free iodine lock, which feeds back on a piezo that puts strain on the fiber source to tune and stabilize the wavelength.

A resonant doubling cavity produces the 280 nm UV light that then passes through a position lock before a system of AOMs to split out, frequency tune, then recombine the different beams (Fig. 3.10). First the $^2\text{P}_{3/2}$ repumper is split off from the BD/BDD path. This beam is shifted -800 MHz by double passing a 400 MHz AOM on the minus first order, which also serves to independently switch this beam. From there it is combined with the $^2\text{P}_{1/2}$ repumper on a beamsplitter, and the two combined with the BD on a second beamsplitter. The $^2\text{P}_{1/2}$ repumper is generated by a VECSEL laser that gets doubled twice, first through a waveguide doubler and then again through a resonant doubling cavity.

The BD is split off from the BDD beamline and shifted +300 MHz with the BD Split AOM. It then passes through its own double pass AOM that serves to both tune the frequency between detection and Doppler cooling configurations (see Sec. 2.5 and 2.8), and individually switch it on/off. As mentioned previously, it is then combined with the two

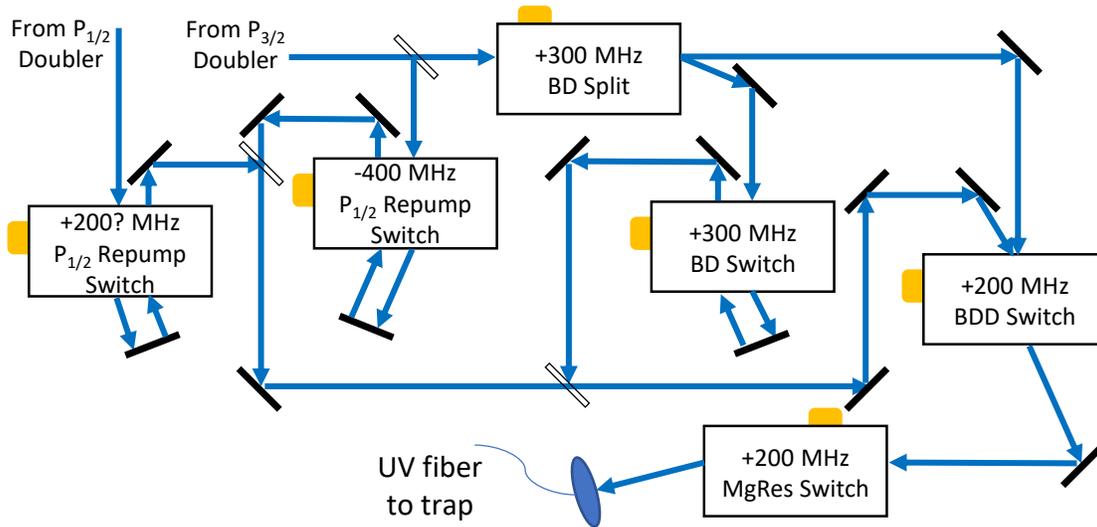


Figure 3.10: AOM setup used to generate our $^{25}\text{Mg}^+$ resonant beams. The $^2P_{1/2}$ repumper is generated by its own laser source, then passes through a 200 MHz double pass AOM that serves both to switch this beam and tune its frequency. It then gets combined with the $^2P_{3/2}$ repumper on a beamsplitter. The other resonant beams (BD,BDD, $^2P_{3/2}$ reumper) all couple to the $^2P_{3/2}$ manifold and are derived from a single source. The output of the UV doubler first has a portion separated for the $^2P_{3/2}$ repumper, which passes through a double pass frequency shifting/switching AOM before being combined with the $^2P_{1/2}$ repumper on a beamsplitter. The rest of the $^2P_{3/2}$ light then goes to the BD Split AOM which separates the BD from the BDD as the plus first and zeroth orders, respectively. The BD then passes through the BD Switch double pass AOM that individually switches this beam and tunes the frequency between cooling and detection configurations. It then gets combined with the two repumpers on a beamsplitter. Meanwhile, the BDD proceeds to pass through a 200 MHz switch. The BD and repumpers come in at an angle such that their zeroth order overlaps with the BDD plus first order, thus recombining all of the beams. They then pass through a shared switch and couple into a single optical fiber to the trap.

Beam	$^2P_{1/2}$ Rep. Sw.	$^2P_{3/2}$ Rep. Sw.	BD Split	BD Sw.	BDD Sw.	MgRes Sw.
$^2P_{1/2}$ Repump	on	off	off	off	off	on
$^2P_{3/2}$ Repump	off	on	off	off	off	on
BD	off	off	on	on	off	on
BDD	off	off	off	off	on	on

Table 3.3: $^{25}\text{Mg}^+$ resonant AOM switching configurations to turn on the various laser beams, as defined in Fig. 3.10. To turn on multiple beams simultaneously, take the logical OR within each column.

repumpers at a beamsplitter.

After the BD is split off from it, the BDD goes through the BDD Switch which serves to independently switch this beam, shift its frequency, and recombine it with the BD and repumpers by overlapping its plus first order with their zeroth order. At this point all $^{25}\text{Mg}^+$ resonant beams are copropagating. They pass through a single shared switch to ensure high extinction when off, then couple into an optical fiber to the trap. After the fiber a pickoff directs a portion of the light to a logarithmic photodetector for monitoring all beam powers and stabilizing the intensity of the BD detection beam to give consistent fluorescence histograms.

Table 3.3 gives the conditions for turning on individual $^{25}\text{Mg}^+$ resonant laser beams. As with similar tables in Sec. 3.6, multiple beams can be switched on simultaneously by taking the logical OR of the relevant rows within each column.

3.7.2 $^{25}\text{Mg}^+$ Raman lasers

Our $^{25}\text{Mg}^+$ Raman beams are a few hundred GHz blue detuned from the $^2P_{3/2}$ manifold. A commercial laser system produces ~ 1.8 W of 560 nm light by starting with a diode laser that goes through a tapered amplifier and then internal resonant doubling cavity. We then take the visible output and double it again with a homemade UV doubling cavity to produce up to 600 mW of 280 nm light, though this level was usually unstable. More reliably we work with 300-400 mW of UV light to minimize doubler maintenance.

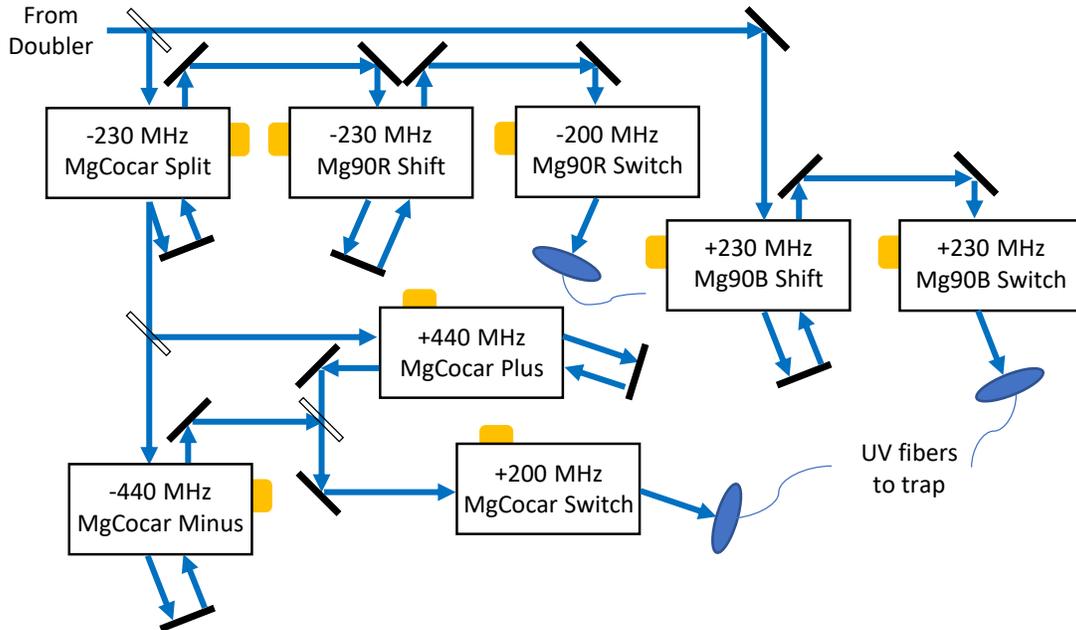


Figure 3.11: AOM setup used to generate our $^{25}\text{Mg}^+$ Raman beams. The Mg90R and Mg90B beams are first separated at a beamsplitter. Next, the Mg90R beam passes through the MgCocar Split AOM that splits off the MgCocar beam as the zeroth order and retroreflects the Mg90R as the minus first order. The Mg90R continues on through another doubler pass frequency shifter and then a single pass switch before coupling into an optical fiber to the trap. The Mg90B beamline passes through a single double pass frequency shifter, this time shifting the frequency up instead of down, then through its own single pass switch and fiber. The MgCocar beamline gets split into two arms by a beamsplitter, each of which passes through a double pass AOM to shift their frequencies in opposite directions. The two arms then get recombined on a beamsplitter, sent through a single pass switch AOM, and then coupled into a fiber.

Fig. 3.11 gives a sketch of the AOM system used to generate the different Raman beams. The UV doubler output gets split into the Mg90B and Mg90R beamlines by a beamsplitter. First, the Mg90R passes through the ~ 230 MHz MgCocar Split AOM where the MgCocar splits off as the zeroth order. The minus first order is retroreflected for a second pass, providing part of the frequency difference between the two motion sensitive Raman arms. It then continues on to a second double pass AOM driven at ~ 230 MHz (Mg90R Shift) to further shift the Mg90R frequency. Finally, it passes through a -200 MHz switch AOM before coupling into an optical fiber to the trap.

The Mg90B beamline passes through one ~ 230 MHz double pass AOM to shift its frequency, then also through a 200 MHz switch AOM before coupling into an optical fiber to the trap. This switch AOM can have two tones injected in it to drive MS interactions. The MgCocar Split, Mg90R Shift, and Mg90B Shift AOMs are all controlled by the same control system DDS and provide the frequency tuning for motion sensitive Raman transitions.

Bridging the ~ 1760 MHz gap between the $^{25}\text{Mg}^+$ qubit states is unfeasible with a single double pass AOM, so we take a different approach with the $^{25}\text{Mg}^+$ copropagating Raman beams than the $^9\text{Be}^+$ beams and instead generate the frequency difference with two double pass AOMs. Each provides a ~ 440 MHz shift on each pass, one using the plus first and the other the minus first order. These also provide frequency tuning when used for shelving, covering a frequency range from the $|3, 3\rangle \leftrightarrow |2, 2\rangle$ transition at approximately 1520 MHz to the $|3, -1\rangle \leftrightarrow |2, 0\rangle$ transition at approximately 1873 MHz. The final shelving steps from $|3, -1\rangle$ to $|2, -2\rangle$ and $|2, -1\rangle$ are done with global microwaves, which can still preserve spectator qubits since it does not involve either of the qubit levels. They then get recombined on a beamsplitter before going through a shared switch AOM before coupling into a fiber to the trap.

The two arms undergo significantly long spatial separation before recombining, unlike the $^9\text{Be}^+$ copropagating Raman beams, making them much more susceptible to phase noise from differential fluctuations. This could be problematic in long quantum algorithms, and

Beam	Cocar Sp.	CoCar -	CoCar +	CoCar Sw.	90 Sh.	90 Sw.	Co Sh.	Co Sw.
Mg90B	off	off	off	off	off	off	on	on
Mg90R	on	off	off	off	on	on	off	off
MgCocar	off	on	on	on	off	off	off	off

Table 3.4: $^{25}\text{Mg}^+$ Raman AOM switching configurations to turn on the various laser beams, as defined in Fig. 3.11. To turn on multiple beams simultaneously, take the logical OR within each column.

likely necessitate a more robust setup. Therefore they are only used when multiple $^{25}\text{Mg}^+$ ions carry information and individual addressing is needed. The only experiment in this thesis that uses these beams is the quantum gate teleportation experiments described in Chapter 5. There they were used for individually addressed shelving and conditional π -pulses, neither of which depend on the laser phase, provided it is stable over the duration of the pulse.

As with $^9\text{Be}^+$, each of these beams has a pickoff after the fiber to provide feedback for pulse shaping and noise eating (Fig. 3.5). Table 3.3 gives the conditions for turning on individual $^{25}\text{Mg}^+$ Raman laser beams. As with similar tables in Sec. 3.6, multiple beams can be switched on simultaneously by taking the logical OR of the relevant rows within each column.

3.8 Microwaves

Microwaves can be used to drive global single qubit rotations (Sec. 2.7). The microwave drive is generated by DDS controlled by the master FPGA (Sec. 3.5). With the hfgui system, a single DDS was used for both species, switching between two different lines for the two species. The $^9\text{Be}^+$ microwave line was then quadrupled, with appropriate filters. The $^{25}\text{Mg}^+$ microwave line was octupled.

The ARTIQ system has a dedicated DDS for each species to simplify the wiring. It also has a high clock frequency of 2.7 GHz, allowing direct generation of our required microwave frequencies without additional doubling. This is done for $^{25}\text{Mg}^+$ microwaves, however one of

our ${}^9\text{Be}^+$ shelving transitions ($|2, 0\rangle \rightarrow |1, -1\rangle$) has a transition frequency close to half the clock frequency (Sec. 2.2). To avoid having a nearby Nyquist image, for ${}^9\text{Be}^+$ microwaves we use a single doubling stage after the DDS.

Each microwave line then passes through a high-power amplifier and then a stub tuner to optimize impedance matching for the necessary range of transitions. Then, each of them goes to a dipole antenna near, but external to, the vacuum chamber. This comes with a few implications: First, the microwave field is global across the trap array, however not with a uniform Rabi rate. This means that it can only be used for uniform rotations when all ions are in the same trapping zone, or if ions in other zones do not have population in the transition being driven. Second, being in the far field the microwave gradient is negligible at the trap, meaning that they cannot be used to drive spin-motion coupling transitions. Finally, the copper resonator only has limited openings for optical access (Fig. 3.1), providing significant shielding of the microwave field. This means that we have to take a very brute force approach to delivering the microwaves to the ions with high power and likely taking an indirect path to the ions. For example, in microwave Ramsey experiments simply placing your hand in the wrong location near the trap can destroy the Ramsey signal. This usually is not a problem, except for when doing initial beam alignment by hand. An integrated microwave line in future traps could avoid all of these problems and allow for significantly faster microwave transitions [Ospelkaus et al., 2008].

Chapter 4

Quantum Logic Gates

In this chapter we present a brief introduction to important quantum states, gates applied to them, and notation for representing a quantum program/circuit. This provides the necessary background to understand the ideal behavior of experiments presented in subsequent chapters. We limit this discussion to pure quantum states, as that is all that is needed for this understanding. Noisy simulations, for example as presented in Fig. 6.7, were performed using the density matrix formalism, but the simulation details are not necessary to interpret their results. We refer readers interested in the density matrix formalism to Ref. [Nielsen and Chuang, 2002] for an introduction.

4.1 State Definitions

States and operators in quantum information processing are commonly represented in matrix form. Qubits are encoded in effective two-level systems, which hence require two basis states to fully represent. These are defined as

$$|\uparrow\rangle = |0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |\downarrow\rangle = |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (4.1)$$

Arbitrary superpositions of these two basis states can be written in the form

$$\cos(\theta/2) |\uparrow\rangle + e^{i\phi} \sin(\theta/2) |\downarrow\rangle \quad (4.2)$$

where θ is the polar angle from the positive Z-axis and ϕ the azimuthal angle from the positive X-axis of the Bloch Sphere shown in Fig. 4.1. Another important basis is defined along the X-axis of the Bloch sphere:

$$\begin{aligned} |\uparrow\rangle_X = |+\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\rangle + |\downarrow\rangle) = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{pmatrix} \\ |\downarrow\rangle_X = |-\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\rangle - |\downarrow\rangle) = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} \end{pmatrix}. \end{aligned} \quad (4.3)$$

Here the subscript X specifies the axis that $|\uparrow\rangle$ and $|\downarrow\rangle$ are referenced to. The computational Z basis is referenced implicitly by a lack of subscript. In our implementations of trapped ion qubits, the Z-axis of the Bloch sphere is given by the direction of the external magnetic field. Similarly, Y-axis basis states are given by

$$\begin{aligned} |\uparrow\rangle_Y = |r\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\rangle + i|\downarrow\rangle) = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ \frac{i}{\sqrt{2}} \end{pmatrix} \\ |\downarrow\rangle_Y = |l\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\rangle - i|\downarrow\rangle) = \begin{pmatrix} \frac{1}{\sqrt{2}} \\ -\frac{i}{\sqrt{2}} \end{pmatrix} \end{aligned} \quad (4.4)$$

where r and l are used in analogy to right and left handed circular polarizations of light. Operators that transform single qubit states are described in Sec. 4.2.

Systems of more than one qubit are represented by tensor products, for example

$$|\downarrow\rangle \otimes |\uparrow\rangle = |\downarrow\uparrow\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \quad (4.5)$$

where the elements of the vector correspond to $|\uparrow\uparrow\rangle$, $|\uparrow\downarrow\rangle$, $|\downarrow\uparrow\rangle$, and $|\downarrow\downarrow\rangle$ from top to bottom. Tensor products of n qubits are similarly represented by vectors of length 2^n , and operators on those systems as $n \times n$ matrices. See, for example, important multi-qubit logic gates in Sec. 4.3.

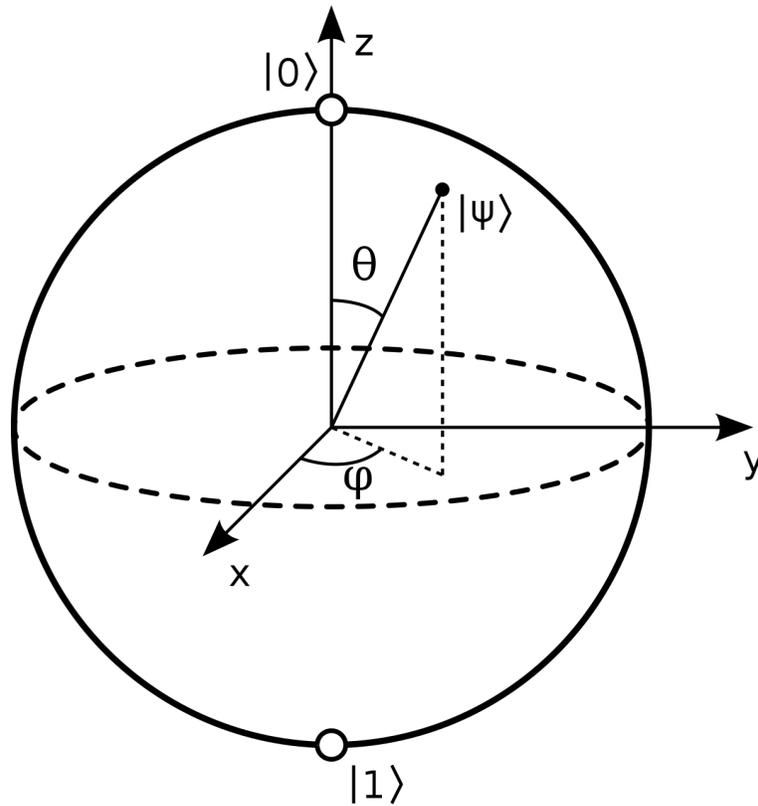


Figure 4.1: Bloch Sphere used to represent qubit states. Z basis eigenstates $|0\rangle = |\uparrow\rangle$ and $|1\rangle = |\downarrow\rangle$ align with the north and south poles, respectively (Eq. 4.1). State $|+\rangle$ aligns with the positive X axis (Eq. 4.3), and $|r\rangle$ with the positive Y axis (Eq. 4.4). Arbitrary states $|\psi\rangle$ are represented as $\cos(\theta/2)|\uparrow\rangle + e^{i\phi}\sin(\theta/2)|\downarrow\rangle$ where θ is the polar angle from the positive Z-axis and ϕ the azimuthal angle from the positive X-axis.

Multi-qubit states that cannot be factored into a tensor product of constituent single qubit states are said to be *entangled*. The four Bell states are entangled and form a common basis for two qubit states, defined as:

$$\begin{aligned} |\Phi^\pm\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\uparrow\rangle \pm |\downarrow\downarrow\rangle) \\ |\Psi^\pm\rangle &= \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle \pm |\downarrow\uparrow\rangle) \end{aligned} \quad (4.6)$$

Common entangled M -partite states include Greenberger-Horne-Zeilinger (GHZ) states [Greenberger et al., 1989]:

$$|GHZ\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle^{\otimes M} + |\downarrow\rangle^{\otimes M}), \quad (4.7)$$

where the exponentiated tensor product symbol corresponds to the tensor product of M qubits in $|\uparrow\rangle$ or $|\downarrow\rangle$, and W states that take the form:

$$|W\rangle = \frac{1}{\sqrt{M}} (|\uparrow\downarrow \dots \downarrow\downarrow\rangle + |\downarrow\uparrow \dots \downarrow\downarrow\rangle + |\downarrow\downarrow \dots \downarrow\uparrow\rangle). \quad (4.8)$$

Alternatively, W states can be defined similarly except with all but one qubit in $|\uparrow\rangle$.

4.2 Single-Qubit Gates

Single qubit gates are operators that act on individual qubits. Any 2×2 Hermitian matrix can be represented as a linear combination of the Pauli matrices

$$\begin{aligned} \sigma_0 = I &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\ \sigma_x = X &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\ \sigma_y = Y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \\ \sigma_z = Z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \end{aligned} \quad (4.9)$$

hence they are frequently used to describe error models and operators on qubits. The basis states described in Eqs. 4.1,4.3, and 4.4 are eigenstates of the Pauli matrix corresponding to their axis. Other important single qubit gates include the Hadamard gate:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (4.10)$$

which transforms between the X and Z bases, the single qubit phase shift gate S (sometimes instead called P)

$$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad (4.11)$$

and another phase shifting gate T

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}. \quad (4.12)$$

In the lab, we implement two types of single qubit rotations. The first rotates along axes in the equator of the Bloch sphere:

$$R(\theta, \phi) = \begin{pmatrix} \cos \frac{\theta}{2} & -ie^{-i\phi} \sin \frac{\theta}{2} \\ -ie^{i\phi} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix} \quad (4.13)$$

where θ is the angle rotated about an axis at an angle ϕ along the equator of the Bloch sphere. The second is simply a rotation about the Z-axis by an arbitrary angle ϕ :

$$R_Z(\phi) = \begin{pmatrix} e^{-i\phi/2} & 0 \\ 0 & e^{i\phi/2} \end{pmatrix}. \quad (4.14)$$

$R(\theta, \phi)$ can be implemented in a variety of ways, including microwaves, motion-sensitive Raman beams, or motion-insensitive Raman beams (see Chapter 2). For example, a microwave rotation described by the Hamiltonian in Eq. 2.7 implements $R(\theta, \phi)$ with θ governed by the interaction time and Rabi rate and ϕ by the corresponding phase. $R_Z(\phi)$ is most easily

implemented by shifting the phases of all subsequent single-qubit rotations by $-\phi$. Our core entangling gate is diagonal in the z-basis (Eq. 4.24), and hence commutes with these phase shifts, so its frame does not need to be updated.

Together, any of the single qubit gates discussed in this section can be generated from these two primitives [Maslov, 2017]. For example, S and T can be generated, up to irrelevant global phases, by R_Z with $\phi = \pi/2$ and $\pi/4$, respectively. Others, like H , require multiple of these primitives to generate (again up to a global phase):

$$H = R(\pi/2, -\pi/2) R(\pi, 0). \quad (4.15)$$

These primitives can also be used to perform the composite pulse sequences discussed in Sec. 2.8 to mitigate errors when mapping population during state preparation and measurement, for example:

$$R(\pi, 0) R\left(\pi, \frac{\pi}{3}\right) R\left(\pi, \frac{\pi}{6}\right) R\left(\pi, \frac{\pi}{3}\right) R(\pi, 0). \quad (4.16)$$

4.3 Multi-Qubit Gates

Harnessing the power of quantum computation requires a way for qubits to interact with each other. As with single-qubit gates, there is a common set of gates for multi-qubit operations. Perhaps the most common multi-qubit gate is the controlled-NOT (CNOT) gate

$$CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (4.17)$$

which, dependent upon the first qubit (control), applies X to the second qubit (target). Paired with single qubit gates, a CNOT can be used to produce Bell states from initially unentangled states. For example,

$$|\Phi^+\rangle = CNOT \cdot (H \otimes I) \cdot |\uparrow\uparrow\rangle. \quad (4.18)$$

The CNOT can also be generalized to involve more controls, for example in the Toffoli gate (also sometimes called CCNOT) where X is applied to the target if both controls are in $|\downarrow\rangle$:

$$\text{Toffoli} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}. \quad (4.19)$$

There are also gates that act symmetrically on the two qubits, for example the controlled-Z (CZ) gate is similar to a CNOT, except that the conditional operation on the target is Z instead of X :

$$\text{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (4.20)$$

This gate is symmetric, since no matter which qubit is labeled the control and which the target, the state $|\downarrow\downarrow\rangle$ picks up a minus sign with the other states untouched. Therefore, in circuit notation (Sec. 4.6) both qubits are assigned control symbols.

Another symmetric two-qubit gate is the SWAP gate

$$\text{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.21)$$

which swaps the states of the two qubits, the related $iSWAP$ gate that swaps the populations but also impacts the phase

$$iSWAP = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (4.22)$$

or the \sqrt{SWAP} gate

$$\sqrt{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1+i}{2} & \frac{1-i}{2} & 0 \\ 0 & \frac{1-i}{2} & \frac{1+i}{2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.23)$$

where $\sqrt{SWAP} \cdot \sqrt{SWAP} = SWAP$.

The symmetric phase gates G_+ and G_- are important gates that arise from our primitive entangling operation (Sec. 4.4):

$$G_+ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.24)$$

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

These two gates have the key features that they are diagonal in the Z -basis, allowing Z rotations to be commuted through them in a circuit, and that the impact on the qubits is symmetric and thus can be implemented with global operations.

4.4 Mølmer-Sørensen Interaction

Our primitive multi-qubit operation is the Mølmer-Sørensen (MS) interaction, which is realized by simultaneously driving symmetrically detuned red and blue sidebands with motion sensitive Raman beams [Sørensen and Mølmer, 1999, Sørensen and Mølmer, 2000]. This gives the interaction Hamiltonian

$$H_{MS} = \hbar \sum_j \Omega_{j,r} \sigma_{+,j} \left(a e^{-i(\delta_{j,r}t - \phi_{j,r})} \right) + \Omega_{j,b} \sigma_{+,j} \left(a^\dagger e^{i(\delta_{j,b}t + \phi_{j,b})} \right) + h.c. \quad (4.25)$$

where j is the ion number, $\Omega_{j,r/b}$ the red/blue sideband Rabi frequency on the j th ion, $\delta_{j,r/b}$ the detuning from sideband resonance, a and a^\dagger the motional creation and annihilation operators (Sec. 2.4), and $\phi_{j,r/b}$ the phases of the sideband drives. This displaces states in the X-basis (with the qubit frame defined by the motion-sensitive Raman beams) through different trajectories in motional phase space. For a system of two qubits, $|++\rangle$ and $|--\rangle$ are displaced while $|+-\rangle$ and $| -+\rangle$ are not (Fig. 4.2) due to counteracting spin-dependent forces. With constant δ and ω these trajectories will be circles in phase space that return to their starting point, forming closed loops whose area is proportional to a geometric phase acquired by $|++\rangle$ and $|--\rangle$.

By running the interaction for $t = 2\pi/\delta$ with $\Omega = \delta/4$ on two qubits, the following truth table is realized:

$$\begin{aligned} |\uparrow\uparrow\rangle &\rightarrow \frac{1}{\sqrt{2}} \left(|\uparrow\uparrow\rangle + e^{-i(\sum_j \frac{1}{2}(\phi_{j,r} + \phi_{j,b}) + \pi/2)} |\downarrow\downarrow\rangle \right) \\ |\uparrow\downarrow\rangle &\rightarrow \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle + |\uparrow\downarrow\rangle) \\ |\downarrow\uparrow\rangle &\rightarrow \frac{1}{\sqrt{2}} (|\uparrow\downarrow\rangle + |\uparrow\downarrow\rangle) \\ |\downarrow\downarrow\rangle &\rightarrow \frac{1}{\sqrt{2}} \left(|\uparrow\uparrow\rangle + e^{i(\sum_j \frac{1}{2}(\phi_{j,r} + \phi_{j,b}) + \pi/2)} |\downarrow\downarrow\rangle \right). \end{aligned} \quad (4.26)$$

From this we see that with the proper choices of initial state and laser phases, any of the Bell states except the singlet ($|\Psi^-\rangle$) can be produced with this interaction. However, unstable interferometric laser phases in our experiment result in $\phi_{j,r/b}$ that fluctuate slowly on the

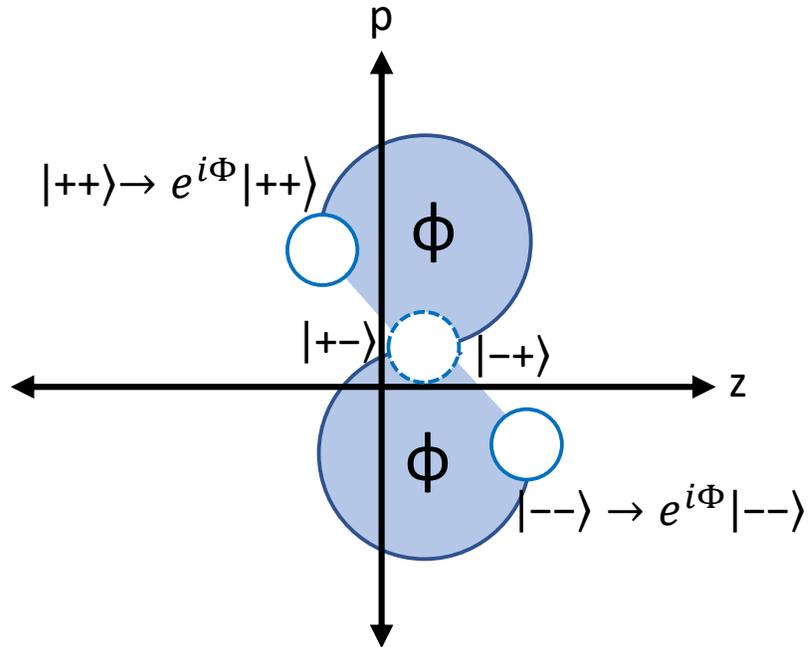


Figure 4.2: Sketch of harmonic oscillator phase space trajectories in a frame rotating at the oscillator frequency of a two ion system under the Mølmer-Sørensen interaction. Here $z = \sqrt{\frac{\hbar}{2m\omega}} (a^\dagger + a)$ and $p = i\sqrt{\frac{\hbar m\omega}{2}} (a^\dagger - a)$. The states $|++\rangle$ and $|--\rangle$ are displaced while $|+-\rangle$ and $|-+\rangle$ are not, causing the former pair to acquire a geometric phase proportional to the area Φ of the enclosed loop in phase space. The starting position is displaced from the origin to highlight insensitivity to the initial motional state.

time scale of the gate operation, so we use composite operations to controllably generate $|\Phi^\pm\rangle$ Bell states (Sec. 4.7). Population evolution for each of these four initial states under the ideal Hamiltonian (Eq. 4.25) is shown in Fig. 4.3.

Our choice of this type of entangling operation is motivated by several factors. First, it is insensitive to initial motional mode temperature to first order in the Lamb-Dicke approximation, which is necessary for high fidelity entangling gates. Second, we encode our ${}^9\text{Be}^+$ qubits in “clock” states that are insensitive to magnetic field fluctuations to first order (Sec. 2.1). This results in seconds long coherence times without any dynamical decoupling. The downside is that field-insensitive qubit states also experience negligible differential Stark shift, making simpler ZZ gates have similarly negligible coupling [Leibfried et al., 2003b]. Such light shift gates can instead be done with quadrupole transitions to D levels [Baldwin et al., 2021]. However, being in the second period of the periodic table, ${}^9\text{Be}^+$ has no easily accessible low lying D levels that this scheme depends on. The MS interaction is not dependent upon differential Stark shifts, and can thus be applied directly to clock qubits. Applying the MS gate to a mixed-species ion crystal works the same, except with the added constraint of needing to calibrate the phase of one species’ MS drive to align with the other [Tan et al., 2015].

For a system of two qubits and driven by motion sensitive Raman beams with an interferometric laser phase ϕ , at the points where the phase space trajectory returns to its starting point the MS interaction implements:

$$MS(\chi, \phi) = \exp\left(-i\frac{\chi}{2}\sigma_\phi \otimes \sigma_\phi\right) \quad (4.27)$$

where $\sigma_\phi = \cos(\phi)X + \sin(\phi)Y$ and χ depends on accumulated geometric phase. As a matrix

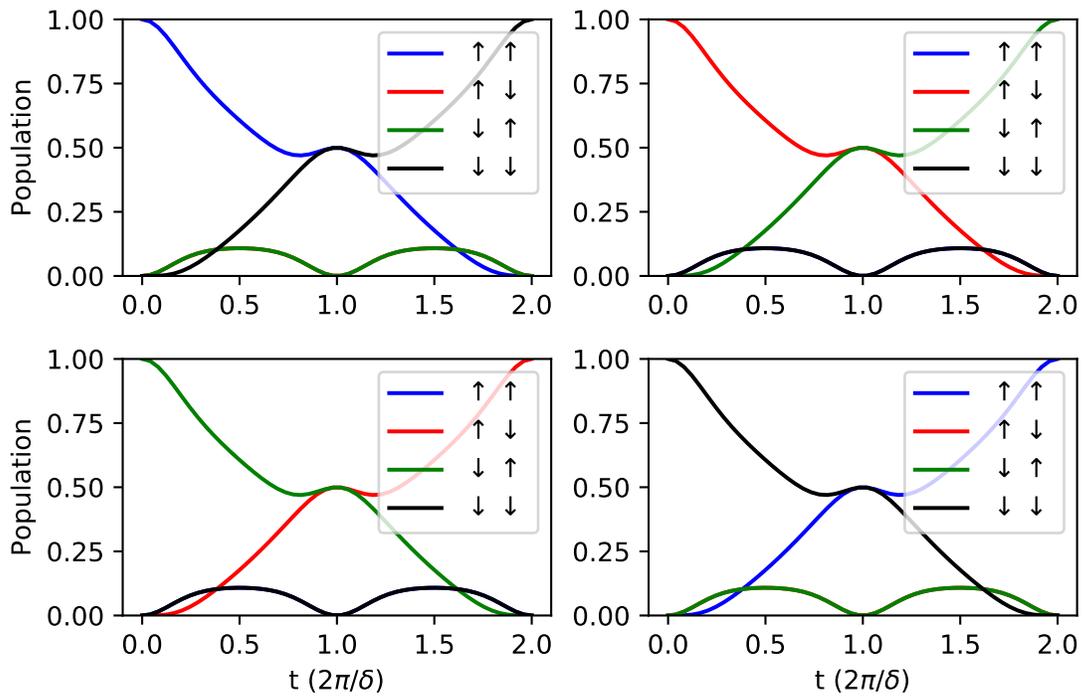


Figure 4.3: Time evolution of a system of two qubits under the Mølmer-Sørensen interaction beginning in $|\uparrow\uparrow\rangle$, $|\uparrow\downarrow\rangle$, $|\downarrow\uparrow\rangle$, or $|\downarrow\downarrow\rangle$.

this takes the form:

$$MS(\chi, \phi) = \begin{pmatrix} \cos(\chi/2) & 0 & 0 & -i e^{2i\phi} \sin(\chi/2) \\ 0 & \cos(\chi/2) & -i \sin(\chi/2) & 0 \\ 0 & -i \sin(\chi/2) & \cos(\chi/2) & 0 \\ -i e^{2i\phi} \sin(\chi/2) & 0 & 0 & \cos(\chi/2) \end{pmatrix}. \quad (4.28)$$

Here $\chi = \pi/2$ corresponding to the typical interaction time for Bell state generation ($t = 2\pi/\delta$) and $\chi = \pi$ corresponding to the “mutually-controlled multiflip” operation described in Sec. 4.5. Our definition differs from Ref. [Maslov, 2017] by including the interferometric laser phase ϕ and a factor of two in χ so that $MS(\pi)$ gives an effective “ π -pulse” from $|\uparrow\uparrow\rangle \leftrightarrow |\downarrow\downarrow\rangle$ and $MS(\pi/2)$ gives an effective $\pi/2$ -pulse (that generates a Bell state). This convention is useful for understanding how the multiflip can be used in Ramsey-type spectroscopy on entangled states as demonstrated in Sec. 6.1. In Sec. 4.7 we discuss how the MS gate can be combined with single qubit rotations to generate entangling gates that are insensitive to slowly fluctuating interferometric laser phases (on the time scale of a single composite gate) in our motion-sensitive Raman beams.

4.5 Mutually-Controlled Multiflip

If we instead run the MS interaction for twice the typical entangling gate time (or a single loop with twice the enclosed area), we observe an interesting feature where the initial state $|\uparrow\rangle^{\otimes n}$ returns to itself if n is odd, but flips to $|\downarrow\rangle^{\otimes n}$ if n is even [Leibfried, 2017, Kienzler et al., 2020]. This generalizes to any initial state with $|\psi\rangle \rightarrow |\psi\rangle$ if n is odd and $|\psi\rangle \rightarrow X^{\otimes n} |\psi\rangle$ if n is even (in a frame defined by the laser, setting $\phi = 0$). The initial state $|\uparrow\rangle^{\otimes n}$ was chosen for illustrative simplicity. If an even number of ions participate, the net effect is $X^{\otimes n}$. If an odd number participates, the net effect is $I^{\otimes n}$. For an even number participating, the matrix takes the form of a $2^n \times 2^n$ matrix with zeros everywhere except

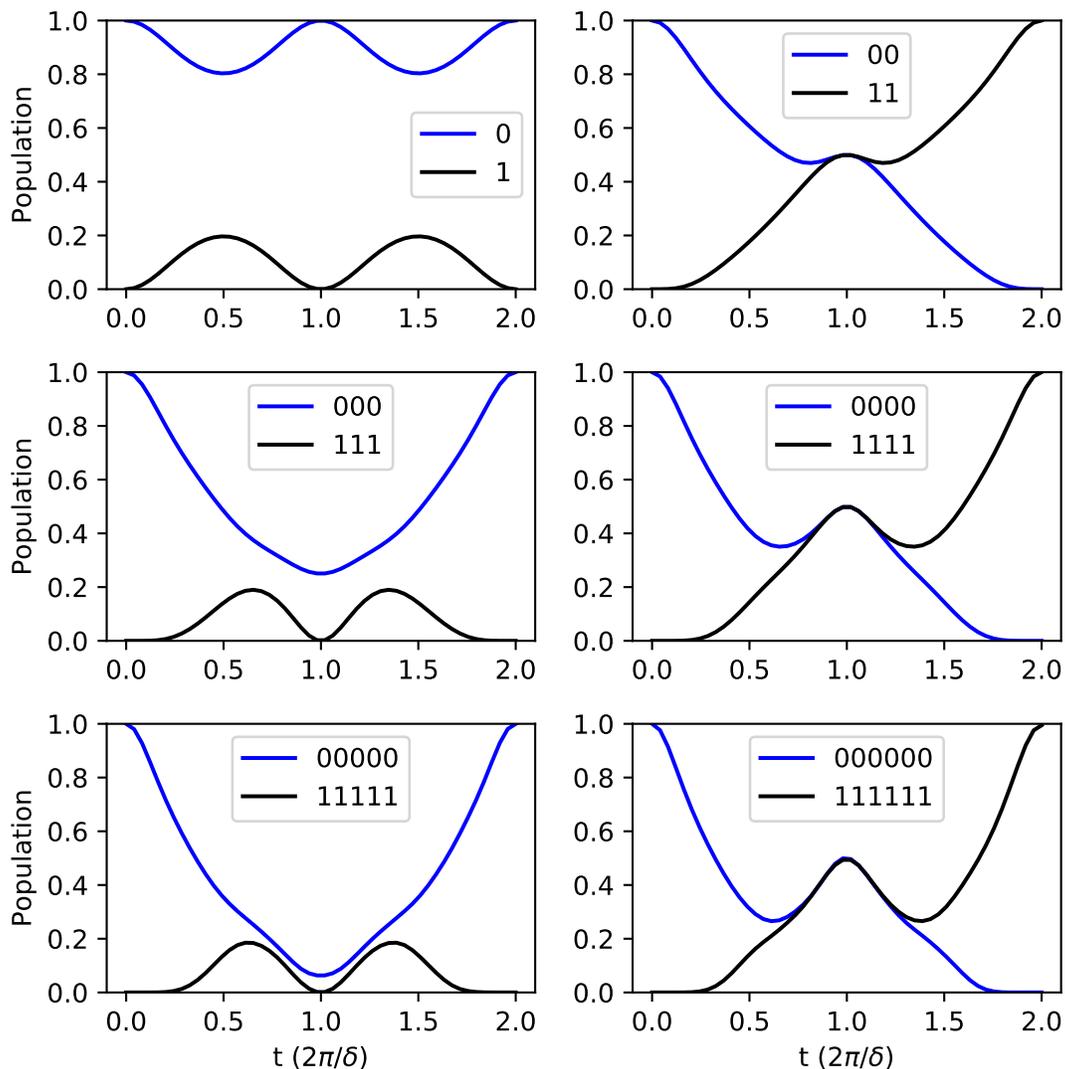


Figure 4.4: Time evolution of systems of $n = 1 - 6$ qubits under the Mølmer-Sørensen interaction. For simplicity only the states $|0\rangle^{\otimes n}$ and $|1\rangle^{\otimes n}$ are shown. Odd numbers of qubits (left column) return to their starting state $|0\rangle^{\otimes n}$ while even numbers of qubits (right column) flip to $|1\rangle^{\otimes n}$. The net effect is the identity for any odd number of ions participating in the interaction, and a global bit flip $X^{\otimes n}$ if the number is even, independent of the initial state. We call this operation the “mutually controlled multiflip” or just “multiflip” for short because a global bit flip is mutually conditioned on the number of ions participating in the interaction.

for the antidiagonal elements, which are all magnitude one. For two qubits, we get:

$$MS(\pi, \phi) = MS(\pi/2, \phi)^2 = \begin{pmatrix} 0 & 0 & 0 & e^{-2i\phi} \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ e^{-2i\phi} & 0 & 0 & 0 \end{pmatrix} \quad (4.29)$$

Implemented this way, the multiflip can be used to detect qubit leakage from the computational levels, as demonstrated in Ref. [Stricker et al., 2020]. In Chapter 6 of this thesis, we demonstrate how this operation can be used for quantum logic spectroscopy by scanning a common offset detuning for one of the species [Kienzler et al., 2020]. This inherits the temperature insensitivity of the MS interaction, and its entangling nature enables enhanced frequency sensitivity.

4.6 Quantum Circuits

Our composite gate operations are more easily visualized using circuit notation, which we briefly introduce in this section to better explain the experiments in later chapters. In quantum circuits, time proceeds from left to right and qubits occupy individual rows. Kets at the start of these rows signify initial states, and kets at the end signify output states. Single qubit operations are represented by boxes that span only single rows. For example, the following circuit prepares $|0+\rangle$ and applies $X \otimes Z$, resulting in $|1-\rangle$.

$$\begin{array}{ccc} |0\rangle & \text{---} \boxed{X} \text{---} & |1\rangle \\ |+\rangle & \text{---} \boxed{Z} \text{---} & |-\rangle \end{array} \quad (4.30)$$

Multi-qubit gates are represented by vertical lines joining various symbols on the involved qubits' rows. For example, the following circuit applies a CNOT, controlled-Z, and SWAP gates, respectively:

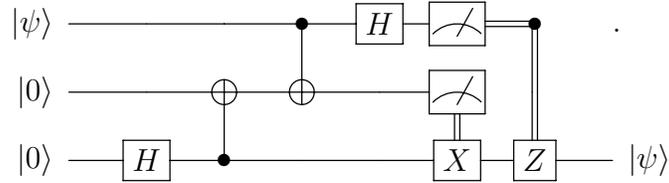
$$\begin{array}{c} \text{---} \bullet \text{---} \bullet \text{---} \times \text{---} \\ | \\ \text{---} \oplus \text{---} \bullet \text{---} \times \text{---} \end{array} \quad (4.31)$$

Here the top qubit is the control of the CNOT (solid dot) and the bottom qubit is the target (circled plus sign).

Measurements are represented by a small sketch of an analog dial, and classical bits of information corresponding to the result by a double line:


(4.32)

These classical bits can traverse rows of the circuit to condition subsequent operations, as can be seen in this example circuit for quantum state teleportation [Bennett et al., 1993]:


(4.33)

In this circuit, an arbitrary quantum state $|\psi\rangle$ is transferred from the first to the last qubit without the two ever directly interacting. This begins by entangling qubits 2 and 3 (H and CNOT) and sending half of the entangled pair to the first qubit's location. There, qubit 1 and 2 are measured in the Bell basis (CNOT and H before measuring both), with the measurement results used to condition single qubit rotations on qubit 3.

4.7 Composite Gates

To compensate for the fluctuating interferometric laser phases in our motion-sensitive Raman beams used to drive entangling operations (Sec. 3.6.2 and 3.7.2), we construct a number of composite gates. Under the assumption that the laser phase at the position of the ions is stable across the composite gate sequence, the composite gates are invariant to fluctuations between shots. Central to all other composite gates is the G_{\pm} phase gate, whose matrix form is listed in Eq. 4.24 and circuit shown in Fig. 4.5a [Lee et al., 2005, Tan et al., 2015]. In this gate, the MS interaction is surrounded by a Ramsey sequence implemented with the same motion-sensitive Raman beams. The two parts of the Ramsey sequence have phases of $\phi \pm \pi/2$ and $\phi \mp \pi/2$, respectively, where ϕ is the unknown interferometric laser

phase. This rotates the $\sigma_\phi\sigma_\phi$ MS interaction along the equator of the Bloch sphere to an effective $\sigma_z\sigma_z$ -type gate that is diagonal in the measurement basis and independent of the laser phase.

$$G = R\left(\frac{\pi}{2}, \phi \mp \frac{\pi}{2}\right)^{(\otimes 2)} \cdot MS(\pi/2, \phi) \cdot R\left(\frac{\pi}{2}, \phi \pm \frac{\pi}{2}\right)^{(\otimes 2)} \quad (4.34)$$

Note that the \pm in the rotation phases alone does not control which of G_\pm is implemented, which depends not only on these phases, but also on the sign of the MS detuning, and for mixed-species gates on the phase of each species drive in motional phase space (motion phase) [Tan et al., 2015, Wan et al., 2019]. To calibrate which of G_\pm we implement, we calibrate the motion phase for the MS gate within the composite gate sequence we are ultimately trying to implement, and for which we know the expected behavior, for example flipping the target in a CNOT. The wrong choice of G_\pm will produce the wrong behavior.

To generate entanglement with a deterministic phase, we surround a chosen G with single-qubit rotations on both ions. Any Bell state can be generated with the right choice of operations (requiring individual addressing in the case of $|\Psi^-\rangle$) and starting state. For example the circuit shown in Fig. 4.5 can be used to generate $|\Phi^+\rangle$ from $|\downarrow\downarrow\rangle$:

$$|\Phi^+\rangle = \left(R_z\left(\frac{\pi}{2}\right) \otimes I\right) \cdot R\left(\frac{\pi}{2}, 0\right)^{\otimes 2} \cdot G_- \cdot R\left(\frac{\pi}{2}, 0\right)^{\otimes 2} |\downarrow\downarrow\rangle \quad (4.35)$$

Similar circuits can be constructed to generate the other Bell states.

Finally, we can construct a CNOT gate (Eq. 4.17) by surrounding the target with single qubit rotations and R_Z rotations on both qubits. With the first qubit as the control and second as the target, this takes the form

$$CNOT = \left(R_z\left(\frac{\pi}{2}\right) \otimes R_z\left(-\frac{\pi}{2}\right)\right) \cdot \left(I \otimes R\left(\frac{\pi}{2}, \frac{\pi}{2}\right)\right) \cdot G_+ \cdot \left(I \otimes R\left(\frac{\pi}{2}, 0\right)\right) \quad (4.36)$$

as depicted in circuit diagram Fig. 4.5c.

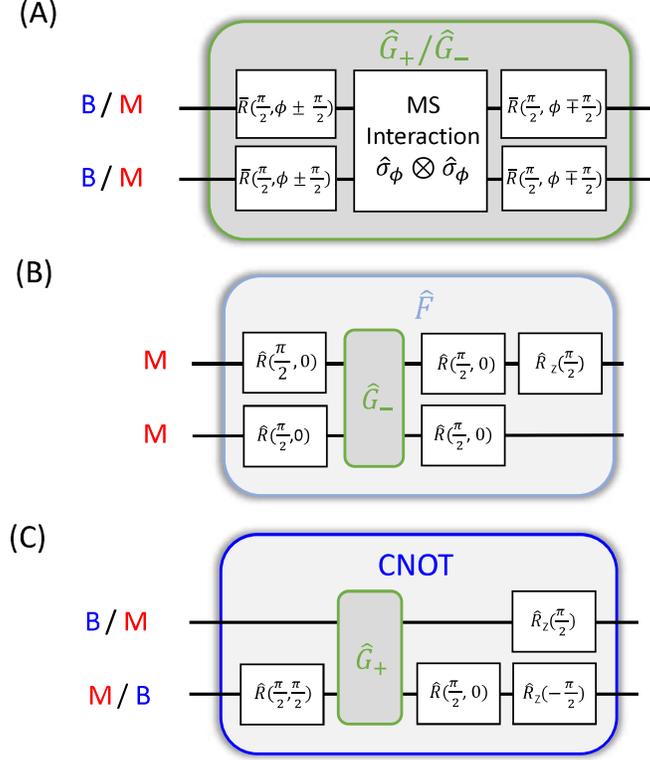


Figure 4.5: Circuits for composite quantum logic gates used in this thesis, reproduced from Ref. [Wan et al., 2019]. (A) To remove dependence on slowly fluctuating interferometric laser phases, we embed our MS interaction in a Ramsey sequence using the same set of lasers to produce phase gates that do not depend on this laser phase as long as it is stable over the composite pulse sequence. This results in the G_\pm operation in Eq. 4.24. The sign depends on the sign of the MS detuning, phases of the surrounding Ramsey sequence, and for mixed-species gates on the phase of the ions displacements in phase space (motion phase). (B) The phase on the $|\Phi^\pm\rangle$ type Bell states generated by the bare MS interaction similarly depends on the laser phase (Eq. 4.26), so we construct the composite operation F to generate $|\Phi^+\rangle$ if applied to $|\downarrow\downarrow\rangle$. (C) CNOT gates are constructed by surrounding G with single qubit rotations on the target, plus additional phase shifts.

4.8 Universal Gate Sets

Certain minimum requirements must be met to perform an arbitrary quantum circuit. For quantum logic gates these sets are called “universal gate sets” [Preskill, 1998]. There are many different sets, each of which must have at least one multi-qubit entangling gate and appropriate single qubit rotations. The CNOT, H , and T gates form one universal gate set. As discussed in Sec. 4.2 and Sec. 4.7, we can generate these gates with our R , R_z , and MS gates, so our capabilities fulfill the requirements for a universal gate set. It is important to note that a universal gate set is only the minimal set needed to implement an arbitrary quantum circuit. Greater flexibility will lead to less overhead in compiling a circuit, as is the case with our R and R_z rotations that can efficiently implement an arbitrary single qubit operation with their flexible rotation angles and phases. This thesis presents only a very basic introduction to these topics as needed to understand the ideal behavior of experiments presented in subsequent chapters. For more detailed information about quantum logic gates, John Preskill’s lecture notes [Preskill, 1998] or the textbook *Quantum computation and quantum information* by Michael Nielsen and Isaac Chuang [Nielsen and Chuang, 2002] are excellent resources.

Chapter 5

Quantum Gate Teleportation

In this chapter, we describe the implementation and results of our quantum gate teleportation demonstration published in Ref. [Wan et al., 2019].

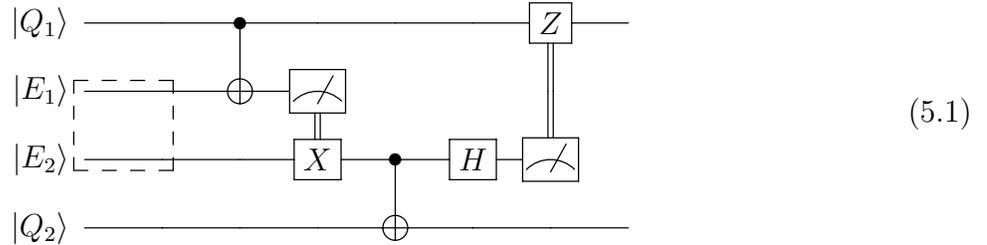
5.1 General Algorithm for Quantum Gate Teleportation

In classical computing, intensive tasks are often done not by faster and faster processors, but by distributing the task across a network of separate processing cores. It is likely that at some point similar modularization will be necessary to keep expanding the capabilities of quantum computation. As discussed in Sec. 4.8, distributing computation across multiple quantum processing nodes requires a two-qubit gate between qubits in separate nodes for universal quantum computation. Quantum gate teleportation accomplishes this without direct interaction between the two qubits, optimally requiring only local operations and classical communication (LOCC) and an entangled resource pair split between the two qubit locations (e.g. one of the Bell states listed in Eq. 4.6). It thereby implements a composite CNOT gate between separated qubits without requiring any direct interactions between them.

The circuit shown in Eq. 4.33 depicts the steps for quantum state teleportation, where an arbitrary unknown quantum state is transferred from one node to another using a split entangled pair and two classical bits of information as resources. Naively, one could imagine combining Alice and Bob's remote quantum processors in the following manner: Alice tele-

ports her qubit to Bob, Bob performs a two-qubit gate between the teleported qubit and one of his own, and then teleports Alice’s half of the output back to her. This accomplishes the goal of distributing quantum computation across two distinct processors, and in fact bears many similarities to the original circuit proposal for quantum gate teleportation put forth by Ref. [Gottesman and Chuang, 1999]. Their proposal has similar resource requirements and Bell-basis measurements on each side that condition single qubit rotations on the other, however it avoids the time ordering of the naive approach discussed above.

We will not go into the details of that proposal here, since a year later a better (and optimal) circuit was discovered by Ref. [Eisert et al., 2000]. This protocol requires only a single split entangled pair, one CNOT gate in each processor, and two bits of classical communication that trigger appropriate conditional single qubit rotations. The circuit representation of this protocol is:



where the middle two qubits E_1 and E_2 provide the resource entanglement represented by the dashed box, Q_1 is the control, and Q_2 the target of a teleported CNOT gate (see Eq. 4.17 for a matrix representation). Notably, this circuit is agnostic to how that initial entanglement was generated. It could be done through the typical unitary gates described in Sec. 4.7, photonic networks [Monroe and Kim, 2013], dissipative methods [Lin et al., 2013a], or Hilbert space engineering [Lin et al., 2016], among others.

The protocol can be understood as follows: Ion Q_1 is entangled with the resource pair by the first local CNOT. The first detection and conditional X gate serve to disentangle E_1 from Q_1 and E_2 . Hence the initial entanglement between E_1 and E_2 is transferred to Q_1 and E_2 . With information about Q_1 now shared with E_2 , Q_1 is the effective control of the second local CNOT with Q_2 the target. The remaining operations serve to disentangle E_2

from Q_1 and Q_2 , resulting in a composite CNOT with Q_1 as the control and Q_2 the target.

Early demonstrations of this circuit had been done probabilistically in photonic systems, where the conditional operations were implemented through passive optical elements and post selection [Huang et al., 2004, Gao et al., 2010]. More recently, a deterministic teleported CNOT gate was performed between two superconducting cavity qubits by means of an entangled resource pair of transmons, using either Fock state encodings or a logical encoding capable of detecting and correcting photon loss [Chou et al., 2018]. Deterministic quantum state teleportation was demonstrated previously with trapped ions [Barrett et al., 2004, Riebe et al., 2004], however gate teleportation has been long coming. In this chapter, we present our results on deterministic quantum gate teleportation between two ${}^9\text{Be}^+$ ions by means of an entangled pair of ${}^{25}\text{Mg}^+$ ions. Since the publication of these results, Honeywell Quantum Solutions has also performed quantum gate teleportation between ${}^{171}\text{Yb}^+$ ions in their prototype scalable QCCD processor [Pino et al., 2021].

5.2 Experimental Demonstration

We choose to use two ${}^9\text{Be}^+$ ions as our data qubits, named B_1 and B_2 , respectively. A pair of ${}^{25}\text{Mg}^+$ ions, M_1 and M_2 , serve as the entangled resource pair. These four ions are used in a crystal ordered ${}^9\text{Be}^+{}^{25}\text{Mg}^+{}^{25}\text{Mg}^+{}^9\text{Be}^+$, which we refer to as “BMMB”. In this configuration, the four axial motional mode frequencies are 1.4, 3.0, 4.1, and 4.2 MHz. Background gas collisions or laser drop outs can cause the ions to heat up and decrystallize, potentially scrambling their order once recrystallized. Different crystal orderings will have different normal mode frequencies, which in principle can be probed to verify the desired ordering. However, we find it simpler (and likely faster) to instead perform a deterministic ordering sequence before each shot of the experiment by distorting the trapping potential while cooling.

The key steps to this procedure are shown in Fig. 5.1. We begin by ramping up the axial confinement while simultaneously Doppler cooling on both species. With sufficiently

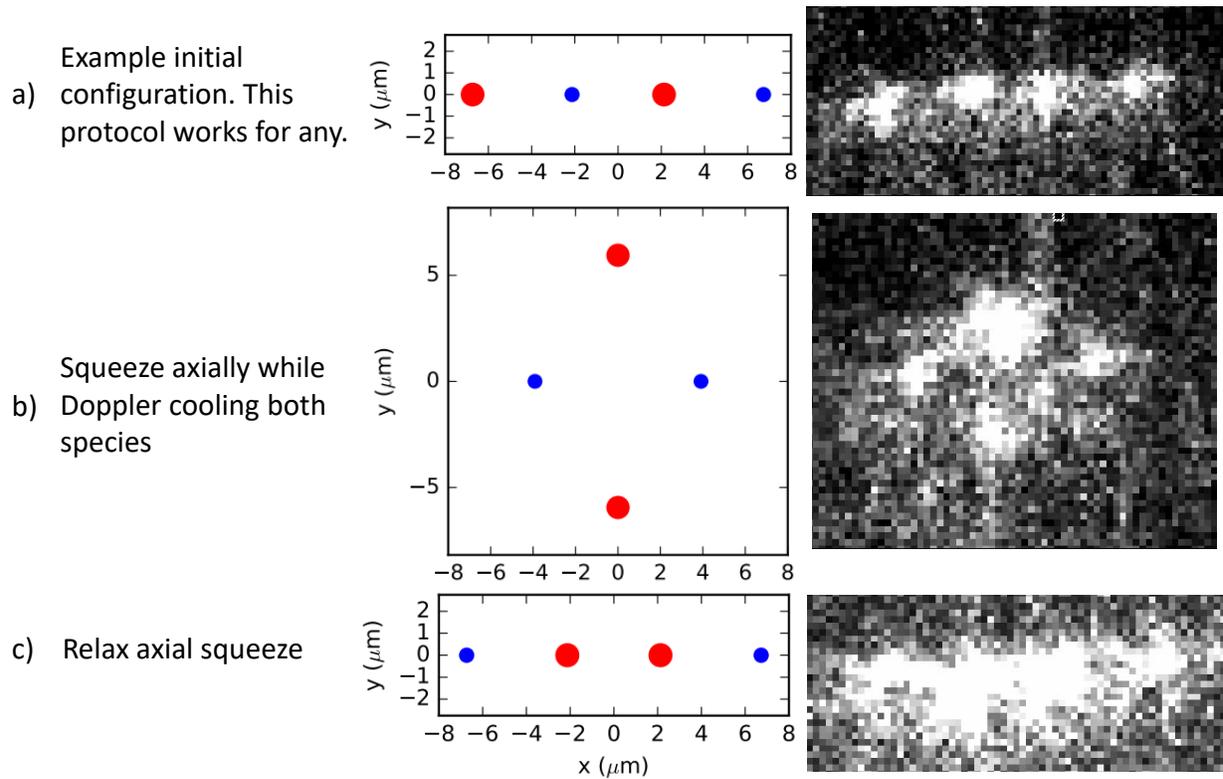


Figure 5.1: Procedure for initializing the order of a BMMB crystal. Beginning in any initial configuration (a), the axial confinement is ramped up while Doppler cooling is applied to both species. A linear crystal is no longer possible within this new potential, causing the two $^{25}\text{Mg}^+$ ions to move off axis due to their lower radial confinement (b). Adiabatically ramping the axial confinement back down while continuing to cool ensures that the $^{25}\text{Mg}^+$ ions remain in the middle, resulting in a consistent final ordering no matter the initial ordering.

strong axial confinement, the four ion crystal can no longer remain in a linear chain. Instead, it falls into a diamond shaped configuration (Fig. 5.1(b)). Here, the two ${}^9\text{Be}^+$ ions are on axis while the two ${}^{25}\text{Mg}^+$ ions off axis due to their weaker radial confinement (which scales inversely with mass). Adiabatically ramping the axial confinement back down, with the cooling still on, results in the two ${}^{25}\text{Mg}^+$ ions remaining in the middle. This procedure yields the same BMMB ordering, regardless of initial configuration.

5.2.1 Shuttling Sequence

Before describing the quantum operations in the circuit, we outline the separation and shuttling sequence used to enable individually addressed gates and detections (Fig. 5.2). Beginning in a single well, the ions are separated by ramping up the quartic term of the axial potential and then inverting the quadratic term [Home and Steane, 2006], resulting in a double well axial potential with one ion of each species in either well. The resulting crystals have axial modes frequencies of approximately 2.1 and 4.5 MHz. All of our laser beams for cooling, coherent control, and detection (Chapter 3) are focused to the laser interaction zone (LIZ) at electrodes 19 and 20 (see Fig. 3.2). The two smaller crystals are then sequentially shuttled back and forth in and out of the LIZ as needed to perform the quantum circuit (Fig. 5.4). Importantly, the two data qubits B_1 and B_2 never directly interact, demonstrating the “teleported” nature of the gate.

Ideally, the center of the quartic potential overlaps with the center of the BMMB crystal when the quadratic part goes through zero during its inversion, resulting in two ions falling into either potential well. In practice, stray fields and trap imperfections will typically require an additional electric field shim superimposed on the separation waveform to correct for differences between the waveform simulation and reality. To calibrate the separation, we begin by scanning an axial field shim superimposed on the separation waveform, then shuttle each well sequentially into the LIZ for detection of both species. Fig. 5.3(a) shows the results of such an experiment, where four different curves are shown corresponding to

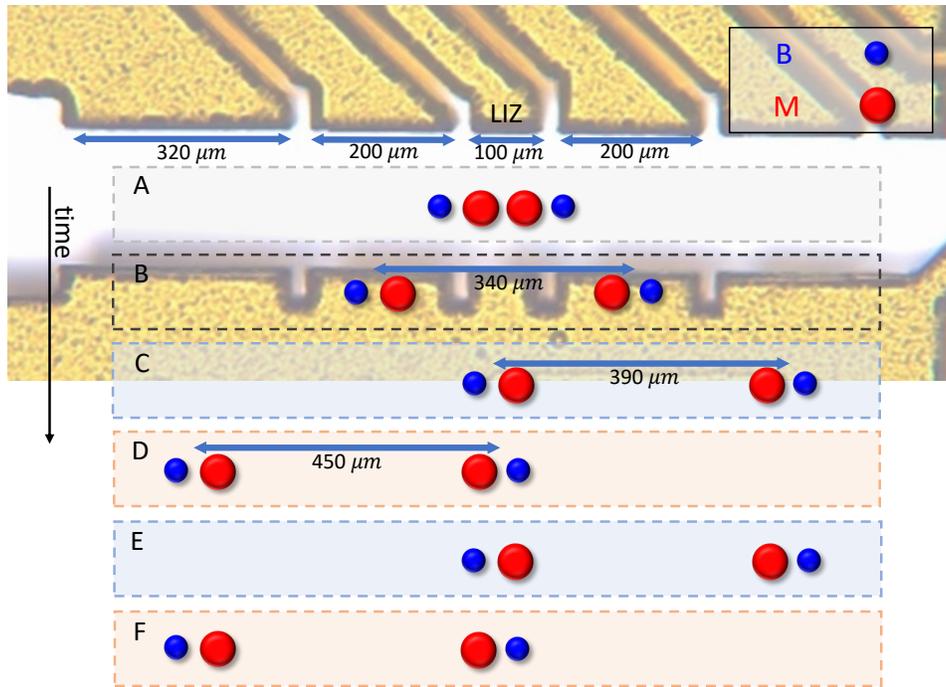


Figure 5.2: Shuttling sequence for quantum gate teleportation. After initially entangling the $^{25}\text{Mg}^+$ resource pair, the BMMB crystal is split into two mixed-species ion pairs by ramping up the quartic term of the potential and then inverting the quadratic part. All lasers are focused to the central narrow electrodes, referred to as the laser interaction zone (LIZ). The two pairs are then shuttled back and forth sequentially into the LIZ to perform the necessary elements of the quantum circuit (Fig. 5.4). Figure reproduced from Ref. [Wan et al., 2019].

photon fluorescence for each combination of species and potential well.

From left to right, we see that initially there are ~ 60 counts for the separate ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ detection windows when the right well has been shuttled into the LIZ, meaning that there are two of each species in that well and none in the other (each bright detection results in ~ 30 counts per ion). Around -0.008 V, the ${}^9\text{Be}^+$ fluorescence in the right well drops to 30 counts, and raise to a similar number in the left well. This indicates that the BMMB crystal has split into B and MMB. We then reach a region where each detection in either well yields 30 counts, meaning that there are one of each species of ion in both wells. Increasing the shim further after that results in three and then four ions in the left well. Drifting stray fields may require this shim to be recalibrated periodically.

Having the correct number of ions in each of the split wells does not necessarily mean that they did not experience significant motional excitation in the separation process. We do not require a precise characterization of the final temperature, just assurance that the crystal is reasonably cold. To accomplish this, when calibrating the separation shim we separate and recombine, and then probe one of the species with a motion-sensitive carrier π -pulse. Cold ions will give better contrast, resulting in a much narrower feature to calibrate to, shown superimposed on the fluorescence scan in Fig. 5.3(b).

When separating a mixed-species crystal, being radially off axis can result in twisting of the crystal and associated heating. Our trap has only a single bias electrode (Sec. 3.2), so we cannot exactly compensate in both wells, however we can do our best to split the difference. To reduce heating from being radially off-axis during separation, we use a similar procedure of separation, recombination, and carrier probe while scanning radial shims provided by the bias electrode and electrodes 19 and 20. These do not provide completely orthogonal shims, so iteration may be necessary.

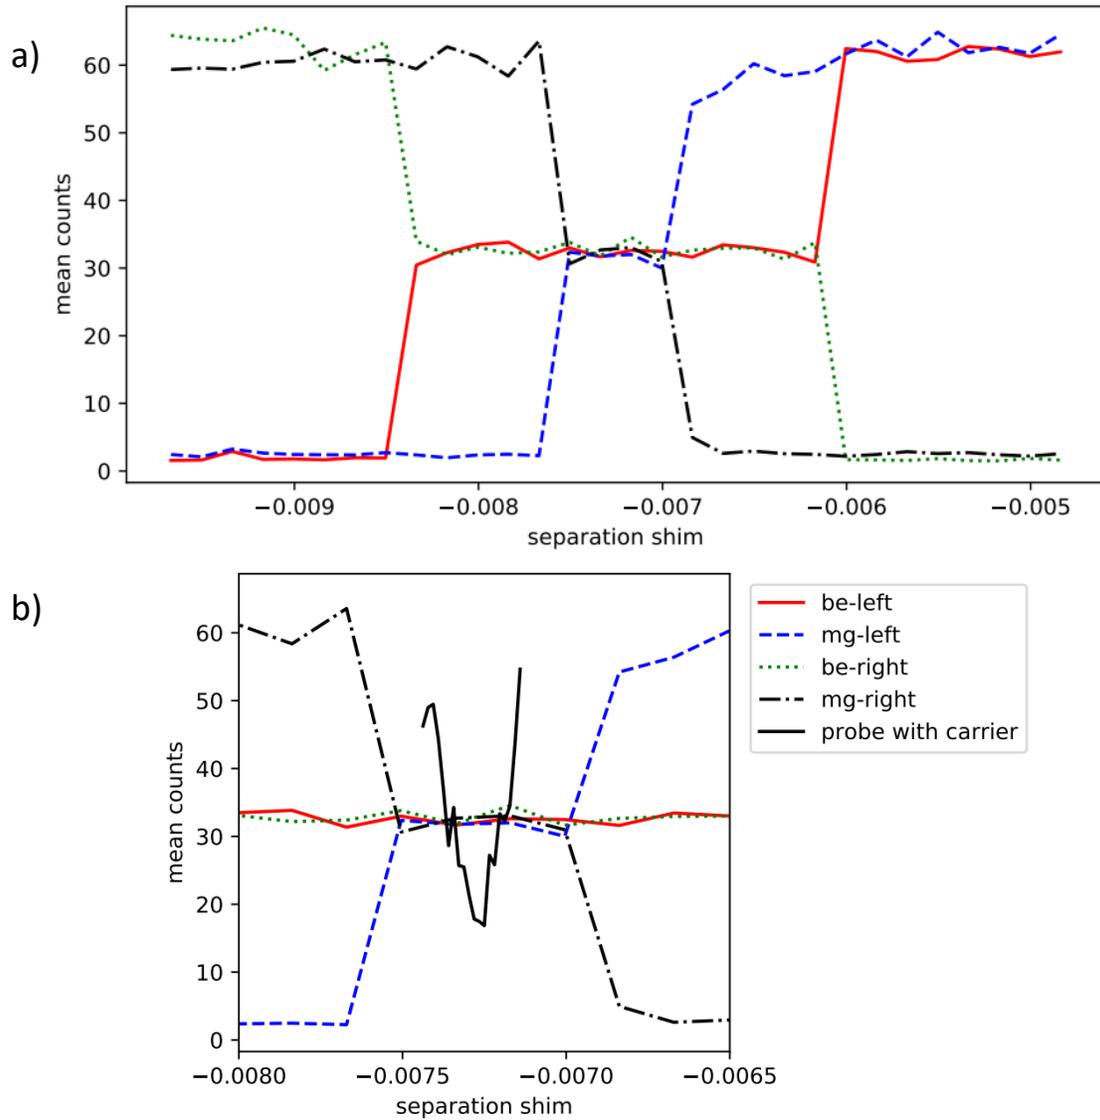


Figure 5.3: Scans for calibrating separation shims. (a) Photon fluorescence in each well for each species after separating and sequentially shuttling into the LIZ while scanning an axial shim superimposed on the separation waveform. Each ion gives roughly 30 counts, showing that the narrow region near -0.0072 results in one ion of each species in each well, as desired. (b) Result of ${}^9\text{Be}^+$ carrier π -pulse with the motion sensitive Raman beams after separating and recombining, giving a finer calibration that also is sensitive to final temperature (since it impacts the carrier flop contrast).

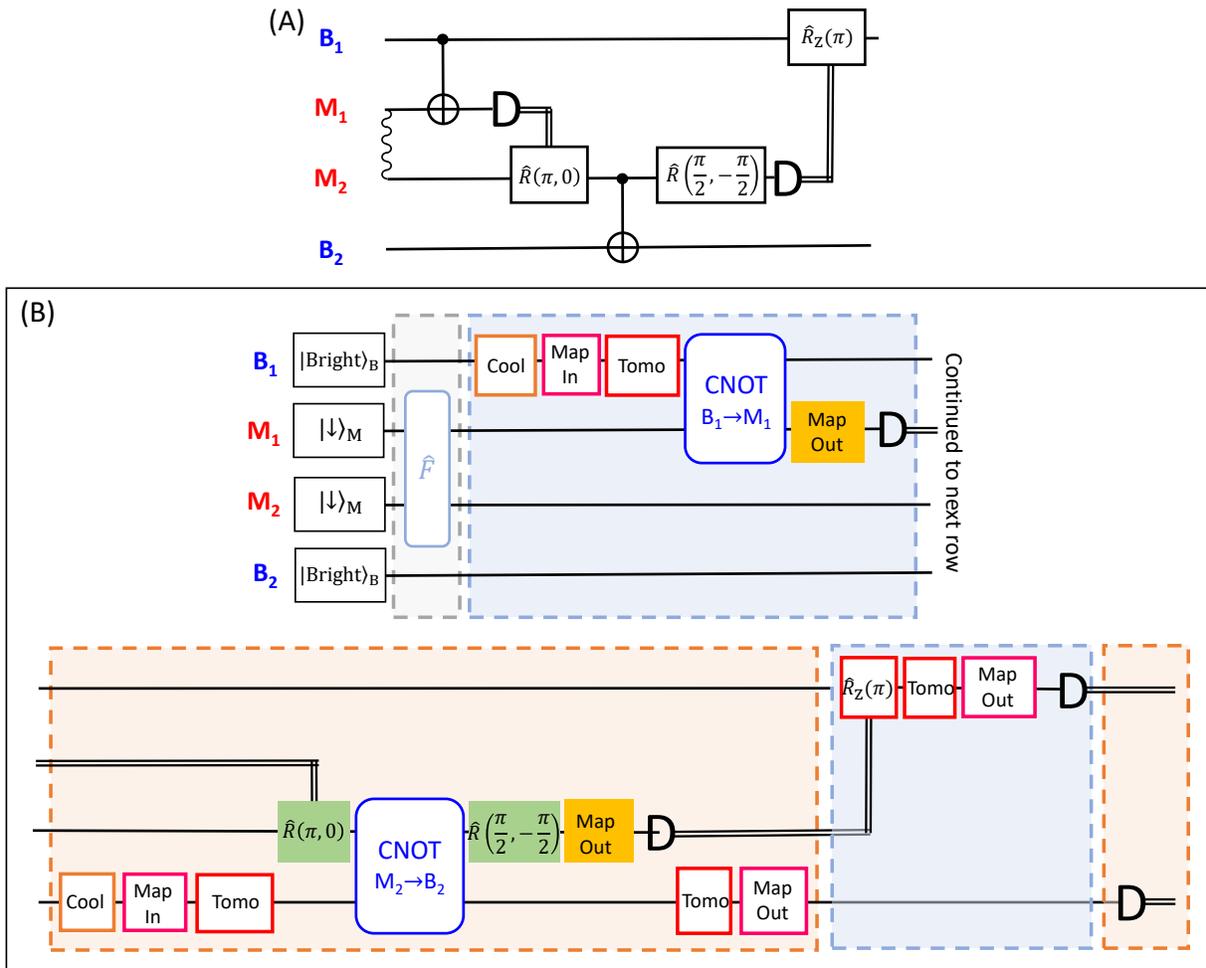


Figure 5.4: Circuit for a teleported CNOT gate. (a) General circuit, beginning with entanglement between M_1 and M_2 signified by the wavy line. The first local CNOT, detection of M_1 , and conditional feedforward operation on M_2 serve to transfer the entanglement between M_1 and M_2 to B_1 and M_2 . This makes B_1 the effective control of the second local CNOT with B_2 the target, the effective operation we want to realize. The remaining operations serve to disentangle M_2 from B_1 and B_2 . (b) Experiment-specific breakdown of how we implement the general circuit in part a). Here $|\text{Bright}\rangle = |2, 2\rangle$; \hat{F} the Bell state generating gate defined in Sec. 4.7; “Cool” Doppler and sideband cooling; “Map in” the process of mapping from the stretch state to the computational basis; “Tomo” a single qubit rotation corresponding to the tomography setting for the input state or measurement basis (Sec. 5.3); and “Map out” the processes of mapping from the computational basis to the measurement basis. Different shaded regions correspond to the different ion positions in the trap as shown in Fig. 5.2. Figure reproduced from Ref. [Wan et al., 2019].

5.2.2 Experiment-Specific Circuit

The exact circuit that we demonstrate differs slightly from that presented in Ref. [Eisert et al., 2000] and repeated in Eq. 5.1. Fig. 5.4(a) shows the circuit that we implement, and 5.4(b) a more detailed breakdown of how we implement it in our experiment. The main difference is that the Hadamard gate on M_2 is replaced with a $\frac{\pi}{2}$ -pulse. This yields equivalent outcomes because it is immediately followed by a detection, which is only sensitive to population and not the relative phase in a superposition. We use the two ${}^9\text{Be}^+$ ions as our data qubits, named B_1 and B_2 , respectively, and the ${}^{25}\text{Mg}^+$ ions M_1 and M_2 as the entangled resource pair.

The experiment begins by initializing the crystal order (Fig. 5.1), then Doppler cooling on both species followed by sideband cooling on ${}^9\text{Be}^+$ while all ions reside in the same potential well (Fig. 5.2(a)). To reduce cooling time while maintaining low final temperature, we first apply continuous sideband cooling, interleaving the four axial motional modes for several rounds. Then, we apply a short sequence of pulsed sideband cooling of ~ 5 pulses per mode. The two ${}^{25}\text{Mg}^+$ ions are then prepared as $|\downarrow\rangle$, while the ${}^9\text{Be}^+$ ions are left in $|\text{Bright}\rangle = |2, 2\rangle$ (see Fig. 2.1). A Bell state generating gate \hat{F} is used on the ${}^{25}\text{Mg}^+$ ions to map $|\downarrow\downarrow\rangle \rightarrow \frac{1}{\sqrt{2}}(|\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle) = |\Phi^+\rangle$. The BMMB crystal is then split into two mixed-species pairs (Fig. 5.2(b)), providing the necessary entangled pair split between two locations as the prerequisite for quantum gate teleportation.

We then shuttle the ions so that the left pair, B_1 and M_1 , is in the LIZ. After shuttling we ground state cool the crystal only on B_1 to avoid destroying the entanglement between the two ${}^{25}\text{Mg}^+$ ions. The initial state of B_1 is then prepared by first mapping into $|\uparrow\rangle$, followed by a single qubit rotation to the desired input state for process tomography. At this point, B_1 is ready for the gate teleportation circuit. Next, we perform the first local mixed-species CNOT gate, with B_1 as the control and M_1 the target. The composite sequence used to implement this gate is described in Sec. 4.7. M_1 is then shelved and measured, with the

photon counts compared against a preset threshold to determine the state. Individually shelving the $^{25}\text{Mg}^+$ ions was accomplished with a hybrid process. Mapping $|\downarrow\rangle \rightarrow |2, 2\rangle \rightarrow |3, 3\rangle \equiv |\text{Bright}\rangle$ is done with the copropagating Raman beams. Shelving to the dark state is also a two step process $|\uparrow\rangle \rightarrow |3, -1\rangle \rightarrow |2, -2\rangle$, however with the second step instead being done with microwaves to avoid needing too large of a frequency range in our copropagating Raman beams. The second step does not touch the qubit levels, so the only population in the other $^{25}\text{Mg}^+$ ion it affects is already lost to leakage error. The result of this detection is later used to condition a rotation on M_2 .

The ions are then shuttled such that the right pair, M_2 and B_2 , are in the LIZ (Fig. 5.2(d)). Again we cool only on $^9\text{Be}^+$ (B_2) to avoid collapsing the state of the $^{25}\text{Mg}^+$ ion (M_2). We then similarly map B_2 to the computational basis and apply a tomography pulse to prepare its input state. Next, a π -pulse is applied to M_2 if M_1 was measured to be in $|\downarrow\rangle$ before performing the second local mixed-species CNOT, this time with M_2 as the control and B_2 as the target. Following that, a $\frac{\pi}{2}$ -pulse rotates the M_2 measurement axis prior to detection. This measurement result is again compared to a threshold, which will later condition another operation. B_2 undergoes a tomography rotation to select the measurement axis and is mapped to the measurement basis (for reasons detailed in Sec. 5.5), but not yet detected.

The left pair is brought back into the LIZ (Fig. 5.2(e)). Here σ_z is applied to B_1 , conditioned upon M_2 being measured as $|\downarrow\rangle$. A tomography pulse then selects the measurement axis for B_1 before it is mapped to the measurement basis and detected. Finally, the right pair is brought into the LIZ (Fig. 5.2(f)) and detected. The gate teleportation experiment is at this point complete. The two $^{25}\text{Mg}^+$ ions are also detected, though only for diagnostic purposes, and the two crystals recombined at the LIZ to prepare them for the next shot of the experiment.

We characterized this teleported gate in a variety of ways, including full process tomography (Sec. 5.3) and consistency checks (Sec. 5.6).

5.2.3 Step Durations

The durations for each step of the experiment are given in Table 5.1, for a total runtime of approximately 21 ms. The vast majority of this time was spent on cooling, which can be improved by faster cooling techniques like EIT [Lin et al., 2013b] or polarization gradient cooling [Ejtemaee and Haljan, 2017]. Our transport was also done adiabatically for simplicity and was very far from perfect. Time spent transporting and time required for recooling can be reduced if care is taken to design transport potentials that separate and move ions quickly while avoiding motional excitation [Bowler et al., 2012, Todaro, 2020].

5.3 Full Process Tomography

In this section we outline the procedure for characterizing the teleported gate. This effort was led by Karl Mayer, who at the time was a graduate student in the group of Scott Glancy and Manny Knill at NIST.

The main way that we characterize the performance of the gate is through full quantum process tomography [Chuang and Nielsen, 1997]. Doing this requires an informationally-complete set of measurements, which we choose to form by preparing the ${}^9\text{Be}^+$ ions to the following set of input states:

$$\begin{aligned} &|\uparrow\uparrow\rangle, \quad |\uparrow\downarrow\rangle, \quad |\uparrow+\rangle, \quad |\uparrow r\rangle, \\ &|\downarrow\uparrow\rangle, \quad |\downarrow\downarrow\rangle, \quad |\downarrow+\rangle, \quad |\downarrow r\rangle, \\ &|+\uparrow\rangle, \quad |+\downarrow\rangle, \quad |++\rangle, \quad |+r\rangle, \\ &|r\uparrow\rangle, \quad |r\downarrow\rangle, \quad |r+\rangle, \quad |rr\rangle, \end{aligned}$$

where $|+\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$ and $|r\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + i|\downarrow\rangle)$ (see Sec. 4.1). The resulting outputs are then measured along each of the nine combinations of measurement axes: XX, XY, XZ, YX, YY, YZ, ZX, ZY, and ZZ (see Fig. 4.1) by rotating bases from X/Y-axes to the Z-axis with appropriate $\frac{\pi}{2}$ -pulses prior to fluorescence measurement. The ${}^{25}\text{Mg}^+$ ions are a fixed resource internal to the process and necessarily always prepared to ideally the same state,

Step	Description	Operation	Duration
	Optical pumping and crystal initialization of the four-ion chain	-	3.2 ms
A	DC on ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ and SP of ${}^{25}\text{Mg}^+$ ions	-	1.3 ms
	SBC on ${}^9\text{Be}^+$ ions	-	5.3 ms
	M_1 - M_2 Bell-state generation	\hat{F}	220 μs
B	Separation of four-ion chain into B_1 - M_1 and M_2 - B_2 in double well	-	570 μs
	Shifting double well	-	230 μs
C	DC, SBC, and SP of B_1	-	2.2 ms
	CNOT on B_1 - M_1	CNOT	280 μs
	Map out and detection of M_1	-	650 μs
	Shifting double well	-	460 μs
	Cooling and SP of B_2	-	2.2 ms
D	Conditional rotation on M_2	$\hat{R}(\pi, 0)$	25 μs
	CNOT on B_2 - M_2	CNOT	280 μs
	Rotation on M_2	$\hat{R}(\pi/2, -\pi/2)$	15 μs
	Map out and detection of M_2	-	650 μs
	Map out B_2	-	220 μs
	Shifting double well	-	460 μs
	DC on M_1	-	200 μs
E	Conditional phase shift on B_1	$\hat{R}_Z(\pi)$	100 ns
	Map out and detection of B_1	-	540 μs
	Detection of M_1	-	180 μs
	Shifting double well	-	460 μs
	DC on M_2	-	200 μs
F	Detection of B_2	-	300 μs
	Detection of M_2	-	180 μs
	Recombination	-	800 μs

Table 5.1: Approximate durations of each step in our quantum gate teleportation demonstration. The following acronyms are used: Doppler cooling (DC), state preparation (SP), and sideband cooling (SBC). As we can see, most of the time is spent on shuttling and associated recoiling. The ${}^{25}\text{Mg}^+$ detections in steps E and F are for diagnostics and are not part of the quantum gate teleportation algorithm. Table reproduced from Ref. [Wan et al., 2019].

their degrees of freedom do not appear in the process tomography.

Together, this gives $16 \times 9 = 144$ different combinations of input states and measurement axes which we call “tomography settings.” Roughly 300 data points were taken for each setting, ordered randomly, in each of two separate datasets. We used the first dataset to test our analysis techniques and establish a standard to be used without modification on the second, which we report as our main result presented below. With experiment maintenance and calibrations, each dataset took roughly two days to acquire.

Because of this large time requirement, several other intermediate checks were used when building up to it. Subsets of the circuit were checked along the way, for example, the circuit can be cut short after the conditional π -pulse on M_2 to check for entanglement between B_1 and M_2 . Later we checked that the right choice of input state (e.g. $|+0\rangle$) produced the expected entanglement in the full process. Once we had enough confidence in the gate to proceed with the full quantum process tomography, we first tested it on a smaller set of measurements to bound the fidelity in the following manner: In a set of eight different measurements, we tested the action of the gate on each combination of Z- and X-basis eigenstates. With the first qubit B_1 as the control, the expected behavior is:

$$\begin{aligned} |\uparrow\uparrow\rangle &\rightarrow |\uparrow\uparrow\rangle, & |++\rangle &\rightarrow |++\rangle, \\ |\uparrow\downarrow\rangle &\rightarrow |\uparrow\downarrow\rangle, & |+-\rangle &\rightarrow |--\rangle, \\ |\downarrow\uparrow\rangle &\rightarrow |\downarrow\downarrow\rangle, & |-+\rangle &\rightarrow |-+\rangle, \\ |\downarrow\downarrow\rangle &\rightarrow |\downarrow\uparrow\rangle, & |--\rangle &\rightarrow |+-\rangle. \end{aligned}$$

If we label the average success probability in the Z and X basis f_1 and f_2 , respectively, then the average process fidelity is bound by [Nielsen, 2002, Pino et al., 2021]:

$$F_{avg} \geq \frac{4}{5}(f_1 + f_2) - \frac{3}{5}. \tag{5.2}$$

This bound is for the average fidelity, which is a slightly different metric than the entanglement fidelity that we report as our main result of this experiment. The two are related by

[Nielsen, 2002]:

$$F_{avg} = \frac{dF_e + 1}{d + 1} \quad (5.3)$$

where d is the dimension of the computational space (in this case $d = 4$).

Here we will only give a brief overview of the analytical methods used during the full quantum process tomography. Details of the methodology can be found in the supplementary materials of Ref. [Wan et al., 2019].

Reference histograms were taken periodically throughout data acquisition to account for imperfect measurements. Input states were assumed to be perfect. Based on the measured reference histograms, maximum likelihood estimation [Fiurášek, 2001] was used to determine the process, whose entanglement fidelity was then computed with respect to the ideal CNOT. This yielded point estimates of $F_{ML,1} = 0.858$ and $F_{ML,2} = 0.851$ for the two data sets. The inferred process for dataset 2 is shown in Fig. 5.5 in the Pauli transfer matrix representation. This representation shows how an arbitrary input density matrix, expressed as a linear combination of Pauli products, is mapped by the process to a similarly constructed output density matrix.

These point estimates say nothing about the uncertainty in these measurements, so we used parametric bootstrapping to obtain confidence intervals. The inferred process was used to randomly generate 2000 synthetic data sets which then had their own processes inferred through maximum likelihood estimation and fidelities again computed with respect to the ideal CNOT. The resulting distribution for dataset 2 is shown in Fig. 5.6 in blue. There is a noticeable downward bias in the mean of the bootstrapped samples compared to the point estimate (red line). This implies a bias in the maximum likelihood estimation procedure, so to approximately correct for the bias the endpoints of the 95% confidence interval (black dashed lines) are reflected about the point estimate. All together, this resulted in 95% confidence intervals of $[0.852, 0.878]$ and $[0.845, 0.872]$ for datasets 1 and 2, respectively. Since dataset 2 was evaluated with a procedure that was fixed in advance by training on

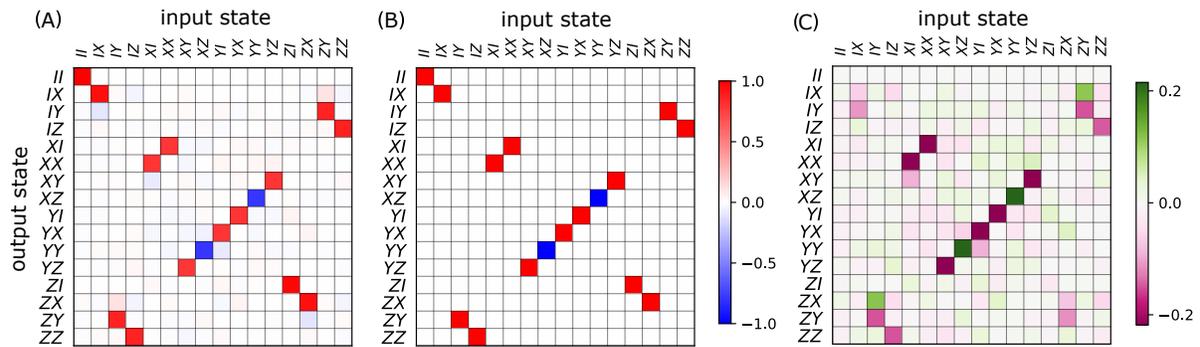


Figure 5.5: Visualization of our remote CNOT gate in the Pauli transfer matrix representation. This representation takes an arbitrary input density matrix, expressed as a linear combination of Pauli products, to an output density matrix expressed similarly. Experimental data (a) compared to an ideal CNOT (b) with their difference in part (c). This data corresponds to a 95% confidence interval of (0.845, 0.872) for the entanglement fidelity.

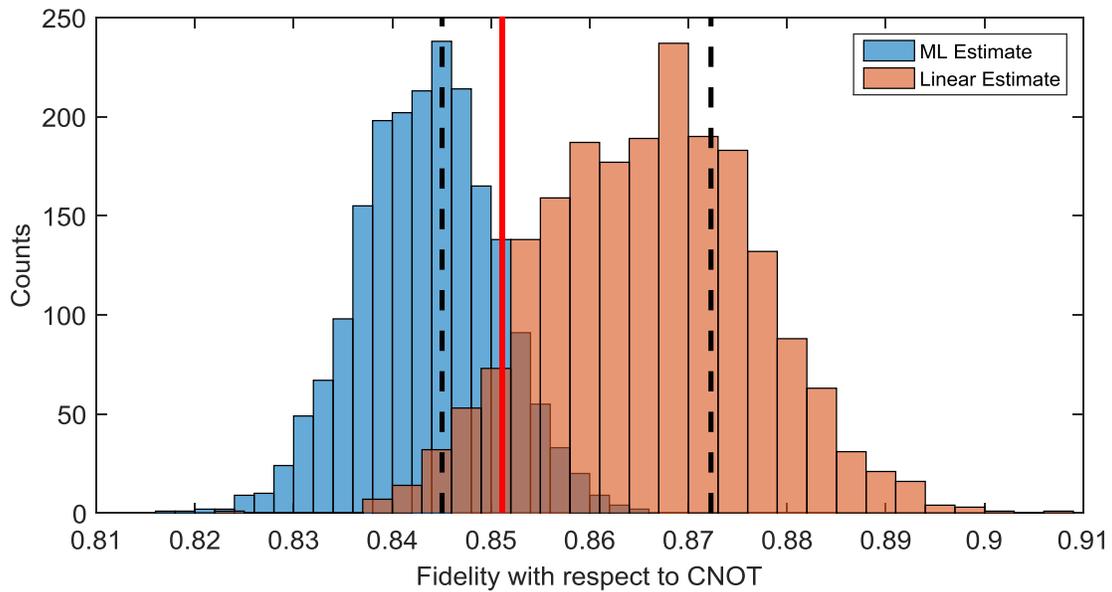


Figure 5.6: Bootstrapped fidelity histograms for our remote CNOT gate demonstration, dataset 2. Maximum likelihood estimation is used to compute a point estimate (red line). Parametric bootstrapping on dataset 2 produces the blue histogram. The bias corrected 95% confidence interval is shown by the dashed black lines. Finally, a linear estimator is used on dataset 2 as a consistency check, with the distribution of non-parametric bootstrapped results generated independently from that of the maximum likelihood bootstrap set shown in orange. Figure reproduced from [Wan et al., 2019].

dataset 1, we consider the results from dataset 2 to be the official estimate of the process fidelity and its uncertainty.

To check the maximum likelihood estimation for consistency with a different evaluation of the data, we compared the results to those obtained by a direct linear estimate of the observed fidelity. Likewise, we generated an additional 2000 non-parametric bootstrapped synthetic data sets to estimate the uncertainty in this estimate and plot their distribution in orange with the maximum likelihood results in Fig. 5.6. The two confidence intervals are consistent with each other, instilling confidence that the maximum likelihood estimation is valid.

5.4 Error Budget

To verify that our teleported CNOT gate works as well as would be expected from its constituent operations, we construct an error budget of the dominant sources shown in Table 5.2. The main sources that we consider are state preparation and measurement (SPAM) errors on both species, error in entangling the two $^{25}\text{Mg}^+$ ions, error in each of the mixed-species CNOTs, coherence of the $^{25}\text{Mg}^+$ ions (decoherence of the $^9\text{Be}^+$ ions is negligible at our chosen magnetic field, see Sec. 2.1), stray resonant light scattering on M_2 and corrupting its state when detecting M_1 , and stray light from cooling on B_2 resonantly scattering on B_1 . If all errors are treated as independent, their collective effect sums up to 16(2)%. This is not necessarily the case, so to get a different, more sophisticated estimate we compute their net effect using a depolarizing model. For a constituent operation \hat{U}_i and associated error rate ϵ_i , the density matrix ρ after that operation transforms as

$$\rho \rightarrow (1 - \epsilon_i) \hat{U}_i \rho \hat{U}_i^\dagger + \epsilon_i \cdot \hat{I}/d_i \quad (5.4)$$

where d_i is the dimension of the computational Hilbert space for the process and \hat{I} the identity. For a two-qubit gate, $d_i = 4$, while for a single qubit rotation $d_i = 2$.

The full algorithm is stepped through with the depolarizing model applied at each step,

Table 5.2: Leading Imperfections for the remote CNOT gate, estimated from measurements that were independent of the datasets used in the process tomography, with 1σ -uncertainty shown in parentheses. The net effect of all sources is estimated through both a simple sum and a depolarizing model (Eq. 5.4) and agrees with the process fidelity within their mutual uncertainties. Table is reproduced from Ref. [Wan et al., 2019].

Source	Error (10^{-2})
SPAM on two ${}^9\text{Be}^+$ ions	1.1(7)
SPAM on two ${}^{25}\text{Mg}^+$ ions	1.5(3)
M_1 - M_2 Bell state	4.0(9)
B_1 - M_1 CNOT	3.0(9)
M_2 - B_2 CNOT	3(1)
Coherence of ${}^{25}\text{Mg}^+$ ions	0.7(3)
Stray light from M_1 detection on M_2	1.1(4)
Stray light from B_2 cooling on B_1	1.2(3)
Sum	16(2)
Depolarizing model	12(1)

taking ϵ_i as the error estimates from Table 5.2. With a complete set of input density matrices (e.g. each possible Pauli product), we can then compute the fidelity of the process with respect to the ideal CNOT. In so doing, we find the total error under the depolarizing model to be 12(1)%, which is consistent with the measured confidence interval using maximum likelihood estimation. This gives us further confidence that our major error sources are accounted for, and that the teleported gate performed as well as could be expected given the known and most detrimental underlying error sources.

We now outline how we determined each of these constituent error rates, from top to bottom in Table 5.2. SPAM errors mostly arise from polarization impurities in the resonant optical pumping and detection beams, and errors in the single qubit rotations mapping between the computational and measurement bases. The error for each species is determined from a set of two experiments. In the first, we prepare a single ion in $|\uparrow\rangle$ and then measure the probability of detecting $|\downarrow\rangle$. The second is the opposite, where we prepare $|\downarrow\rangle$ and measure the probability of detecting $|\uparrow\rangle$. The SPAM error per ion is then taken to be the mean of these two probabilities. In our setup, that gave a total of 1.1(7)% SPAM error for the two ${}^9\text{Be}^+$ ions and 1.5(3)% SPAM error for the two ${}^{25}\text{Mg}^+$ ions.

Imperfection in preparing the ${}^{25}\text{Mg}^+$ Bell state is estimated through the standard method of measuring the population after the entangling gate and the amplitude of the observed parity oscillation when inserting a $\frac{\pi}{2}$ -pulse with a variable phase after the gate [Sackett et al., 2000]. Care has to be taken when correcting for SPAM errors in these experiments. Each qubit contributes 3/2 times their individual SPAM errors to these joint measurements, as derived in Ref. [Ballance, 2014]. Subtracting out the appropriately-weighted SPAM error then yields the estimated Bell state error, which for the two ${}^{25}\text{Mg}^+$ ions was 4.0(9)%.

Estimating the error in the two mixed-species gates would best be done through full process tomography as discussed above, however we settled for an easier to implement estimate for the purpose of this budget. Instead, we use the gates to prepare Bell states and determine populations and parity as discussed above to estimate the CNOT gate fidelity.

In so doing, we estimate errors of 3.0(9) and 3(1)% for the two mixed-species CNOT gates, respectively.

Since the two $^{25}\text{Mg}^+$ ions wait different amounts of time before their states are detected in the experiment, directly accessing their decoherence in a simple experiment is not possible. Instead, we find the $1/e$ coherence time of a single $^{25}\text{Mg}^+$ ion to be 140(30) ms and use that to infer the decoherence error. The uncertainty is primarily to account for drifts in day to day performance, not statistical uncertainty in the measurement. The two ions spend 4.2 ms entangled (see Table 5.1). M_2 still remains entangled with B_1 and B_2 for an additional 3.6 ms after M_1 is measured. The Ramsey contrast of the Bell state should decay at twice the rate of the constituent particles, so we estimate the decoherence error to be equivalent to the loss of Ramsey contrast from a single $^{25}\text{Mg}^+$ ion after $2 \times 4.2 + 3.6 = 12$ ms. From our measured coherence time of 140(30) ms, we infer a total decoherence error of 0.7(3)% for the $^{25}\text{Mg}^+$ ions.

In this way we independently estimate the main errors of the implementation that were anticipated. In the next section we will outline the unexpected impact of stray light scattering.

5.5 Scattered Resonant Light

During our demonstration of quantum gate teleportation, there are two main periods where scattered resonant light used for detection or cooling of a certain ion can prematurely corrupt information encoded in the other ions of the same species (i.e. indirect depumping error). First, detecting M_1 , and second cooling on B_2 . In larger registers running more complex algorithms, we expect this kind of error to also be present, and careful mitigation strategies will be necessary to render it insignificant.

To determine the effect of scattered resonant light during gate teleportation, we perform a series of test experiments shown in Fig. 5.7. The first two, parts (a) and (b), are reference experiments where the ions are prepared as for gate teleportation, and then either pair

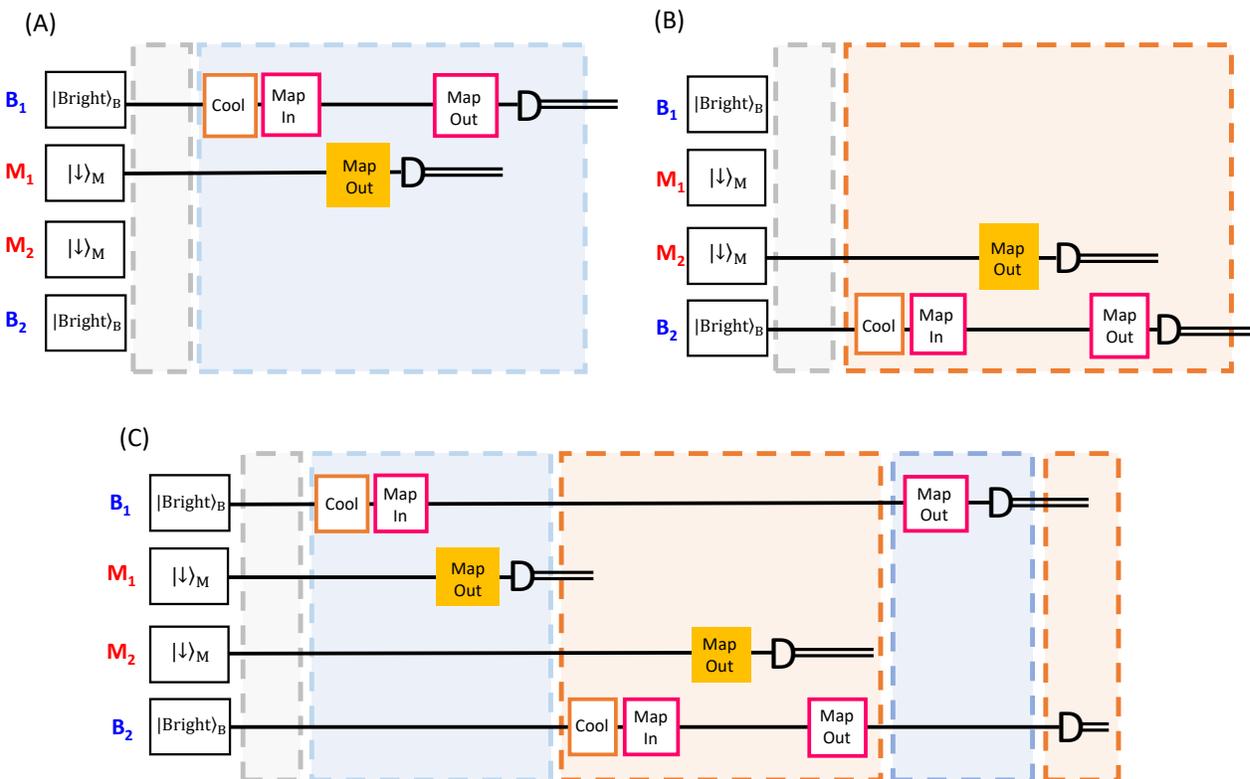


Figure 5.7: Circuits for measuring depumping errors due to scattered resonant light in the quantum gate teleportation experiment. Parts (a) and (b) are reference experiments where mixed-species pairs of ions are prepared, separated, cooled, and detected in the appropriate configurations. Part (c) implements the gate teleportation experiment up until the relevant detections, except without any of the logic gates. The difference between detections in experiment (c) and the corresponding detections in parts (a) and (b) allows us to estimate the depumping errors due to stray resonant light. Results of these estimations are shown in Table 5.3. Figure reproduced from [Wan et al., 2019].

	Ref. (Fig. 5.7(a/b))	Meas. (Fig. 5.7(c))	Difference
M ₁ Bright	97.90(45)	98.00(44)	0.10(63)
M ₁ Dark	99.90(10)	99.70(17)	-0.20(20)
B ₁ Bright	99.70(17)	98.00(44)	-1.70(48)
B ₁ Dark	99.80(14)	99.20(28)	-0.60(32)
M ₂ Bright	98.60(37)	96.36(59)	-2.21(70)
M ₂ Dark	99.80(14)	99.90(10)	0.10(17)
B ₂ Bright	99.50(22)	99.80(14)	0.30(26)
B ₂ Dark	99.80(14)	99.70(17)	-0.10(22)

Table 5.3: Depumping errors measured in the experiments shown in Fig. 5.7. The first column lists the ion detected and which state it was prepared in. Column two gives the percentage of shots where the measured state matched the prepared state, as shown in Fig. 5.7(a) for the first four (M₁ and B₁) and 5.7(b) for the last four (M₂ and B₂). The third column shows the measured disagreement after incorporating the cooling and detection sequences as used in the actual experiment. The final column gives the difference of the middle two, where the difference is attributed to depumping from stray resonant light.

is moved into the LIZ before cooling its ${}^9\text{Be}^+$, mapping it into the computational basis, mapping out and detecting ${}^{25}\text{Mg}^+$, and finally mapping out and detecting the ${}^9\text{Be}^+$. These tests serve as SPAM references for the two mixed species pairs, respectively. In part (c) we run the entire gate teleportation experiment up until the relevant detections, except without any of the gate operations, leaving only shuttling, cooling, mapping into and out of the computational space, and detections. This experiment includes SPAM errors and depumping error induced by stray resonant light. We estimate the scattered light error as the difference between the detections in part (c) and the corresponding detection in part (a) or (b), depending on which ion is considered. The results of this evaluation are shown in Table 5.3, with statistically significant differences highlighted in bold. Each experiment was run after preparing both $|\uparrow\rangle$ and $|\downarrow\rangle$.

From this, we find statistically significant errors on both B₁ and M₂. The error on B₁ results from cooling on B₂, including both Doppler cooling and repumping during sideband cooling. The error on M₂ comes from detecting M₁. We take the average of the differences for both the $|\uparrow\rangle$ and $|\downarrow\rangle$ preparation experiments as our depumping error, finding an error

of 1.1(4)% from stray resonant light on M_2 while detecting M_1 and an error of 1.2(3)% on B_1 while cooling on B_2 .

We can roughly estimate how likely a spectator ion is to absorb a photon that was scattered directly from the target ion in the LIZ and compare this stringent lower limit to the experimentally observed depumping rate. First we assume that the target ion is scattering at a rate of $R_t = \gamma/2$, where γ is the linewidth of the transition. Assuming for simplicity a spherical emissions pattern, the intensity of the scattered light at the position of the spectator ion is:

$$I_s \approx \frac{\gamma}{2} \cdot \frac{hc}{4\pi\lambda z^2} \quad (5.5)$$

where h is Planck's constant, c the speed of light in a vacuum, λ the transition wavelength, and z the separation between the target ion and spectator ion. We approximate the absorption cross section of the spectator ion as $\sigma_s \approx \frac{3}{2\pi}\lambda^2$. The scattering rate of the spectator ion is then:

$$R_s \approx \frac{I_s}{hc/\lambda} \cdot \sigma_s = \frac{3\gamma\lambda^2}{16\pi^2 z^2}. \quad (5.6)$$

Taking $^{25}\text{Mg}^+$ detection as an example, for $\gamma = 2\pi \times 41.3$ MHz (Sec. 2.3), $\lambda = 280$ nm, and a minimum separation z of 390 μm , the probability of at least one scattering event in a 200 μs detection window is roughly 5×10^{-4} .

This intrinsic scattering from the target ion falls far short from being able to produce our measured depumping error, so the scattered light must come primarily from something other than the target ion. Possible explanations include scattering from the trap structure or backreflection from the output window of the vacuum system. Ideally, scatter off the front electrode should not be visible to spectator ions, although scatter off its interior edge could be. If coming from the trap, it is perhaps more likely that the scatter comes off of the back electrode, to which the detection beam comes in at a 45 degree angle. This electrode sits nominally in the shadow of the front electrode, however any finite angle of incidence of

the laser beam relative to the wafer or slight misalignment of the trap wafers themselves would make the back electrode visible to the laser beam. Any of these effects are further complicated by imperfections in the laser beam spatial profile, which we know to be present at least to some extent from the hexagonal nature of the photonic crystal fiber cores. Damage to the fiber tips, dust on mirrors or lenses, etc. can all further degrade the beam profile.

During each detection period we collect roughly 30 photons for a bright ion, however many more are scattered but not collected. Accounting for the solid angle captured by our imaging system and other transmission/reflection losses, on the order of 10^3 photons are scattered by a bright ion. A depumping error of 0.01 then implies a relative resonant light intensity ratio of order $10^3/10^{-2} = 10^5$ between the intended ion and spectator ion in the neighboring potential well. We can make an estimate, under a few assumptions, of the plausible intensity due to back reflection from the output window with standard Gaussian beam calculations. The beam size is given by

$$w(z) = w_0 \sqrt{1 + \left(\frac{\lambda z}{\pi w_0^2}\right)^2} \quad (5.7)$$

where w_0 is the beam waist, λ the wavelength, and z the axial distance from the location of the waist. With this, the intensity at some radial distance r from the center is given by

$$I = R \cdot I_0 \cdot \exp\left(-\frac{2r^2}{w(z)^2}\right) \quad (5.8)$$

where R is the fraction reflected by the output window. Taking $^{25}\text{Mg}^+$ detection as an example, we use $w_0 = 0.025$ mm, $\lambda = 280$ nm, $z = 100$ mm, $R = 0.001$, and $r = 0.390$ mm to find $I/I_0 = 9 \times 10^{-5}$. Unless R is orders of magnitude less than the rough estimate used, it seems that window backreflection alone could easily account for our observed depumping rate.

In this calculation, we assume normal reflection, which in reality is unlikely to be exact. Regardless, when scaling up, a greater number of qubits presents even more targets for stray resonant light, so these concerns will persist. This problem is not unique to our

setup, as similar kinds of errors have been measured highly engineered commercial systems [Pino et al., 2021]. Moreover, it is not unique to ions in general. Neutral atom arrays, for example, typically have even greater densities of scattering centers and will suffer from similar crosstalk. In future systems it may therefore be beneficial to either intentionally tilt the laser beams to avoid this possibility, use carefully angled windows, or measure qubits indirectly through a second species (Chapter 7).

There are two main approaches to reducing resonant light crosstalk. The first is to directly engineer the system to carefully control resonant light throughout the trap. As discussed in the previous paragraph, one easy step is to angle the windows to the vacuum system. Careful control of the beam profile to avoid unwanted side lobes is another engineering step to mitigate stray resonant light. Our UV optical fibers already do a pretty good job at this, however, the geometric structure of the photonic crystal core imprints some hexagonal pedestals to the output mode which results in only 98% overlap with a perfect Gaussian mode [Colombe et al., 2014]. Ions can simply be separated by sufficient distance, however that adds overhead for separating ions for detection that otherwise need to be nearby to interact and is in conflict with the goal of scaling to ever larger numbers and densities of ions.

Another approach is to instead offload all resonant laser tasks to a separate species of ion. There are two key steps to trapped ion quantum information processing that require dissipation to the environment: state initialization and state detection. The first can be broken up into both ion internal state initialization and motional state initialization (i.e. cooling to near the ground state). Sympathetic cooling, where one species is used to cool the collective motion of a mixed-species system, has already long been employed in the QCCD architecture and is widely accepted as a necessary feature of any complete system [Wineland et al., 1998, Kielpinski et al., 2000]. Likewise, measurement can be done indirectly through the second species through various schemes for quantum logic spectroscopy that map the state of the qubit to the separate species readout ion [Schmidt et al., 2005, Tan et al., 2015].

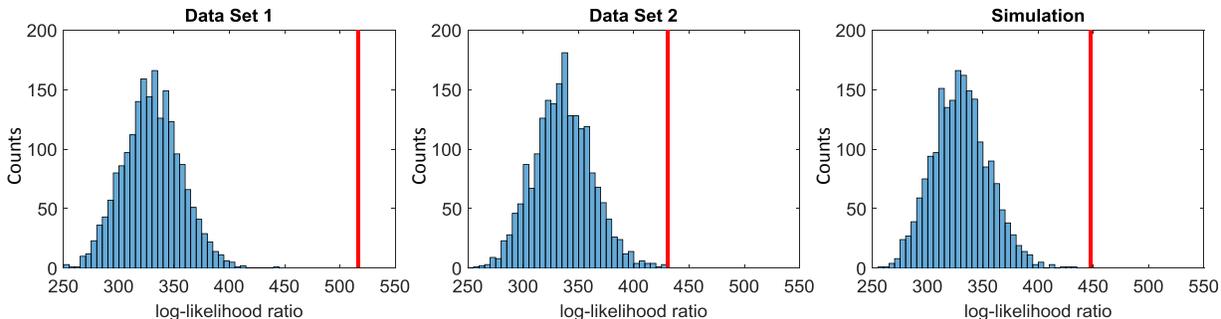


Figure 5.8: Histograms of log-likelihood ratio statistics from parametric bootstrapping compared to the value obtained from the tomography data, measured or simulated (vertical red lines). The first two frames show the experimental results and the third a simulated dataset accounting for the experimental drift identified in Sec. 5.7. The differences between the value from the tomography datasets and the histogram means are 7.1, 3.6, and 4.6 standard deviations for the three panels, respectively. Figure reproduced from [Wan et al., 2019].

In Chapter 7 we propose and demonstrate a promising protocol for this purpose that is resistant to spontaneous emission errors, enabling higher indirect measurement fidelity than what can typically be achieved directly with ${}^9\text{Be}^+$ fluorescence measurements.

5.6 Consistency Checks

Process tomography and fidelity give significant information about the average process, however they do not contain information about how the process may have changed over time. In this section, we give only a brief conceptual overview of the likelihood ratio tests done to check for consistency in the process throughout data acquisition to motivate the search for experimental drifts described in Sec. 5.7. This analysis was again led by Karl Mayer, with details in the supplementary materials of Ref. [Wan et al., 2019].

The general idea is to test the null hypothesis that our experimental results are properly represented as the result of a single consistent quantum process that did not vary with time. A log-likelihood ratio statistic is computed that quantifies the difference in likelihoods that

our experimental data came from a fully unrestricted model compared to a model restricted by our null hypothesis. We used parametric bootstrapping of the observed process to produce a distribution of log-likelihood ratio statistics, as shown in Fig. 5.8. The observed log-likelihood ratio statistic derived from each experimental dataset is then compared to this distribution to determine whether we can reject the null hypothesis. This can be quantified as the difference between the log-likelihood ratio corresponding to the experimental dataset and the bootstrapped distribution mean, scaled by the standard deviation. For the two datasets, this gives differences of 7.1 and 3.6 standard deviations, providing strong evidence to reject the null hypothesis.

This indicates that our quantum process is drifting significantly throughout the data taking process, which is no surprise given the length of time taken and complexity of the setup used. Motivated by this we searched for sources of drift in our experimental setup as described in Sec. 5.7. We found significant drifts in the Rabi rate for the copropagating carrier pulses for ${}^9\text{Be}^+$ single qubit rotations as operated during the data run, and simulated a full tomography dataset with this drift included in addition to the errors listed in Table 5.2. This was modeled as a normal distribution of over-rotation errors during the single qubit rotations in composite entangling gates, state preparation, measurement basis selection, and mapping between the computational and measurement bases. The same analysis techniques as used for the process tomography on the experimental data were applied to the simulated tomography data, producing a histogram of bootstrapped log-likelihood ratio statistics. The resulting log-likelihood ratio statistic differed from the histogram mean by 4.6 standard deviations, from which we conclude that such drifts can account for a major part of the observed discrepancies in log-likelihood ratio statistic in the experimental data.

5.7 Identifying Experimental Drifts

Motivated by the measured discrepancy in the log-likelihood ratio test, we underwent a theoretical and experimental search for a plausible cause. We eventually narrowed in on drifts

in the single qubit rotations on ${}^9\text{Be}^+$ ions driven with our motion-insensitive copropagating carrier beams (Sec. 3.6.2). This arises because both Raman beam tones go through the exact same path to the trap after the optical fiber. The *total* power of the two tones is stabilized with our pulse shaping pickoff photodiode and servo circuitry (Fig. 3.5). However, the *relative* power of those two tones can drift over time. Transition Rabi rates depend on the product of the two, so it is possible for them to drift in our setup while the total power remains constant.

To support that such drifts were actually present during data runs for the gate tomography, we performed a set of simple test experiments. In each, we apply an odd number of $\frac{\pi}{2}$ -pulses to a single ${}^9\text{Be}^+$ ion and then measure the population. This ideally puts the final state along the equator of the Bloch sphere, where population measurements are maximally sensitive to small deviations in the rotation angle, which would occur if the transition Rabi rate were to change. We then simply repeat this experiment, alternating between 1, 3, 5, and 7 $\frac{\pi}{2}$ -pulses, over a period of roughly five hours without recalibrating the pulse time. Calibrations of other parameters that could also drift and impact the results, for example the magnetic field, were still run periodically in a manner similar to what was done during the gate teleportation data runs.

The results of these experiments are shown in Fig. 5.9. We almost immediately observe drifts, which are impactful on a time scale much smaller than the interval that was used for calibrating ${}^9\text{Be}^+$ copropagating carrier pulses during gate teleportation data acquisition. The drift shows over-rotation, at one point drifting back towards the initial calibration before wandering off again. As expected, the more pulses the greater the deviation from measuring 50% of the population to be bright. Also, the 1 and 5 pulse experiments drift one way, while the 3 and 7 pulses drift the other way. This is consistent with what we would expect for persistent over- or under-rotations. An over-rotation on the downward part of the loop around the Bloch sphere ($4n+1$ $\frac{\pi}{2}$ -pulses) will lead to less bright population than expected, while on the upward part of the loop on the other side ($4n + 3$ $\frac{\pi}{2}$ -pulses) it will lead to more

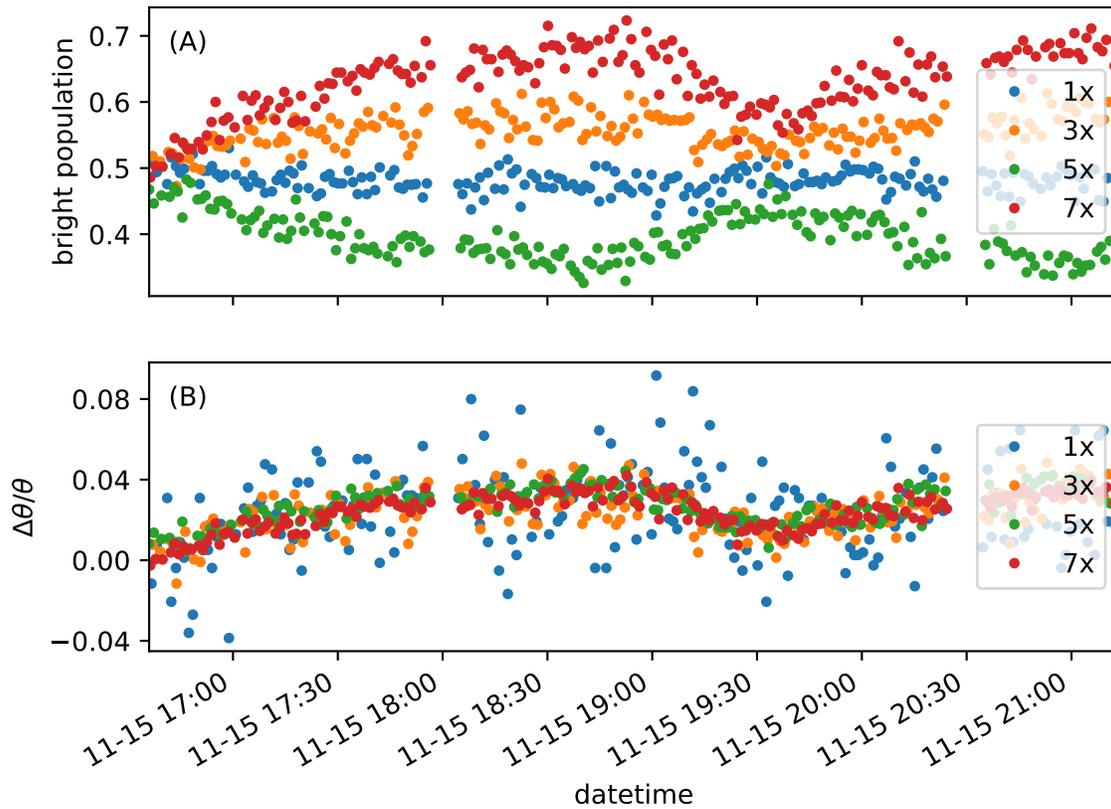


Figure 5.9: Drifts in ${}^9\text{Be}^+$ copropagating carrier $\frac{\pi}{2}$ -pulses Part (a) shows the drift in bright population after simple experiments where 1, 3, 5, or 7 $\frac{\pi}{2}$ -pulses were applied to a single ${}^9\text{Be}^+$ ion. Ideally the final state will lie along the equator of the Bloch sphere for each of these experiments, resulting in an average bright population of 0.5. (b) From these results, we estimate the fractional change in rotation angle needed to produce them. Gaps in the data were for reloading of ions lost due to background gas collisions. Figure reproduced from [Wan et al., 2019].

bright population than expected.

In Fig. 5.9(b) we take the observed drifts in measured populations in part (a) and infer the fractional change in rotation angle $\delta\theta/\theta$ that would lead to such a measurement. As mentioned above, being near the equator of the Bloch sphere provides maximum sensitivity to rotation angle drifts, but on the other hand will lead to more noise in the rotation angle inferred from the population measurement. This effect is observed in the blue data points where only a single $\frac{\pi}{2}$ -pulse was applied, however each experiment appears to follow the same general drift. From this we conclude that the rotation angle in our ${}^9\text{Be}^+$ copropagating carrier pulses during data acquisition could have drifted by up to 4% under the same circumstances as those present during the test.

Our theory collaborators added such a drift to their circuit simulation to generate the synthetic dataset and applied the log-likelihood ratio test described in Sec. 5.6. They observed a 4.6 standard deviation discrepancy (Fig. 5.8), from which we conclude that this drift could have significantly contributed to our experimental data's failed consistency checks.

After discovering this drift, we made an attempt to servo it out. Referring back to Fig. 3.8, after separating from the ${}^9\text{Be}^+$ motion-sensitive Raman beamline the copropagating carrier beamline first passes through a single pass switch AOM, and then through a double pass AOM where both the zeroth and plus first orders are retroreflected and coupled into the optical fiber. We shape these pulses by measuring the intensity after the fiber with a pickoff and feeding back to the RF amplitude driving the switch AOM (see Fig. 3.5 for more details). With the goal to stabilize the relative power of the two copropagating beams, we sampled their beatnote after the fiber on a fast photodetector. The measured beatnote amplitude was used to feedback on the RF amplitude driving the double pass AOM.

This did not work well because the beatnote amplitude also depends on the power delivered to the fast photodetector, causing the two servos in these feedback loops to fight each other. We therefore abandoned this plan for the time being, however we outline here

some possible future improvements that might be made. One option is simply to use a more sophisticated control loop than the simple PI servos that can handle this kind of nonlinear behavior. Another might be to sample the two tones while they are separated in the double pass AOM setup. We could then servo the ratio of the powers, making the total power divide out and hopefully avoid the conflict of this servo fighting with the pulse shaping. This is no small task, as the distance the beams are separated is kept very short to minimize potential relative phase drifts. Sampling before the lens allows for easier separation of the two beams since they are still diverging, however stably fitting in the required elements without compromising divergent path length could be challenging. Perhaps a better option is to make the first of the two retroreflecting mirrors (90 degrees to each other) partially transmissive. This samples the beams without requiring any additional elements in the sensitive region, at the expense of the two beams being parallel and slightly harder to isolate from each other.

Note that this second idea would only work if the source of the drift is the AOM diffraction efficiency. If the source instead introduces different drifts in fiber coupling efficiency, using the beatnote after the fiber would be required. This could come from changes in temperature gradient where the two beams are separated in the retroreflector, or even from a temperature drift after the double pass AOM if the two beams are not perfectly aligned (which is certainly true to some extent).

Finally, the observed drift was slow, so these pulses could always just be calibrated more frequently. However, if we have learned anything at all from this gate teleportation experiment, it is that calibrations add up to a significant overhead in time and introduce the problem of the apparatus drifting away from earlier calibrations faster than one can revisit them. This introduces another imperfection that is usually not taken into account when fault tolerant thresholds are estimated. Making various components passively stable eases such concerns. As we will outline in the following section, minimizing calibrations to only happen when needed will be useful for any large scale QCCD architecture to be successful. This limits system up time, but could be improved through smarter calibration routines.

5.8 Validators and Dependencies

Running such a complex experiment that incorporates the key elements of the QCCD architecture requires numerous parameters to be precisely calibrated. Maximizing the amount of time the system can spend running experiments, rather than being calibrated, requires a careful approach to those calibrations. More complex operations depend on the success of their components, so there is a natural hierarchy of calibrations. For example, perhaps the most far reaching calibration is the magnetic field. Discrepancies between the assumed and actual magnetic field will lead to frequency shifts of energy levels due to their magnetic field dependence (see Eq. 2.1). This leads to driving fields being off-resonant and slippage in phase tracking. Properly calibrating a microwave π -time, for example, depends on doing the calibration at the correct magnetic field. We call such relationships *dependencies*, where each operation except one (magnetic field strength) has its own set of dependencies.

One way to calibrate an algorithm would be to simply start at the most basic parameter, and then calibrate others according to the hierarchy of their dependencies. In the initial build up, this is the natural way to do things. However, when trying to maximize the duty cycle of data acquisition in the final measurement run, it is a rather poor choice since it does not factor in the rates at which different parameters drift. Instead, we took a more top-down approach where we constructed test experiments we call *validators* that test the fidelity of a certain operation. Each operation has its associated validator(s) along with its dependencies. If the validator passes a predetermined threshold fidelity, then we deem all of that operation's dependencies to be sufficiently well calibrated. If not, a calibration of that operation and validations of each of its dependencies are triggered. These short experiments allow us quick ways to test that a calibration is still valid, without having to rerun it in its entirety.

As an example, we can test all of the pulses used during state preparation and shelving by mapping from the optically pumped bright state to the maximally dark state and back. This experiment has all of the individual transition π -pulses as dependencies. If the final

bright state fidelity passes the chosen threshold, all of the involved pulses are then validated. On the other hand, if it fails to pass threshold, then it triggers the validators of each of its dependencies.

Another example might be a ${}^9\text{Be}^+$ copropagating carrier $\frac{\pi}{2}$ -pulse that serves as one of the dependencies for a CNOT gate (Sec. 4.7). Its dependencies include the magnetic field and pulses mapping between the computational and measurement bases. We used two experiments as validators. The first measured the probability of returning to $|\uparrow\rangle$ after applying four $\frac{\pi}{2}$ -pulses. The second measured the probability of ending in $|\downarrow\rangle$ after six pulses. Exceeding both thresholds was required for passing the validator. Failing either or both triggers recalibration of the pulse parameters, in this case the transition frequency, pulse duration, and phase offset between subsequent pulses (to account for Stark shifts). Failure will also then trigger validations of the dependencies.

Using this framework, data acquisition proceeded as follows. The complete gate teleportation algorithm has a set of dependencies like the CNOTs and Bell state generating gates. Each dependency has a time period since its last validation over which it is still assumed to be valid. If all dependencies passed their validations within their assigned time periods, we proceed with data collection in chunks of roughly 300 shots corresponding to a given tomography setting. If the validation for a dependency has expired, we first trigger its validators. However, each of those validators has their own set of lower level dependencies that must have been passed within a chosen time frame before it can be run. A validator was not executed until all of its dependencies passed their validations within the allotted time frames, all the way down to the most basic operations for which there are no lower level dependencies.

Finally, each individual shot during data collection was preceded by a fluorescence precheck that verifies that the desired number of ions are present and reasonably cold. If a given block of 300 shots had more than 1% of its prechecks fail, it was discarded. This resulted in the tomography settings having variable numbers of shots, from 297-300.

5.9 Outlook on Quantum Gate Teleportation

We now conclude this chapter with outlook on possible improvements to quantum gate teleportation and related considerations for implementations of the QCCD architecture.

5.9.1 Utility within the QCCD Architecture

The utility of quantum gate teleportation, together with a means to generate entanglement between separated ion traps, is straightforward. However, it can find use even *within* a QCCD array, though the motivation is a bit more subtle.

Specific to the QCCD architecture that we investigate, quantum gate teleportation can be used for more than just connecting multiple distinct QCCD modules. In particular, if a gate needs to be done between two distant qubits, it may be quite costly to transport them together to perform the gate. Meanwhile, the rest of the processor could be on hold waiting for these steps to complete. Quantum gate teleportation offers a way to mitigate this dead time from distant shuttling as sketched in Fig. 5.10. Rather than waiting for the data qubits to be brought together, an entangled pair of resource qubits can be prepared and distributed ahead of time, thus eliminating the shuttling induced latency for gates between distant qubits. This, of course, comes with the overhead of more gates and the additional entangled pair of qubits, so hardware specifics will determine when this is worth it. Moreover, carefully mapping the desired quantum circuit to the specific hardware constraints can help minimize spatial distance between interacting qubits.

Our experiment is the most simple proof of principle demonstration in our trap that we could come up with, and hence does not accurately reflect how gate teleportation will be used in the long run. Regardless, it serves as an excellent synthesis of many of the key features that are necessary for realizing a large-scale QCCD processor. Namely, it involves separation and shuttling transport waveforms, individually addressed single qubit gates and detections, same- and mixed-species two-qubit gates, and real-time conditional operations.

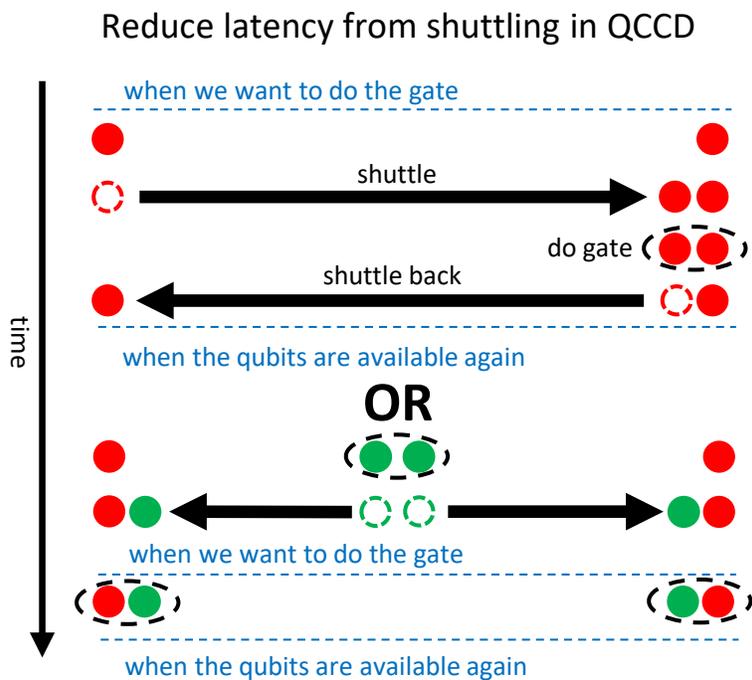


Figure 5.10: Performing a two-qubit gate between distant qubits in the QCCD architecture can be done by shuttling the two ions together, however this comes at an expensive time cost as other parts of the algorithm might need to wait for the qubits to be shuttled. This latency can be greatly mitigated through quantum gate teleportation by predistributing an entangled pair to the two locations. This can be done ahead of time, in principle requiring no additional dead time.

5.9.2 QCCD Electrode Structure

The first improvement that could be made to our trap that should be considered in future QCCD designs is the ability to independently compensate radial micromotion in each trapping location. Our single bias electrode (Sec. 3.2) prevents us from compensating in more than one location at a time. A segmented bias electrode closer to the trap would overcome this problem.

Though all of our gates happened in our single laser interaction zone, in this experiment we calibrated the two mixed-species gates separately. The two trapping configurations had nominally the same axial trap frequencies, though they differed slightly in practice. Future experiments in this trap that use multiple trapping potentials should strive to shim the wells to have the same mode frequencies and ion positions to possibly allow for a gate to be applied in either orientation with the same calibrated parameters. Working with different crystal configurations, e.g. BM and MB would complicate this, since the ions will need to sit in different positions relative to the laser beams. This is not insurmountable, as the beam center can in principle be aligned to the average position in each configuration.

Another limitation of our trap is the small number of electrode segments in our main experiment region. In this demonstration, certain electrodes contributed greatly to the axial confinement of each of the two locations. This makes it difficult to individually tune the parameters of each well, including trap frequency, ion position, and micromotion compensation. Future traps would greatly benefit from additional electrode segments between trapping locations, giving more degrees of freedom to tune these parameters.

5.9.3 Roles for Each Species

Careful examination of our circuit (Fig. 5.4) reveals a barrier to embedding it in a larger quantum algorithm. After separation and shuttling, but before the local mixed-species CNOT gates, we cool on the ${}^9\text{Be}^+$ ions. Hence any quantum information they might

have brought into the operation would be destroyed. This highlights the shortcut we took in this proof of principle demonstration, in that the second species was also used for long-lived qubits. If instead the same species is used for both the qubits and the entanglement resource, the second species can be reserved as dedicated coolant ions. This prevents needing to store quantum information in a field-sensitive qubit (and associated decoherence losses), and allows the data qubits to bring in arbitrary quantum information from earlier in a larger algorithm. Such a demonstration has since recently been performed by Honeywell Quantum Solutions, where they used $^{171}\text{Yb}^+$ ions for all four qubits and had dedicated $^{138}\text{Ba}^+$ coolant ions [Pino et al., 2021].

The QCCD architecture depends on trapping two species, however there seems to be little reason to encode information in both. Only one can make use of a field-independent qubit transition, which will lead to significant decoherence in the other ($^{25}\text{Mg}^+$ in our case). Sympathetic cooling is also an ever present requirement, and a gate needing to be done between qubits of separate species with encoded information leaves no option for cooling after bringing them together. It is hence much more easily handled with a dedicated species.

The role of the second species, however, should not be limited only to cooling. Scattered resonant light was also a significant problem in this experiment. Quantum logic spectroscopy (QLS) between species is an invaluable tool for qubit readout that would completely remove the threat of spectator qubit depumping during detection [Schmidt et al., 2005, Tan et al., 2015].

Chapter 6

Applications of the Mutually-Controlled Multiflip

Spectroscopy of atomic and molecular systems has been a powerful tool for testing theories like quantum electrodynamics and physics beyond the standard model [Baron et al., 2014, Ludlow et al., 2015, Ahmadi et al., 2018]. The vast array of species provides a wide range of transitions within their electronic structures with a similarly wide range of interesting properties. For example, some transitions might have very narrow linewidths for use in atomic clocks, others with easily calculable transition frequencies to compare measurements with theory, and yet others with cycling transitions that enable efficient laser cooling and state readout. The last of these is particularly important, since these high precision experiments must work at relatively low particle numbers to avoid being tainted by things like collision broadening. This makes direct detection of the absorption/emission of most interesting transitions difficult, so the ability to map between the interesting transition and a cycling transition is extremely valuable.

Such a mapping between a transition of interest and a cycling transition *within* the electronic structure of an atom or molecule is typically done with lasers or microwaves. However, not all species of atom or molecule have a readily accessible cycling transition to map to. Quantum logic spectroscopy (QLS) is a powerful tool for ions that extends that mapping *between* two ion species that share a common mode of motion, thereby enabling readout of ion species that would otherwise be intractable. In essence, this combines the favorable properties of both species, where we call the species with the transition of interest

the spectroscopy ion (SI) and the species with the well controlled cooling/detection the logic ion (LI).

QLS has been put to judicious use in a wide range of experiments. We list here a portion of them to give a flavor for the utility of this approach: $^{27}\text{Al}^+$ ions have an attractive clock transition with an 8 mHz natural linewidth and a small blackbody radiation shift (a typically large systematic shift in optical clocks), but its cycling transition requires too short of a wavelength to be practical [Rosenband et al., 2007]. Cotrapping an $^{27}\text{Al}^+$ ion with a $^{25}\text{Mg}^+$ ion has since enabled the most accurate atomic clock to date, with a systematic uncertainty below 10^{-18} [Brewer et al., 2019]. This kind of control enables precision tests of general relativity and geodesy [Chou et al., 2010, McGrew et al., 2018]. By comparing to another clock species, it then also allows for tests of variations in fundamental constants like the fine-structure constant [Rosenband et al., 2008, Beloy et al., 2020]. QLS is also being employed for state preparation and detection of individual molecular ions [Wolf et al., 2016, Chou et al., 2017, Chou et al., 2020] or highly charged ions [Micke et al., 2020].

The traditional QLS protocol shown in Fig. 6.1 transfers information directly through the shared motion [Schmidt et al., 2005]. Assuming perfect preparation of the LI and motional ground state, but an arbitrary unknown SI state, the initial state of the system is $(\alpha|\downarrow\rangle_{SI} + \beta|\uparrow\rangle_{SI})|\downarrow\rangle_{LI}|0\rangle$. A blue sideband π -pulse on the SI then injects a quantum of motion only if it is in $|\downarrow\rangle$, resulting in $|\uparrow\rangle_{SI}|\downarrow\rangle_{LI}(\beta|0\rangle + \alpha|1\rangle)$. At this point the quantum information that was in the SI has been transferred to the collective motion of the ions. The information transfer to the LI is then completed with a red sideband π -pulse, giving $|\uparrow\rangle_{SI}(\beta|\downarrow\rangle_{LI} + \alpha|\uparrow\rangle_{LI})|0\rangle$.

As briefly introduced above, this technique has been put to great use, however it does come with its share of limitations. First, it explicitly depends on starting in the motional ground state, which is never quite true in practice, leading to relatively low readout fidelities. Second, the SI information is destroyed in the measurement process. These shortcomings can be ameliorated by modifying the traditional QLS procedure to instead take place between

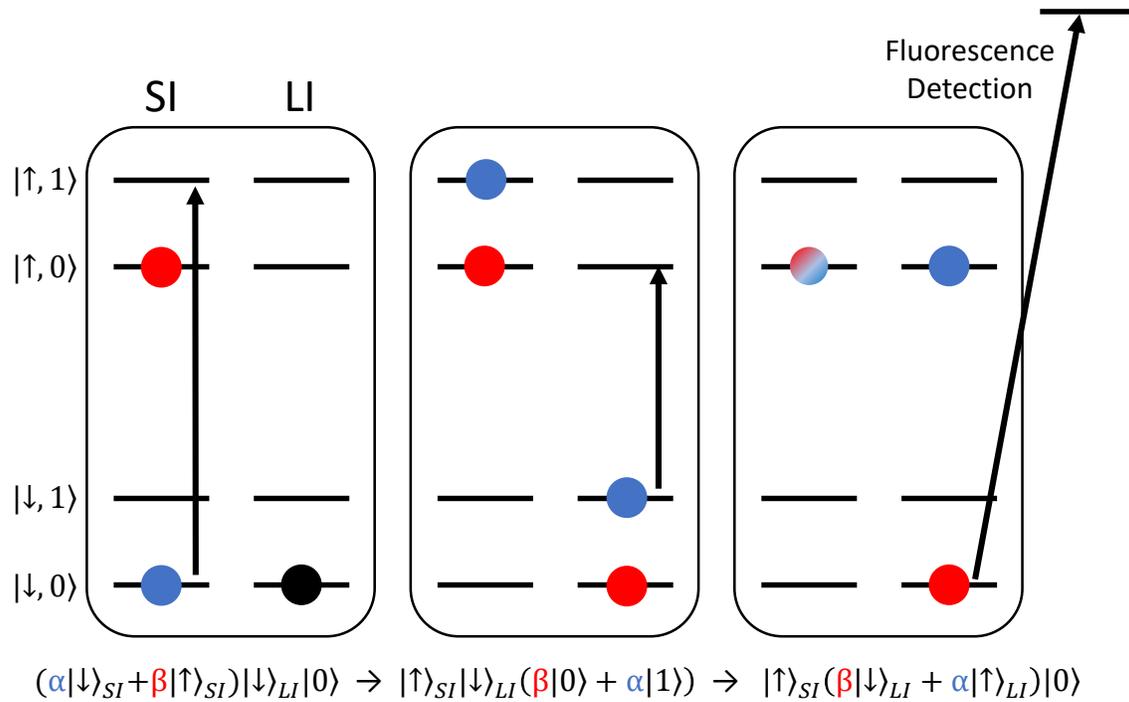


Figure 6.1: Protocol for traditional quantum logic spectroscopy. Beginning in the motional ground state, information about an arbitrary SI state is transferred to the shared motion with a motion adding sideband π -pulse. This information is then transferred to the logic ion with a motion subtracting sideband π -pulse.

one of the SI qubit states and an auxiliary state $|aux\rangle$ to produce a quantum non-demolition (QND) measurement (Sec. 7.2). For example, QLS can be done on the $|\downarrow\rangle \rightarrow |aux\rangle$ transition to detect the SI qubit state without touching $|\uparrow\rangle$ (hence the non-demolition). The QND measurement can then be repeated many times with Bayesian analysis to significantly improve readout fidelity [Hume et al., 2007]. In the initial demonstration in Ref. [Hume et al., 2007], $|aux\rangle$ is a relatively short lived ($300 \mu\text{s}$) excited state, so it had with high probability decayed back to $|\downarrow\rangle$ before the next round of detection. If using hyperfine ground states, such decay will not happen, so either an alternate QLS mapping must be used (Sec. 6.2) or an expanded protocol involving multiple probes per round (Chapter 7).

Repetitive QND measurements were used to demonstrate approximately 99.94% readout fidelity of an $^{27}\text{Al}^+$ clock [Hume et al., 2007], however this also comes with a cost. Fluorescence detection of the LI involves scattering many photons, which with a finite Lamb-Dicke parameter leads to recoil heating that must be recooled. The required recoiling to the ground state decreases the duty cycle of the experiment with increased time being spent on detection. In atomic clocks, decreasing the duty cycle of the clock probe increases the undesirable Dick effect, related to aliasing of the local oscillator noise [Dick, 1987, Santarelli et al., 1998].

It is thus desirable to have alternate QLS protocols that are not highly sensitive to the motional ground state fidelity. As with traditional QLS, we can turn to the wealth of knowledge developed for trapped ion quantum information processing through temperature-insensitive quantum logic operations. Towards that end, significant research has gone into geometric phase gates that are first-order insensitive to initial temperature and have been used to perform entangling gates with up to 99.92(4)% fidelity in same-species crystals [Gaebler et al., 2016] and up to 99.8(1)% in mixed-species crystals [Hughes et al., 2020]. Such gates can be adapted for QLS by surrounding them with appropriate single qubit rotations, for example by constructing a CNOT gate between $^9\text{Be}^+$ and $^{25}\text{Mg}^+$ ions (Sec. 4.3) [Tan et al., 2015].

A similar approach was proposed in Ref. [Mur-Petit et al., 2012], where a ZZ interaction is surrounded by a basis rotation on the LI to map the resulting phase shift to a population measurement. In this chapter, we describe experiments that use the XX-type interaction with what we call the multiflip (Sec. 4.5), removing the need for wrapper pulses and combining the spectroscopy probe and readout pulses into a single operation. Sec. 6.1 describes how the multiflip can be used for spectroscopy, including data presented in Ref. [Kienzler et al., 2020]. Sec. 6.2 will then discuss some of the ways the multiflip can be used in quantum information processing. In particular, it will focus on its use for indirect state preparation and readout, thereby avoiding entirely the errors from scattered resonant light observed in our quantum gate teleportation demonstration (see Table 5.2 and Sec. 5.5). I will also discuss how the multiflip can be used to measure error syndromes for quantum error correction.

6.1 Multiflip as a Tool for Spectroscopy

In this section we present results from Ref. [Kienzler et al., 2020]. Recall from Sec. 4.5 that the behavior of the multiflip depends on the number of ions participating in the interaction. For an even number of participating ions, the net effect of the multiflip is a global bitflip, while for an odd number the net effect is the identity (up to a global phase). Consider then, for example, an odd number of SIs and a single LI. The LI is assumed to always be participating, requiring well calibrated control. If the SI drive is also on resonance, then an even number of ions participate in the interaction, resulting in a global bitflip that is detected on the LI. On the other hand, if the SI drive is far off resonance, then only the single SI participates in the interaction, being odd in number, resulting in no bit flip. In the intermediate regime near (but not on) resonance the behavior is more complicated as the interaction transitions between these two behaviors (see Fig. 6.4).

Implementing the multiflip is the exact same as the MS interaction, except for two key differences. First, we introduce an additional common detuning δ_S for the two SI tones. Detunings from resonance for the two species' drives are then $\delta_{LI} = \pm(\omega_m + \delta_{MS})$ and

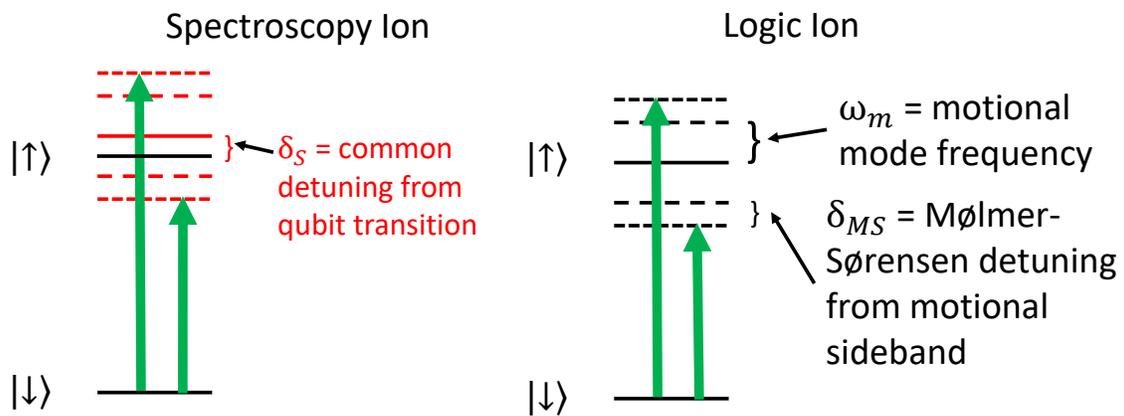


Figure 6.2: Sketch of energy levels when using the multiflip for spectroscopy. As with the typical MS interaction (Sec. 4.4), ideally symmetrically detuned red and blue sidebands are driven on both species. For spectroscopy, we introduce an additional common detuning δ_S to the SI drive as the difference between the actual and assumed qubit frequency.

$\delta_{SI} = \pm(\omega_m + \delta_{MS}) + \delta_S$ (Fig. 6.2). A search for a spectroscopy line can then be done by scanning δ_S . The second difference is to enclose twice the area as the typical entangling gate loop in phase space. This can be done by simply running the interaction with the entangling gate parameters but for twice the typical entangling gate time ($2 \times 2\pi/\delta_{MS} = 4\pi/\delta_{MS}$), thereby looping around twice in phase space, or by adjusting the parameters to enclose the doubled area in a single loop (for example by increasing the Rabi rates by a factor of $\sqrt{2}$).

Both Rabi-type and Ramsey-type spectroscopy can be done with the multiflip. In Rabi-type spectroscopy, the multiflip is applied as described in a single pulse. Scanning the frequency will then trace out a line shape (Fig. 6.4(a)-(b)). Unlike typical Rabi spectroscopy, the more complicated nature of the driving fields gives a messier lineshape, though still well behaved near resonance and ideally with full contrast. The multiflip can instead be used for Ramsey-type spectroscopy by dividing it into two equivalent “effective $\frac{\pi}{2}$ -pulses,” with a free precession period in between. Each of these should return to the origin in phase space to avoid residual spin-motion entanglement.

The multiflip is robust to quite a few experimental imperfections. In both types of spectroscopy, it inherits the temperature insensitivity of the MS interaction to first order in the Lamb-Dicke approximation. Moreover, the center of the resonance does not depend on the initial states of either species (only that they are in the manifold being probed), although the specific line shapes do. Therefore preparing the SIs to a pure state is not necessary when initially searching for an unknown transition. In Sec. 6.1.1 we demonstrate robustness to these two parameters, as well as to the SI Rabi rate (which may not be well known when searching for a transition). Then, in Sec. 6.1.2 we show how the entangling nature of the MS interaction driving the multiflip gives rise to a quantum enhancement in frequency sensitivity.

6.1.1 BM Crystal Multiflip Results

Performing experiments on a ${}^9\text{Be}^+ - {}^{25}\text{Mg}^+$ (BM) crystal requires first initializing it to a well defined order. As with the BMMB crystal used in the quantum gate teleportation experiment (Chapter. 5), background gas collisions can reorder the crystal. With a BM crystal, reordering should not result in different normal mode frequencies, however in reality trap asymmetries can lead to small shifts. In addition, this shifts the position of the ions in the laser beams, potentially changing transition Rabi rates and π -times. This led us to calibrate two separate mixed-species CNOT gates in the gate teleportation experiment for the BM and MB crystals.

Therefore, before each shot of an experiment with the BM crystal we also run a crystal order initialization sequence as described in Ref. [Home, 2013] and depicted in Fig. 6.3, with minor modification. Due to the crystal asymmetry, the procedure is slightly more complicated than the symmetric BMMB ordering sequence, however it also simply involves ramping static potentials up/down to take advantage of the differential radial confinement of the two species. Deterministic crystal ordering is done with the following four steps:

- (1) A radial shim pushes the ions off the radial micromotion null. Different ion masses gives different radial confinement, so the ${}^{25}\text{Mg}^+$ ion is pushed out further than the ${}^9\text{Be}^+$ ion.
- (2) A quadrupolar twist is applied to the neighboring DC electrodes, so that the ${}^{25}\text{Mg}^+$ ion always falls on the chosen side, independent of initial order.
- (3) The radial shim is relaxed, bringing the ions back to the micromotion null.
- (4) The quadrupolar shim is relaxed, returning to the initial potential well.

In practice, we found it beneficial to superimpose an additional axial squeeze during step three, and then unsqueeze during step four. In principle, this should not be necessary.

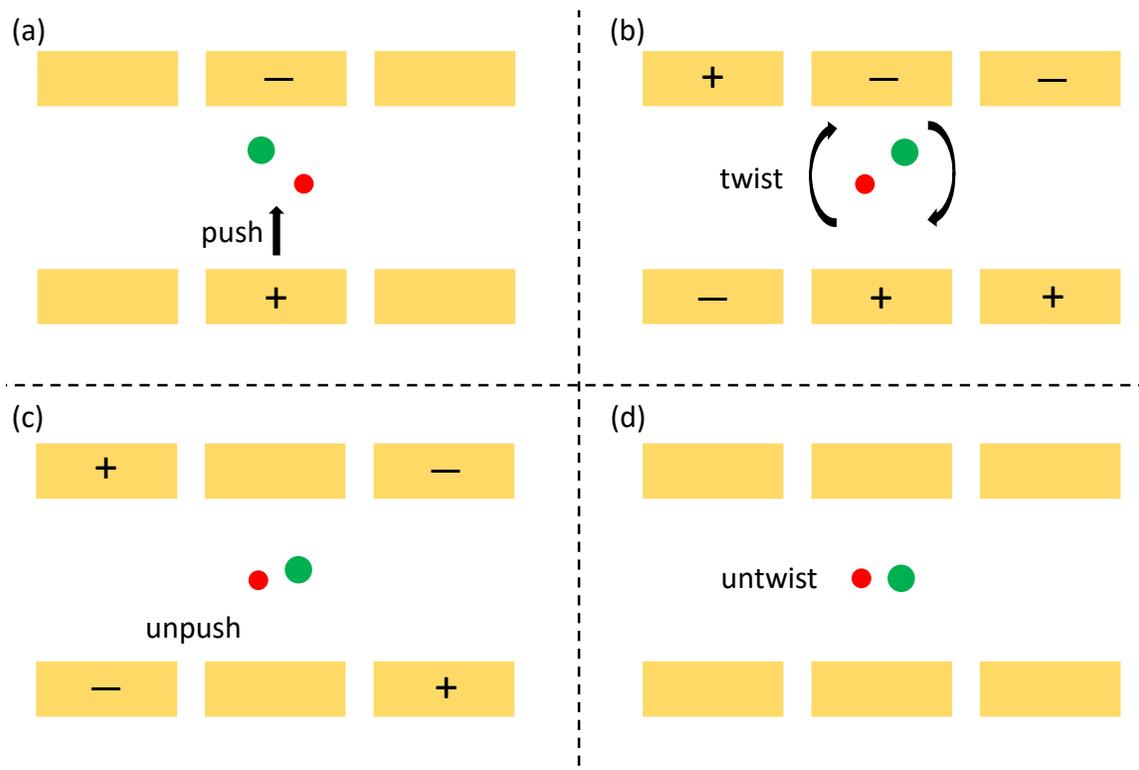


Figure 6.3: Initializing the order of a BM crystal from an arbitrary order. (a) A radial shim is first used to push the ions off the radial micromotion null. Due to the different masses, and hence different radial confinements, the heavier $^{25}\text{Mg}^+$ ion is pushed out further. (b) A quadrupolar shim on the neighboring electrodes is used to twist the crystal to ensure that the $^{25}\text{Mg}^+$ ion falls on the correct side. (c) The radial shim is relaxed to bring the ions back to the micromotion null. (d) Finally, the quadrupolar twist is relaxed. In these experiments, the ions were located at electrodes 19 and 20 (see Fig. 3.2).

We begin with the basic behavior of the multiflip on the BM crystal in near ideal conditions (i.e. gate parameters are all well calibrated and the ions are near the motional ground state). First, we demonstrate the line shape of Rabi-type spectroscopy by preparing both ions to $|\uparrow\rangle$ before applying the multiflip while scanning δ_S (Fig. 6.4(a)). As expected, far off resonance, both ions are measured to be in $|\uparrow\rangle$. However, as we approach $\delta_S = 0$, we see both ions flip to $|\downarrow\rangle$. The lineshape is much more complicated than for typical Rabi spectroscopy, most notably with the pronounced side lobes. These arise when $|\delta_S| \approx \delta_{MS}$, bringing one of the two sideband tones into resonance. The asymmetry in these lobes arises from the fact that motion-adding and motion-subtracting sidebands behave differently when the motional mode is initialized near the ground state.

Fig. 6.4(b) shows similar data, except with the SI prepared in $|\downarrow\rangle$. The center of the resonance does not change, although the surrounding lineshape does. In particular, the sidelobes are mirrored due to the reversal of which SI sideband is motion adding/subtracting for preparing $|\downarrow\rangle$ instead of $|\uparrow\rangle$. As suggested by these two experiments, though not shown explicitly, the center of the LI resonance is independent of the initial SI state provided its population is within the manifold being probed, including mixtures and superpositions.

Numerical simulation fits of these Rabi-type experiments are shown as partially transparent solid/dashed lines of the corresponding color for the SI/LI in Fig. 6.4. The free parameters in these fits are the interaction time t_{MS} , gate detuning δ_{MS} , the average Rabi rate for the two species $\Omega_{MS} = \sqrt{\Omega_{MS,SI} \cdot \Omega_{MS,LI}}$, and a parameter c quantifying the Rabi rate imbalance where $c = \frac{\Omega_{MS,LI}}{\Omega_{MS}} = \frac{\Omega_{MS}}{\Omega_{MS,SI}}$. We characterize the Rabi rates this way because in our calibration procedure we scan one of the species' Rabi rate to optimize closing the loop, which may leave some residual imbalance. If working with multiple ions of each species, as will be discussed in Sec. 6.1.2, this can be avoided by calibrating the MS gate loop areas individually. Note that the relative phase of the drives for the two species (motion phase) still requires a joint calibration. Varying t_{MS} might be unexpected, but was done as a shortcut to approximately account for differences in pulse shape between the experiment

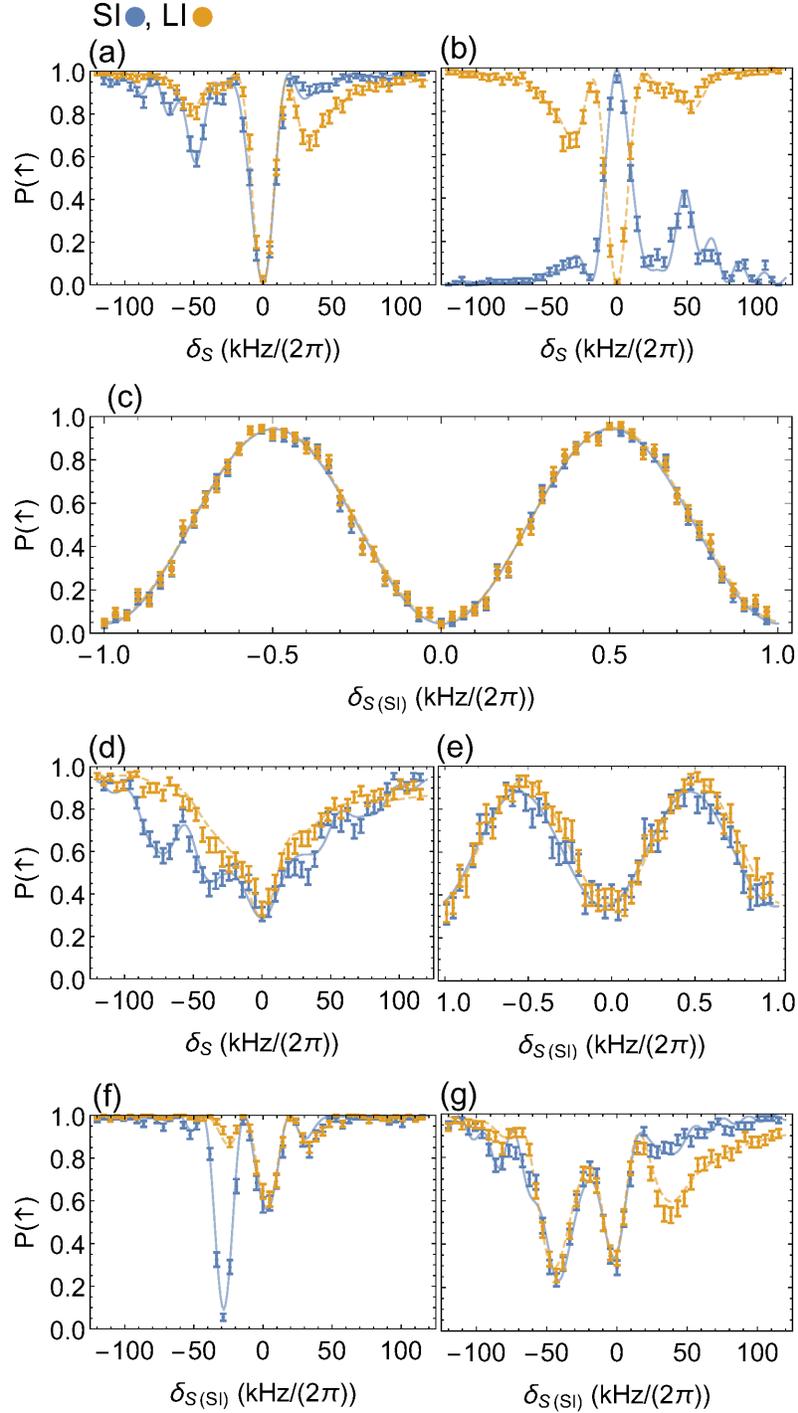


Figure 6.4: Multiflip results for various experiments on a BM crystal. SI ($^{25}\text{Mg}^+$) results are shown in blue and LI ($^9\text{Be}^+$) results in orange. Numerical fits are shown by solid/dashed lines of the same color. (a) Rabi-type spectroscopy after ground state cooling and preparing $|\uparrow_{SI}\uparrow_{LI}\rangle$. (b) Rabi-type spectroscopy after ground state cooling and preparing $|\downarrow_{SI}\uparrow_{LI}\rangle$. (c) Ramsey-type spectroscopy after ground state cooling. (d) Rabi-type spectroscopy after Doppler cooling only (compare to part (a)). (e) Ramsey-type spectroscopy after Doppler cooling only (compare to part (c)). (f) Rabi-type spectroscopy after ground state cooling but with $\Omega_{SI} = \Omega_{MS}/2$. (g) Rabi-type spectroscopy after ground state cooling but with $\Omega_{SI} = 2\Omega_{MS}$. Figure reproduced from [Kienzler et al., 2020].

and our simulations. Simulations assumed square pulses, but experimentally we shape our Raman beam pulses to round out the edges and reduce pulse bandwidth (see Sec. 3.5). Thus the numerical fit approximates the pulse that smoothly ramps on and off as a square pulse of roughly equal area.

Basic Ramsey-type spectroscopy is shown in Fig. 6.4(c), where the multiflip is divided into two equal-length effective $\frac{\pi}{2}$ -pulses with 1 ms of free evolution between them. The first pulse maps the ions into a GHZ state (Eq. 4.7) $|\uparrow\uparrow\rangle + e^{i\phi}|\downarrow\downarrow\rangle$ where ϕ depends on the interferometric laser phases of the motion-sensitive Raman beams driving the interaction. As long as these phases are stable over the whole experiment, the outcome of the whole Ramsey sequence does not depend on ϕ . After free evolution the second pulse maps any phase accumulated to the qubit populations. Ideally, the only phase accumulation happens due to detuning between the SI and its driving fields. The LI is not assumed to be ideal, so LI dephasing and its mitigation will be discussed in Sec. 6.1.4.

Fitting sinusoids to the results gives LI contrast of 0.89(1) and a phase offset of $\Delta\phi = 0.01(1)$ between the two species' oscillations. The phase offset is consistent with zero, but as we will discuss (and remedy) in Sec. 6.1.5, certain errors can lead to appreciable line pulling effects that result in a non-zero phase. The reduction in contrast is also more than we would expect from errors in the map in/out pulses. We attribute the additional reduction to phase fluctuations in the Raman beams performing the spectroscopy, which are known to be interferometrically unstable due to the drastically different paths the beams take to the ions.

We now demonstrate the robustness of the multiflip by intentionally introducing certain errors. First, we show robustness to initial temperature by repeating the Rabi- and Ramsey-type experiments after only Doppler cooling the ions ($\bar{n} \approx 3.3$ for the in-phase mode that the gate is driven on). In Fig. 6.4(d) we see that there is still an appropriately centered response, albeit with reduced contrast and a broader line shape. Likewise, we perform Ramsey-type spectroscopy with fringe contrast of 0.57(2) and a phase offset $\Delta\phi = 0.21(4)$

between the two species. We attribute the large reduction in contrast to our fairly high Lamb-Dicke parameters of approximately 0.17 and 0.29 for ${}^9\text{Be}^+$ and ${}^{25}\text{Mg}^+$ respectively, leading to larger fluctuations in the Rabi rate from Debye-Waller factors [Wineland et al., 1998]. Simulations show that for the same initial thermal distribution, contrast of roughly 97% can be achieved with Lamb-Dicke parameters of 0.1 for both the SI and LI.

The phase shift is significant this time, arising from loop-closing errors due to the thermal state of the motion, large Lamb-Dicke parameters (and associated Rabi rate fluctuations), and errors in pulse parameter calibrations. Detailed simulations of these effects, as well as ways to mitigate them, are discussed in Sec. 6.1.5.

Finally, if searching for an initially unknown transition, the associated Rabi rate may not be precisely known. For the multiflip to be useful for this task it must therefore be at least somewhat robust against differences between the assumed and actual Rabi rates. To test this, we do ground-state cooled Rabi-type spectroscopy, except we set the SI Rabi frequency to half (Fig. 6.4(f)) or twice (Fig. 6.4(g)) the known value. From the data we see that in each case the LI still shows a resonance at $\delta_S = 0$, although with reduced contrast and with amplified side lobes when the transition is overdriven. This is not a big problem for an initial search, since once any kind of signal is observed the gate parameters can be tuned to narrow in on the transition resonant frequency and π -time.

Interestingly, when the SI is underdriven one of the sidelobes becomes its strongest resonance. Since the ions are cooled to the ground state, this happens only on the motion adding side of the frequency scan. The weaker MS coupling then allows this fringe to be more pronounced than in the ideal case (compare Fig. 6.4(a) and 6.4(f)). In contrast, when overdriving the SI the increased MS coupling causes the SI and LI spectra to overlap over a larger frequency range (Fig. 6.4(g)).

6.1.2 BMMB Crystal Multiflip Results

As we saw in Sec. 6.1.1, the multiflip is robust to a number of experimental imperfections. In this section, we outline and perform a basic proof of principle demonstration of the multiflip on a larger system of ions (BMMB crystal). We choose this crystal because we have experience in working with it from our quantum gate teleportation experiment (Chapter 5). The Ramsey fringe periodicity does not depend on the number of LIs, but with multiple of them their disagreement can be used to filter out erroneous trials that would otherwise reduce fringe contrast, and in some cases correct the error in post processing to make the trial still usable. For example, with two LIs trials where their detection results disagree can be filtered out through postselection, since an error on the LIs unrelated to the SIs must necessarily have happened at some point in the experiment. Ideally they are all prepared to the same state, and then either flip or are unaffected. With three or more LIs, a majority vote can correct the signal without needing to discard it. In the result of a tie, the trial can still be discarded as in the two LI case.

First we tested the protocol with Rabi-type spectroscopy, however the behavior is more complicated when we have even numbers of both LIs and SIs instead of an odd number of each. Ideally, with odd numbers of each the final LI state would depend on whether or not the SIs participated in the interaction. Instead, with our BMMB crystal with the SIs both on resonance and far off resonance we have an even number of ions (4 vs 2) participating, so in both cases we expect a global bitflip on the LIs. This makes probing a single point potentially ambiguous, however when tracing out the lineshape the more complicated intermediate regime still gives a line shape that allows one to find $\delta_S = 0$ (Fig. 6.5(a)). The shape is akin to a dark resonance superimposed on a broader line for the LI, which would be the only observable signal if the SI cannot be interrogated directly.

The bipartite nature of the MS interaction can provide at best a factor of two reduction in linewidth compared to a single SI [Shaniv et al., 2018]. The greater benefit to additional

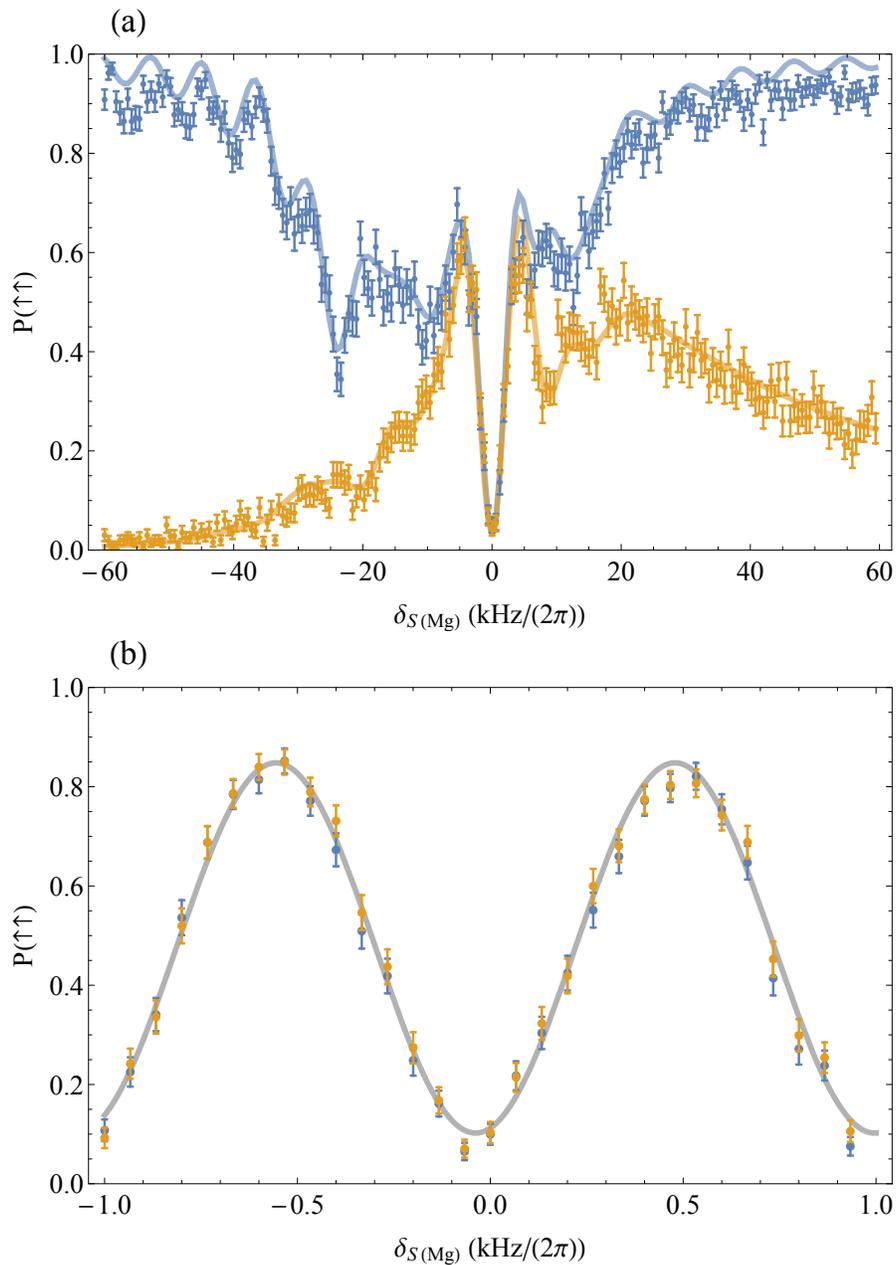


Figure 6.5: Multiflip spectroscopy results in a BMMB crystal, with the SI signal in blue and LI signal in orange. Rabi-type spectroscopy is shown in (a). Here the LI signal on resonance and far off resonance is similar, however the more complicated intermediate dynamics still give a discernible resonance feature. Performing the spectroscopy on a system of odd number SIs and LIs would avoid this ambiguity. Ramsey-type spectroscopy is shown in (b) with a free evolution time of $500 \mu\text{s}$. The same fringe periodicity is observed in half the free evolution time used in the BM case (Fig. 6.4(c)), consistent with the expected scaling with the number of SIs. Figure reproduced from [Kienzler et al., 2020].

SI) comes from allowing an entangled state to evolve during free evolution in Ramsey-type spectroscopy. This arises from the entangling nature of the MS interaction, where the map in pulse takes an even number of ions to a GHZ state of the form (ignoring normalization) $|\psi\rangle = |\downarrow\downarrow\dots\downarrow\rangle + e^{i\phi_0} |\uparrow\uparrow\dots\uparrow\rangle$ where ϕ_0 depends on the phases of the laser beams. For free evolution time τ and detuning δ from the local oscillator, this N particle state evolves to $|\psi\rangle = |\downarrow\downarrow\dots\downarrow\rangle + e^{i(\phi_0+N\delta\tau)} |\uparrow\uparrow\dots\uparrow\rangle$. The map out pulse then translates the accumulated phase to the ion populations, oscillating as a function of frequency with an N -fold reduction in period. Ideally, this gives an N -fold increase in frequency sensitivity (i.e. an N -fold increase in maximum slope), achieving what is called the “Heisenberg limit” where the precision of the experiment is fundamentally limited by the uncertainty principle [Leibfried et al., 2004b]. With N independent and unentangled two-level systems, sensitivity instead only increases as \sqrt{N} [Ludlow et al., 2015].

The first experiment towards Heisenberg-limited spectroscopy with trapped ions was performed on three ions of the same species [Leibfried et al., 2004b]. We build on this by extending the principle to a mixed-species crystal, giving the benefit of indirect quantum logic readout. Results for Ramsey-type spectroscopy on our BMMB crystal are shown in Fig. 6.5(b), where a free evolution time of $500 \mu\text{s}$ was used. Comparing with Fig. 6.4(c), we observe the same periodicity, but with half the free evolution time, which is consistent with what we would expect from Heisenberg scaling. In this proof of principle demonstration, imperfections in the operations and driving fields lead to reduced contrast that prevented realizing the full potential of this scaling. The fitted contrast was 75(2)%. We estimate the fidelity of our Ramsey pulses that map into and out of the GHZ state by measuring the GHZ state fidelity after mapping in, finding it to be 0.937(6). Doubling this error to account for both pulses falls well short of explaining the Ramsey contrast. We again attribute the discrepancy to fluctuations in the interferometric laser phase during free evolution. Finally, the phase shift between fitting the SI and LI populations was found to be 0.03(3), which is not significant in this experiment.

One important point to note is that the ideal Heisenberg scaling is limited in practice by increased decoherence of entangled states. For a given free evolution time, decoherence will be greater for an entangled state of N entangled SIs compared to individual SIs, reducing the full N -fold increase in sensitivity that would otherwise be achieved in the absence of decoherence. In fact, the minimum frequency uncertainties attainable for N particles is the same for maximally entangled Ramsey-type spectroscopy (like the multiflip) as for N particles under uncorrelated traditional Ramsey spectroscopy [Huelga et al., 1997]. This comes from balancing the trade offs between increased resolution from longer Ramsey time and decoherence. With the uncorrelated single SI superposition states decohering less quickly than entangled states, they are able to probe longer before decoherence ultimately kills the signal. However, this is not all to say that using entangling spectroscopy protocols like the multiflip is not without value. Though the minimum attainable frequency uncertainties are the same, entangling protocols are able to reach that minimum N times faster. This is no small benefit. Although trapped ion atomic clocks hold the current record for precision [Brewer et al., 2019], working with such small particle numbers limits how fast they can reach such stability compared to neutral atom optical lattice clocks. Therefore, entanglement can remedy one of the main weaknesses of clocks based on single ions.

6.1.3 Multiflip Use Cases

As shown in the previous sections, the multiflip has a variety of useful properties, however not every property is amenable to every application. For example, in state of the art trapped-ion clocks, great care has to be taken to characterize the ion motion and the associated time dilation shift to avoid limiting the clock's accuracy [Chen et al., 2017]. If the ions are left in a relatively warm thermal state of motion, for example by only Doppler cooling as in Fig. 6.4(d)-(e), the resulting shot-to-shot uncertainty in the time dilation shift may become dominating and hinder clock performance. For this reason, it is likely that any high precision or high accuracy clock will still require ground state cooling. Likewise,

warm motional modes will bring Debye-Waller factors that lead to loop closing errors and associated line shifts. Though this is not a fundamental limitation, as will be shown in Sec. 6.1.5, any clock will benefit from passive stability instead of relying on an error signal that cancels out such effects. Any high performance clock using this protocol will still likely need to carefully ground state cool and calibrate gate parameters to avoid residual spin-motion entanglement from loop closing errors. Realizing the full benefit of multiple SIs in such a clock would also require the Ramsey-type protocol.

Pushing the frontiers of accuracy and precision is not the only application for ion clocks, however, so the temperature resiliency of this protocol may still give benefit in other applications where time dilation shifts are not a main limiter. The robustness of the protocol may make it attractive to portable clocks, where the lack of a highly controlled laboratory environment leads to other uncertainties and shifts becoming dominant.

Finally, precision clocks are only one of many spectroscopy applications. The multiflip can also be used to search for an unknown atomic or molecular ion transition. It is in this application where the robustness of the multiflip is most favorable. In this context, the transition Rabi rate is unlikely to be known precisely, making the existence of identifiable features when over/under driving the transition of utmost importance (Fig. 6.4(f)-(g)). Although the exact center of these features may be ambiguous, once the approximate location of the transition is located pulse parameters can be further optimized. Likewise, in an initial search maximum contrast is not required, so the repetition rate can be significantly increased by removing the need to ground state cool.

6.1.4 Impact of Logic Ion Properties and Their Mitigation

In the Ramsey-type protocol, the ions evolve as a combined SI and LI GHZ state. However, our labeling of SIs and LIs is purely artificial. The physics that governs the system makes no such distinction, so being in a GHZ state we have no way to distinguish between phase accumulated by the SIs or the LIs. This can be turned into a feature in creating

“designer atoms” [Roos et al., 2006] or atomic combination clocks [Akerman and Ozeri, 2018] where both species contribute to the timekeeping to reduce certain systematic effects.

For high accuracy clocks, influence of the LIs is potentially crippling. This can come in two main ways: either decoherence or phase accumulation from detuning of the LI drive from resonance. The former reduces contrast (thus limiting the duration of the probe and fringe narrowness), while the latter leads to a shift of the Ramsey fringes. There are several approaches to mitigating these effects. First, since it is the *absolute* LI frequency stability that matters, it is advantageous to pair an SI with an ~ 1 PHz optical clock transition with a ~ 1 GHz hyperfine LI. Taking the ratio of the two transition frequencies, the LI *fractional* frequency uncertainty can be of order 10^6 times worse before having comparable impact on the clock’s absolute frequency stability. That alone may be enough in some clock applications, however with current state of the art ion clocks now reaching down below 10^{-18} fractional frequency uncertainty, this still places a rather stringent 10^{-12} fractional frequency uncertainty requirement on the LI.

If further suppression of LI dephasing is needed, dynamical decoupling can be used to refocus any LI phase drifts. Using such pulse sequences, the coherence time of a $^{171}\text{Yb}^+$ ion was pushed to roughly one hour [Wang et al., 2021]. Finally, shifts from the LI driving fields themselves can be mitigated through “hyper-Ramsey” spectroscopy where the duration, phase, and frequency of the second Ramsey pulse are tailored to cancel the shifts from the Ramsey pulses themselves [Yudin et al., 2010, Huntemann et al., 2012], or “autobalanced” Ramsey spectroscopy where the free evolution time alters between long and short probes, the difference of which can be used to suppress the shift [Sanner et al., 2018].

6.1.5 Line-Pulling Effects and Their Cancellation

When working with state of the art optical clocks, one must be wary of potential shifts and uncertainties from the interrogation pulses themselves. One potentially damaging form of shift arises when the loops in phase space driven by the MS interaction (Fig. 4.2) do not

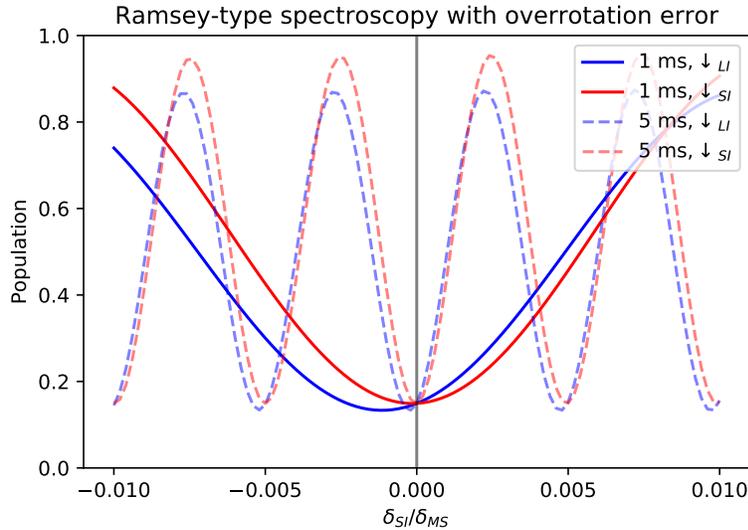


Figure 6.6: Example shift of fringes in Ramsey-type multiflip spectroscopy due to loop closing errors. Here the two pulses are applied for 1.15 times the ideal pulse length. The final $|\downarrow\rangle$ populations are shown in red for the SI and blue for the LI. Solid lines correspond to a Ramsey wait time of 1 ms and dashed lines to 5 ms. In all cases the loop closing error leads to shifts in the LI resonance and an accompanying reduction in contrast, however longer Ramsey wait times result in lower absolute frequency shifts. X-axis is a unitless detuning ratio.

fully close, resulting in residual spin-motion entanglement. Simulations suggest this as an explanation for the shift between SI and LI Ramsey fringes in Fig. 6.4(e). Fig. 6.6 shows a clear shift in LI Ramsey fringe in an otherwise ideal simulation. To exaggerate the effect for illustrative purposes, a pulse time 1.15 times the ideal was chosen, likely far exceeding any realistic experimental error.

Failure to close the motional phase space loops can arise from imperfections in the pulse length, MS detuning, or Rabi rate calibrations. Fluctuations or drifts in the intensity or frequency of the driving fields can also lead to this error. Perhaps most important is the motional state of the gate mode. Any variance in the motional state will lead to shot-to-shot fluctuations in the Rabi frequency (Eq. 2.12), for example due to imperfections in ground state cooling or if operating at the Doppler temperature. This will result in loop closing

shifts and associated frequency uncertainty. Even in a very well controlled clock operating near the motional ground state, that cooling is never perfect and the crystal can heat during free precession (without the ability to recool), inevitably limiting the stability of the clock at some level. It is therefore necessary to construct a method for removing these shifts.

Typical optical clocks interrogated through Ramsey spectroscopy operate by probing on both sides of the central fringe by shifting the phase of the second pulse by $\pm\frac{\pi}{2}$, thus probing at the points of greatest slope/sensitivity of the population signal. Any shift of the local oscillator relative to the clock ions will shift one of the populations up and the other down, informing the control loop which direction to shift the local oscillator to maintain a balanced population signal and therefore resonance. Such operation will cancel out some, but not all, of the shift from loop closing errors. Full cancellation of this shift can instead be achieved by varying an additional parameter, either the SI input state or the direction of their MS detuning. Combining this with the phase shift of the second pulse gives a total of four different measurements, from which an error signal can be constructed that is insensitive to shifts from loop closing error. Fig. 6.7(a)-(d) show how each of these four measurements themselves are shifted off resonance, but collectively give an unshifted error signal when appropriately summed in part (e).

6.1.6 Multiflip Variants

So far in this chapter we have discussed a proof-of-principle demonstration of the multiflip protocol with our apparatus, which is really designed for quantum computing, not precision spectroscopy. For that reason, in this section we discuss variants to the general approach described above that may be better suited for experiments optimized for spectroscopy.

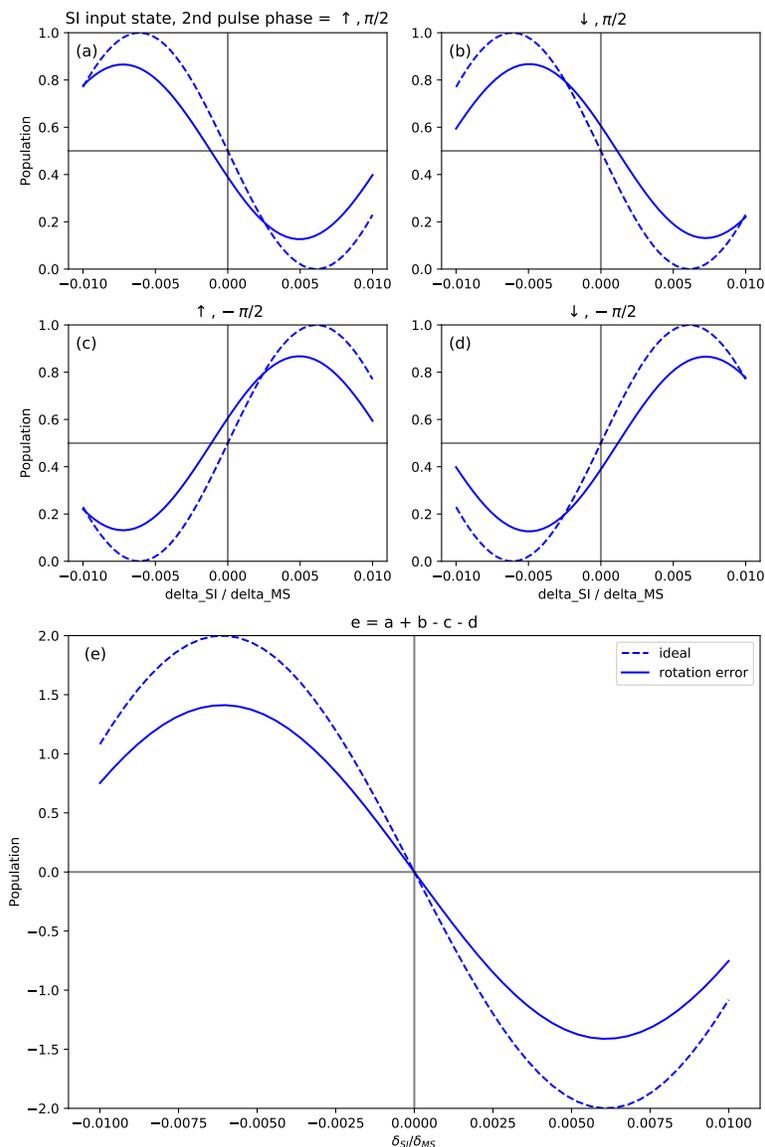


Figure 6.7: Four measurements (a)-(d) used to construct an error signal (e) whose center is insensitive to loop closing errors. Parts (a)-(d) show Ramsey fringes (1 ms wait time) under all combinations of SI input state and phase of the second pulse relative to the first ($\pm\frac{\pi}{2}$). The ideal signals are shown by dashed lines, with solid lines showing the signal with pulse lengths 1.15 times the ideal length. We notice that c-d have opposite slope near resonance, so we construct a four-measurement error signal by taking $a+b-c-d$, shown in part e. The resulting error signal is insensitive to the shift from loop closing errors. Moreover, it ideally has twice the amplitude of the two-measurement error signal in traditional Ramsey spectroscopy [Morinaga et al., 1989], resulting in the same amount of information per shot of the experiment. X-axes are unitless detuning ratios.

6.1.6.1 ZX Multiflip

In our proof of principle multiflip demonstration, we used the MS interaction because that is what our quantum computing experiment is designed for, being typically focused on quantum computing with field-insensitive hyperfine qubits. However, for the purpose of readout, the exact interaction on the internal states is not limited to XX-type interactions. What is critical is that it couples the internal state of the ions to the motion in a well defined way.

A Z-type interaction [Leibfried et al., 2003b] could also be used, which may be the natural choice if used in a state-of-the-art ion clock that uses an optical transition. Our choice of the X-type MS interaction is primarily dictated by the fact that Z-type interactions are not compatible with field insensitive qubits due to their negligible differential Stark shift. An optical qubit has no such restriction, so a Z-type interaction on the SI with its simpler laser setup may be more desirable. If the choice of LI is compatible, a Z-type interaction can also be driven on it, however the interaction will need to be surrounded by an LI basis rotation to convert the induced phase flip to a measurable bit flip. For reasons discussed in Sec. 6.1.4, a field-insensitive hyperfine qubit driven by the MS interaction would still likely be a better choice, and avoids the need for additional wrapper pulses.

We call this hybrid the “ZX multiflip” for the type of interaction driven on the SI and LI, respectively. It follows the same rules developed above for the “XX” multiflip, except the participation parity now instead conditions a phase flip on the SI, with the LI still seeing conditioned bit flips. This modification extends to larger numbers of ions following the same rules. In Fig. 6.8 we simulate the basic behavior on a system of one SI and one LI. Parts (a)-(b) show that the LI flips as expected when the SI is in the manifold being driven. In this case the SI is in a mixture, indicating that the ZX multiflip works independent of the state of the SI as long as its population is in the qubit manifold. If the SI is instead prepared to some auxiliary state $|s\rangle$ outside the manifold being driven, we see in parts (c)-(d) that

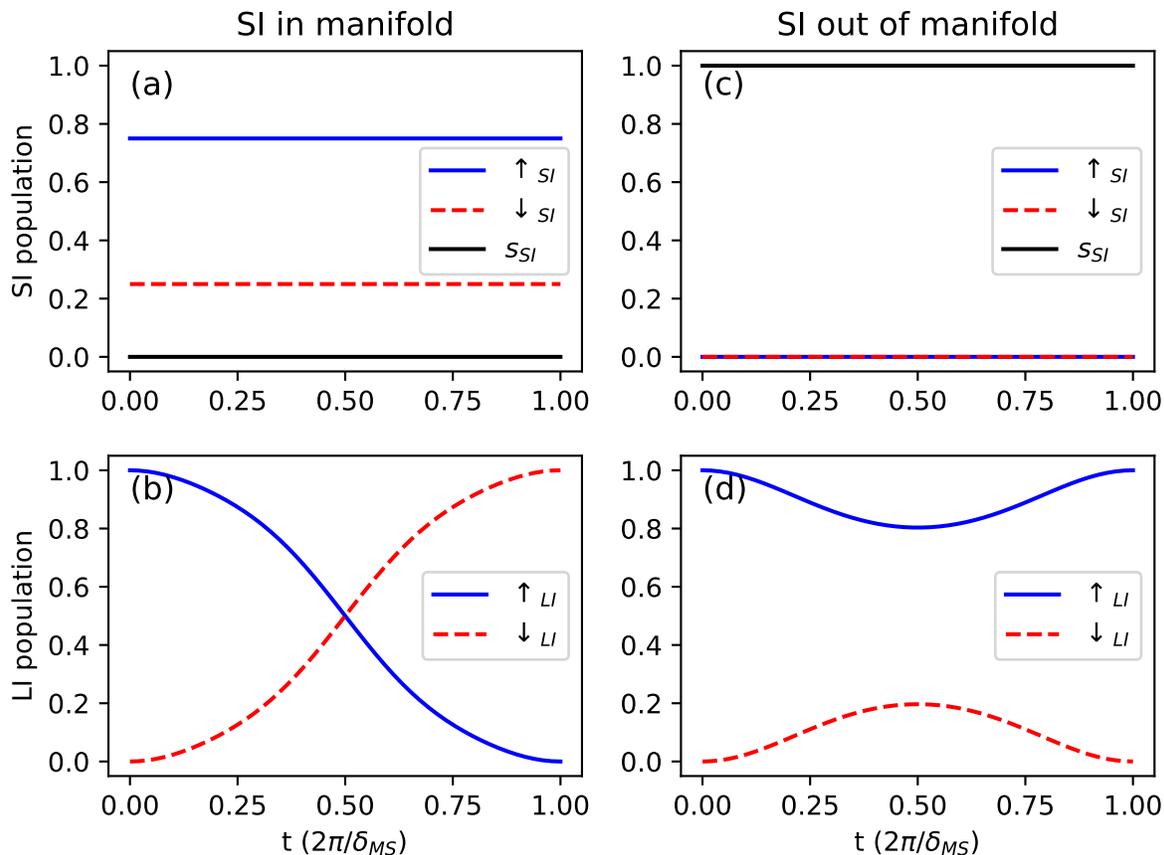


Figure 6.8: Multiflip variant where a Z-type interaction is instead driven on the SI, with the MS interaction still applied to the LI, here simulated on a system with one of each type of ion. Subfigures (a)-(b) plot the behavior when the SI is in the manifold being driven. Shown is a simulation with a mixed distribution of $|\uparrow\rangle$ and $|\downarrow\rangle$. The LI state flips as expected and is not affected by the average over LI states since an even number of ions are participating. Panels (c)-(d) show the evolution when the SI is instead out of the manifold being driven, in some auxiliary state $|s\rangle$, we see that the net effect on the LI is the identity. This concept can be extended to greater numbers of ions. If the SI is in a superposition of two states, one inside and one outside of the manifold, the multiflip will result in the LI population being in a superposition of qubit states whose populations correspond to the SI state populations (because those SI amplitudes govern whether the LI will flip).

the net effect on the LI is the identity. Superpositions of SI states both inside and outside of the qubit manifold will result in an LI superposition, whose amplitudes correlate with the amplitudes of the SI states inside/outside the manifold.

6.1.6.2 Multiflip For Clock Readout Only

Entangling the LI(s) with the SI(s) during the multiflip combines the spectroscopy probe and quantum logic readout operations into a minimal number of laser pulses, however this does not come without its drawbacks. Mainly, as discussed above, any phase accumulation from the LIs during free evolution is indistinguishable from the desired SI signal. Sec. 6.1.4 discusses various ways to mitigate such effects, however the additional overhead required may be deemed too costly for some experiments. In such experiments, the multiflip is not entirely without use. It can still be used for readout only in the following manner: After the clock probe with the final state of the SIs ready to be measured, the multiflip can be driven on an auxiliary transition that involves only one of the qubit states. Provided there is high confidence that population did not start in the auxiliary state, the multiflip informs whether the SIs were in the involved qubit state (or if not then they were in the other). Such a modification of the protocol still provides the temperature benefits of phase gates while it avoids possible contributions of an LI phase to the signal, but comes with the expense of additional laser pulses.

6.2 Multiflip in Quantum Information Processing

We will now shift focus to how the multiflip can be used in quantum information processing. In the previous sections we discussed how it can be used in spectroscopy experiments. While high frequency precision is required for clock applications, requirements on readout fidelity are much looser than in quantum computing applications. First, in Sec. 6.2.1 we discuss how the multiflip can be used for qubit state parity measurements (whether the number of qubits in $|\uparrow\rangle$ is odd or even), which requires transforming the participation

parity (how many qubits are in the manifold that participates in the multiflip) that can be measured with the multiflip into a measurement of qubit state parity. Such multi-qubit measurements are key ingredients in quantum error correction (QEC) [Nielsen and Chuang, 2002, Devitt et al., 2013]. Next, in Sec. 6.2.2 we show how qubit state parity measurements using the multiflip could be used to demonstrate the repetition quantum error correction code. Finally, in Sec. 6.2.3 we outline a protocol for completely nonresonant control of a qubit using the multiflip for both state preparation and measurement, targeted at avoiding the potentially crippling scattered resonant light errors observed in the quantum gate teleportation experiment (Sec. 5.5).

Classical error correction is a useful starting point for understanding quantum error correction. Digital logic, represented by 1's and 0's, is inherently quite robust to noise due to the binary nature of the signals, for example as represented by voltages in a circuit with a threshold to distinguish the two bit states. With a large separation between states, strong noise is required to introduce an error. If further protection is desired, a simple way to achieve this is simply encoding information redundantly in a “logical bit” that is composed of three physical bits in the classical repetition code with states: $0_L \rightarrow 000, 1_L \rightarrow 111$. In this encoding, if a single bit flips the logical information can be recovered through a majority vote. The corrupted state can be restored to the most likely code state before being exposed to more noise in the further course of the calculation. With three bits per logical bit, this encoding can correct for a single error. The notion of “fault tolerant” error rates arises from the observation that the probability of preserving the logical bit will be higher if the error rate on physical bits is below a certain threshold that depends on the code [Nielsen and Chuang, 2002]. Assuming uniform error probability p across all three bits, double errors will happen with probability $3p^2(1-p)$ and triple errors with probability p^3 . The threshold for improvement over the base error rate of a single bit is then the nontrivial solution to $3p^2(1-p) + p^3 = p$, yielding $p_{thresh} = 1/2$. At this threshold error rate, double or triple errors, which result in a logical error, occur with equal probability as an error on a bare physical bit.

QEC codes have analogous thresholds, however typically with much lower threshold error rates due to the number of operations required to measure information about a quantum system without destroying the information that it encodes.

Naively extending the classical repetition code to a quantum code by using a tensor product of multiple copies of a quantum state fails on multiple fronts. While tensor products $(\alpha|0\rangle + \beta|1\rangle)^{\otimes n}$ of an arbitrary but known input state can be prepared, this is not useful for error correction. Local noise impacts copies differently, and even if all noise were global, quantum measurements are projective and probabilistic, yielding a single bit of information. It would therefore take measuring an infinite number of copies to fully characterize the continuous parameters α and β . Moreover, taking an arbitrary but unknown input state and preparing such a tensor product is impossible due to the no cloning theorem [Wootters and Zurek, 1982]. Although the naive extension to quantum information fails, an analogous code can still be constructed. Rather than try and use multiple copies of a qubit, information should be distributed across multiple entangled qubits, for example as $\alpha|0_L\rangle + \beta|1_L\rangle \equiv \alpha|000\rangle + \beta|111\rangle$. This will allow for one bit flip ($\sigma_x = X$) error to be corrected. An analogous code composed of X-basis eigenstates can correct for a phase error ($\sigma_z = Z$), and by concatenating the two encodings, where one encoding is performed on three logical qubits of the other code, we arrive at Shor's original proposal [Shor, 1995].

Quantum measurements are still projective, so we cannot measure individual qubits to determine whether an error happened. Instead, one must measure stabilizers, each a product of multiple Pauli matrices, to infer on which qubit (if any) an error has occurred. The two stabilizers for the bit-flip repetition code introduced above are IZZ and ZZI . Since both logical qubit states are +1 eigenstates of each of these operators, the operators can be measured without learning anything about, and hence projecting, the encoded information. Likewise, if any individual X error occurs, the resulting state will still be an eigenstate of these operators, however with different sets of eigenvalues that uniquely identify which qubit suffered the error. The results of these measurements can be interpreted with the inset table

in Fig. 6.10.

6.2.1 Measure Qubit State Parity with the Multiflip

As simulated in Fig. 4.4, the behavior of the multiflip depends on the parity of the number of ions participating in the interaction. This has several useful applications. Leakage errors refer to an ion leaking to a separate physical state outside the two chosen to encode the qubit. If this happens, that ion will no longer participate in the multiflip, which can be detected by the lack of a spin flip on a readout ion prepared to a known participating state. This has been demonstrated nicely in correcting leakage errors from a qubit that is part of a surface error correction code plaquette [Stricker et al., 2020]. Such errors are particularly damaging to fault-tolerant quantum computation and require additional overhead to correct [Fowler, 2013, Suchara et al., 2015, Brown and Brown, 2019]. Promising alternatives to the circuit based approaches in the preceding references use hardware-based repumping techniques to convert leakage errors to Pauli errors that can be corrected by typical codes [Hayes et al., 2020, McEwen et al., 2021].

QEC codes work by encoding a single logical qubit across multiple entangled physical qubits to achieve the needed redundancy for error correction [Nielsen and Chuang, 2002, Devitt et al., 2013]. Stabilizer codes are characterized by a set of stabilizer operators, each a tensor product of Pauli matrices, for which each logical qubit state is an eigenstate with eigenvalue $+1$. If an error ($X =$ bit flip, $Z =$ phase flip, $Y =$ both) acting on a qubit leads to a state where at least one stabilizer has eigenvalue -1 , this error is detectable by the code. Moreover, if the bit string of eigenvalues is unique, the error is correctable. Crucially, measuring stabilizers reveals no information about the logical qubit. Therefore, measuring stabilizers is key to identifying correctable errors without destroying the encoded quantum information. The simplest circuit for measuring two-qubit state parity using only pairwise

gates is



Protocols based on the MS interaction are powerful for this application, allowing for entangling gates on more than two qubits, and have been used to repetitively measure parities of two ${}^9\text{Be}^+$ ions through a ${}^{40}\text{Ca}^+$ readout ion before applying appropriate corrections to stabilize any of the four Bell states (Eq. 4.6) [Negnevitsky et al., 2018].

The multiflip offers a promising alternative approach to that taken in Ref. [Negnevitsky et al., 2018]. Fig. 6.9 gives an example of how it can be used to measure the ZZ state parity of a pair of qubits, for example in the repetition code (Sec. 6.2.2). Natively the multiflip measures participation parity, which can be converted into a state parity by shelving one of the qubit states to an auxiliary state $|a\rangle$ prior to the multiflip. The parity is then detected by a spin flip, or lack thereof, on a syndrome ion. If the syndrome ion measurement indicates that the data qubits flipped during the multiflip, a conditional corrective pulse returns them to their initial state. Finally, the auxiliary state is unshelved.

Straightforward extensions to more qubits proceed similarly, possibly changing the interpretation of the syndrome ion signal depending on their number. Likewise arbitrary Pauli products, including any stabilizer, can be measured by appropriate basis rotations before/after measuring Z products as in Fig. 6.9.

The matrix form of the multiflip shown in Eq. 4.29 depends on the Raman beam interferometric phase. This phase dependence will persist for any even number of participating ions, while for an odd number it does not (yielding just the identity up to a global phase). While typically stable over the course of a single composite gate, for example a CNOT as constructed in Eq. 4.36, this phase cannot be assumed to be stable over the course of a long experiment employing QEC. Fortunately, the conditional π -pulse in Fig. 6.9 removes this phase dependence if driven with the same motion sensitive Raman beams, making the net

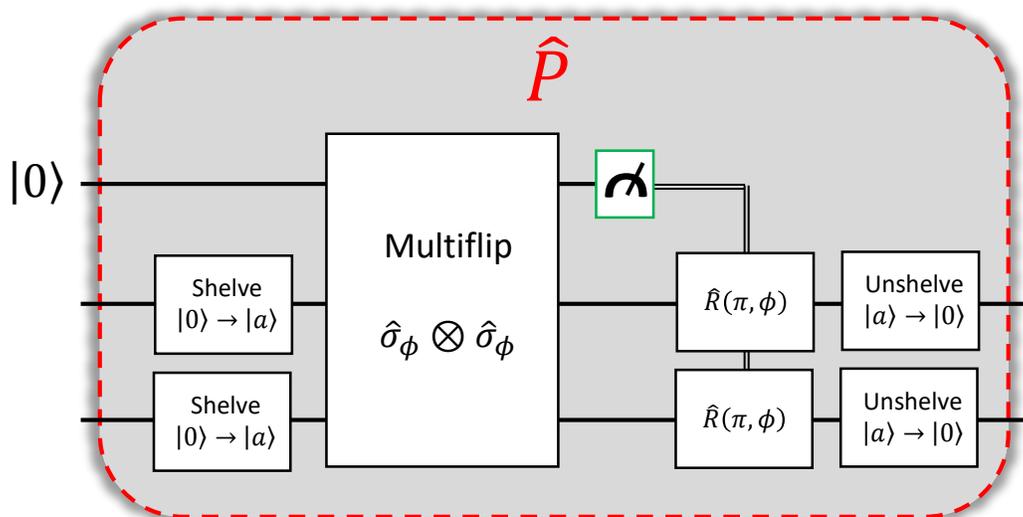


Figure 6.9: Circuit for measuring state parity through the multiflip. A measurement of the parity of ions participating in the interaction is converted to a qubit parity measurement by shelving one of the qubit states to an auxiliary state $|a\rangle$. The multiflip maps the state parity to the syndrome qubit, which is subsequently measured. Conditioned on this measurement a global π -pulse is applied to the qubits if the measurement result indicates that they were flipped by the multiflip. Finally, the shelved population is returned to the qubit manifold. Other Pauli products can be measured by appropriate basis changes before/after the circuit or choice of interaction axis. Extensions to additional ions are trivial, only changing the interpretation of the outcome on the syndrome qubit and condition for the global π -pulse.

behavior of the circuit in Fig. 6.9 the identity on the qubits for either result. This assumes that interferometric phase is stable over the course of the qubit state parity measurement. If the phase is not sufficiently stable over the course of the operation, for example due to latency in detecting the syndrome qubit, the procedure can still be used by replacing the bare multiflip with a composite gate that removes the phase dependence. Analogous to the construction of the phase gates G_{\pm} , the multiflip can be embedded in a Ramsey sequence that uses the same set of motion sensitive Raman beams as itself. This would then be further surrounded by $\frac{\pi}{2}$ -pulses from a phase-stable source that rotate back to the original basis.

In fault-tolerant QEC, it is vitally important that any errors that do happen do not cascade uncontrollably [Gottesman, 2010]. Measuring error syndromes can easily cause cascading if circuits are not constructed carefully. Being a many qubit interaction, using the multiflip to measure stabilizers is vulnerable to such errors. However, that does not mean it is devoid of use cases in fault-tolerant QEC. First, it can be used to verify an entangled cat state of syndrome qubits in Shor-type QEC [Shor, 1996]. Second, it could be used to measure the two-qubit gauge operators in the Bacon-Shor code, and from them the stabilizers, where correlated errors on both qubits result only in a change of gauge and not a logical error [Bacon, 2006]. Third, in flag-based syndrome extraction it can be used to simplify the second round of bare syndrome extraction that gets triggered if an error is detected, in which case further errors cannot be tolerated anyway [Chao and Reichardt, 2018]. Finally, other applications may emerge as the field continues to evolve.

6.2.2 Protocol for the Repetition Quantum Error Correcting Code in a BMBMB Crystal

Earlier in my graduate studies I proposed demonstrating repetitive QEC through mixed-species quantum logic on a BMBMB crystal, using state parity measurements as described in Sec. 6.2.1. Ion crystal lifetimes were bad enough while demonstrating quantum gate teleportation (Chapter 5) that we decided to set this idea to the side for the time

being. It may be an interesting experiment to demonstrate in a future cryogenic system with better vacuum and consequently longer ion lifetimes, so we briefly outline a protocol for implementing it in a small region of a QCCD with a single laser interaction zone.

We envision testing the IZZ/ZZI repetition code on a BMBMB crystal with three ${}^9\text{Be}^+$ ions (Q1-Q3) encoding the logical qubit and two ${}^{25}\text{Mg}^+$ ions as syndrome qubits S1 and S2. Fig. 6.10 shows the circuit for testing this code in our system, overlaid by the transport and separation pattern of the five qubits to implement the code, with the constraints of a single laser interaction zone and being unable to split the crystal into more than two subparts. The logical state is first prepared to the special case $|+L\rangle = \frac{|0_L\rangle + |1_L\rangle}{\sqrt{2}}$ with global operations. Error is then optionally introduced, for example through a global microwave pulse of small angle such that double errors are unlikely. Q1 and S1 are shuttled off to the side, then the multiflip is used to measure IZZ (see Sec. 6.2.1) through S2. The crystal is then recombined before S2 and Q3 are separated off to allow ZZI to be measured. With these two error syndromes, the inset table shows what corrective pulses to apply, if any. Since we cannot split the crystal to isolate Q2 from the others, we use a set of 2 or 3 addressed π -pulses to implement an effective 2π pulse on the qubits without errors and a 1π or 3π pulse on the corrected qubit. Finally, the logical Z basis can then be measured directly, or the encoding (Fig. 6.11) can be repeated to measure in the X basis. The latter would show, if the qubits are measured to be in $|000\rangle$, that no Z errors were introduced on top of correcting any X errors.

Preparing $|+L\rangle$, a GHZ state in the X basis, can be done with only a global MS interaction. For an even number of ions, the MS gate will prepare a Z-basis GHZ state, while for an odd number of ions it will prepare an X-basis GHZ state (as is the case here). Hence we follow the MS interaction with a global $\pi/2$ pulse to rotate into a Z-basis GHZ state (dashed blue box in Fig. 6.11). However, in using our motion-sensitive Raman beams, the phase of this GHZ state will depend on the fluctuating interferometric laser phase. In Fig. 4.5(b), additional wrapper pulses are used to produce a Bell state with a controlled

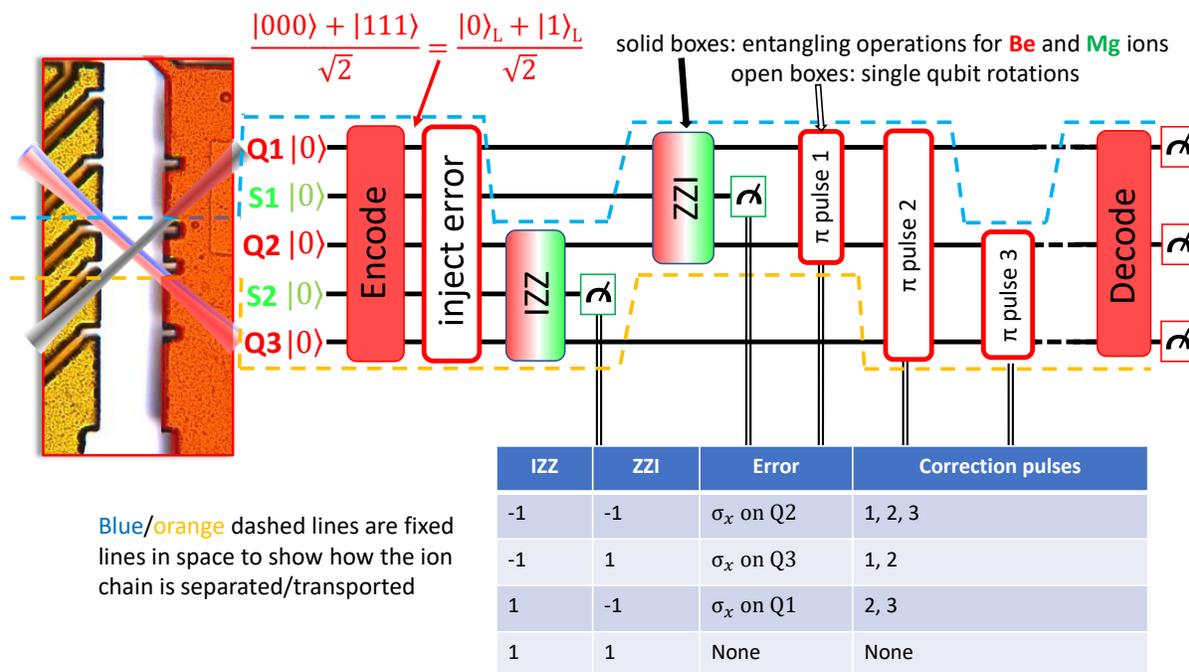


Figure 6.10: Circuit for demonstrating the repetition QEC code with a BMBMB crystal. Qubits Q1-Q3 are shown in red, and the two syndrome qubits S1-S2 in green. Dashed blue and orange lines represent spatial separation between the subset of ions undergoing operations in the laser interaction zone (LIZ) and the remaining qubits. The qubits are first encoded into $|+_L\rangle$ while all are in the LIZ. Q1 and Q3 are subsequently split off with one of the syndrome qubits to allow IZZ and ZZI to be measured, respectively, through the multiflip-based protocol in Fig. 6.9. If the trap does not have enough electrodes to easily split the crystal into three parts, a set of up to three π -pulses are used to correct an identified error, as decoded in the inset table. This applies either a π - or 3π -pulse on the corrupted qubit and 2π -pulse on the uncorrupted qubits. The encoding procedure can be repeated to decode at the end, where the qubits returning to $|000\rangle$ demonstrates that any X errors were corrected and no Z errors were introduced by the process. Trap and laser beams are not shown to scale.

phase. Something similar is proposed in Fig. 6.11 to control the phase of the resulting GHZ state. We surround the GHZ state generator with a Ramsey sequence using the same set of motion-sensitive Raman beams, removing the effect of the fluctuating phase and rotating the basis, analogous to the gate \hat{G} in Fig. 4.5(a). Finally, the basis rotation is undone, except this time with a phase stable source, to produce the desired GHZ state.

Together this provides the recipe for demonstrating the repetition code through mixed-species quantum logic. While this code only corrects for X errors, with our phase stable encoder/decoder one could demonstrate the ability to correct these errors without introducing Z errors, making it promising for concatenation into a complete Shor code. This would serve as an initial demonstration of the utility of mixed-species quantum logic in an error correction context, however would be but a starting point. Further improvements can certainly be made, particularly with regards to fault tolerance, as the same considerations discussed at the end of Sec. 6.2.1 apply here.

6.2.3 Protocol For Preparation and Readout of Hyperfine Qubits Without Resonant Light

Operating a quantum computer purely through coherent unitary operations is not possible. Some form of dissipation to the environment is needed to remove entropy during preparation of both the qubit and motional states and for error correction. Notably, state preparation of syndrome qubits for QEC or resource qubits for quantum gate teleportation are more conveniently done mid-circuit. Likewise, coupling to the environment with resonant light is necessary for state readout, again mid-circuit for the same use cases. However, this resonant light can have deleterious effects if it reaches unintended ions.

As discussed in Sec. 5.5, scattered resonant light imparted measurable infidelity on the teleported CNOT gate. Such errors can be reduced by better engineering of the light fields, vacuum system, and trap structure, however at some level even intrinsic scatter off the ions themselves becomes an issue. In the near term, spectator ions can simply be

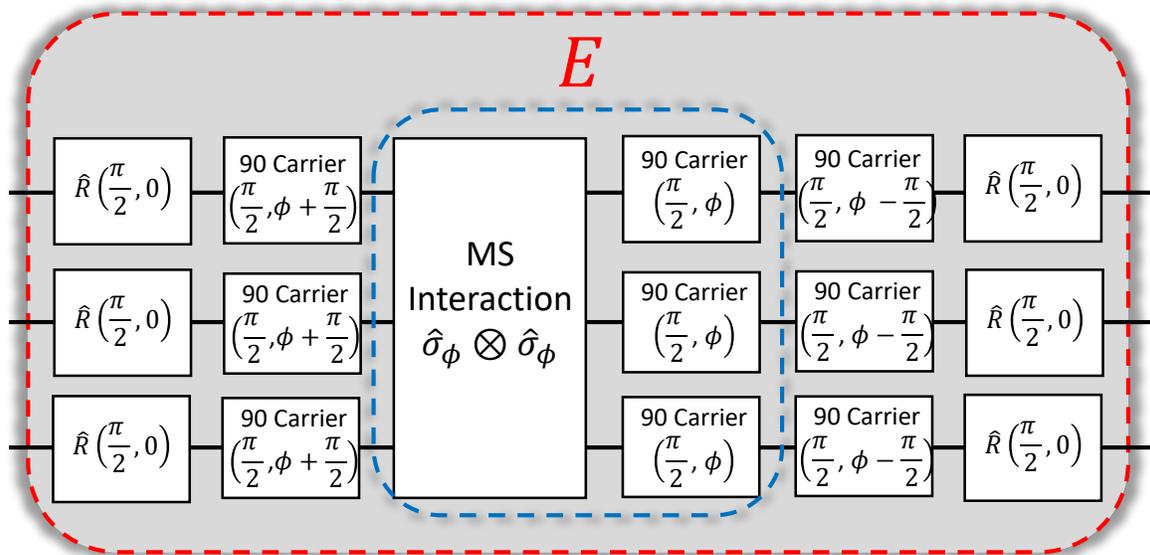


Figure 6.11: Preparing a state in the logical X basis requires control over the phase of the entangling interaction used to generate it. Like two-qubit entangling operations or CNOT gates (Sec. 4.7), the fluctuating interferometric phase must be accounted for. Operations in the central dashed blue box will produce a GHZ state, however with a fluctuating phase. Like two-qubit composite gates, this can be removed by embedding the operation in a Ramsey sequence using the same set of motion-sensitive Raman beams. The desired operation is then achieved by an additional basis change at the beginning and end of the sequence, driven by a phase-stable source. ‘90 Carrier’ refers to single qubit rotations using the same motion sensitive Raman beams as those driving the MS interaction.

shuttled sufficiently far away from detection zones, however this imposes limitations on how densely ions can be packed when scaling to larger systems, and at some point the required distance and overhead may become impractical. These type of workarounds can certainly give great benefit with small systems in the short run, however with larger systems in the long run it may be advantageous to instead trade additional overhead for looser engineering requirements. This concept is not new to quantum information processing; it is well established that at some point it becomes more practical to implement QEC to extend the logical depth of circuits than to insist on better and better engineered (and thus higher fidelity) operations.

While resonant light is necessary in some form, with two species of ion we can choose *which* is addressed with resonant light. We can therefore offload all tasks that require resonant light, such as laser cooling and qubit state initialization to the separate species readout ion. The multiframe can be used to indirectly measure a qubit through a readout ion by shelving one of the qubit states, performing the multiframe, and then measuring whether the readout ion flipped. If the qubit was in the shelved state, the readout ion will not flip, while if the qubit was in the unshelved state the readout ion will flip. This can potentially be repeated for increased fidelity (Sec. 7.2) [Hume et al., 2007], however the multiframe makes inefficient use of Raman beam power (see Sec. 7.4) so other more efficient techniques may be desirable if using lasers. In Chapter 7 we will discuss a more powerful protocol for indirect measurement with Raman beams by defining measurement subspaces that give additional tolerance to dominant sources of error [Elder et al., 2020]. Discrimination of subspaces will project the state of the ion after detection into those subspaces, which is perfectly fine for indirect measurement, but not as useful for indirect preparation of a pure fiducial state. Here we therefore outline a protocol for indirect state preparation using the multiframe, which, while costly for laser based protocols, is very promising if the multiframe can be driven with near-field microwave gradients on the qubit [Ospelkaus et al., 2008] (the readout ion can be driven through whichever means).

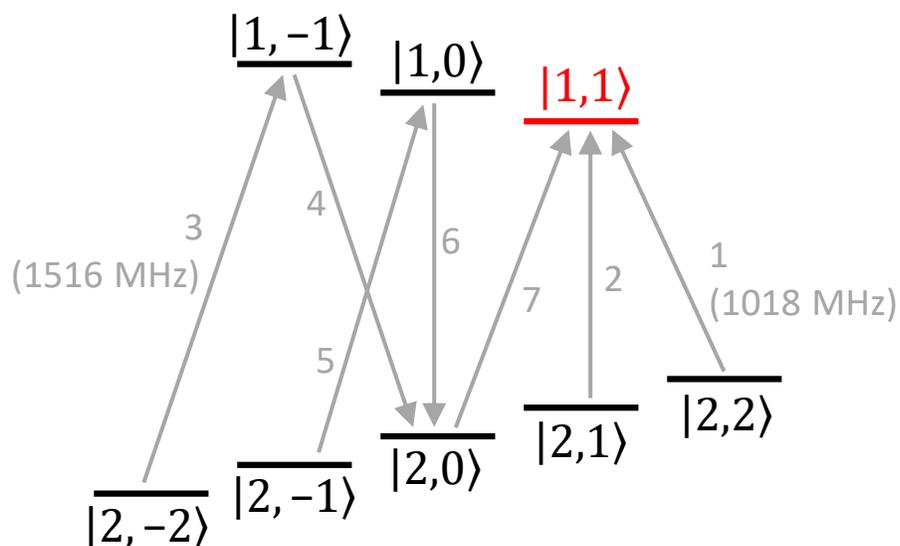


Figure 6.12: Beginning from an unknown state anywhere in the ${}^9\text{Be}^+ 2S_{1/2}$ manifold, a target initial state can be prepared with reasonably high fidelity (though still limited by the motional ground state fidelity), by ground state cooling with a separate species cooling ion, then iteratively driving motion adding sidebands on the ${}^9\text{Be}^+$ ion and then cooling that motion away. This serves to increase the chance of successfully projecting into the desired state with the protocol described in Fig. 6.13. Numbers show one possible ordering of the sidebands. Alternatively, if fast microwave transitions are available, iterative microwave mapping between all of the other states and a single sideband transition and could be used.

Indirect preparation can most simply be done on a system of one qubit and one readout ion. Before using the multiflip for projective state preparation, it is first useful to bring most of the qubit population to a particular known state. This can readily be achieved by ground state cooling with the readout ion, then iteratively driving motion adding sidebands on the qubit (to make population transfer ideally unidirectional), and then cooling any injected motion away with the readout ion. Since at most one quantum of motion is injected during this process, a single motion-subtracting sideband and repumping on the readout ion will suffice. Iterating this process across the entire ground state manifold will then bring most of the population into the chosen initial state. Fig. 6.12 shows one possible ordering for pumping to $|1, 1\rangle$. Depending explicitly on the motional ground state fidelity, this process only serves to increase the probability of subsequent high-fidelity projective measurements yielding the desired results.

The behavior of the multiflip depends on the participation parity of ions in the interaction. It can therefore be used to detect whether the qubit is in the two states connected by the resonant transition. This will projectively prepare an uncontrolled arbitrary superposition of those two states, however we need to prepare a pure state to be useful for quantum computation. Two applications of the multiflip are therefore necessary, as outlined in Fig. 6.13 and discussed below.

Assuming population begins distributed throughout the ground state manifold. While the iterative pumping in Fig. 6.12 will bring much of the population to the target initial state, in this case $|1, 1\rangle$, finite temperature will result in appreciable population being left behind. With a single qubit and single readout ion, the multiflip will measure whether the qubit is in the chosen transition for the gate, which we chose to be $|2, 0\rangle \leftrightarrow |1, 1\rangle$. In other words, if the qubit is in a state connected by that transition it will participate in the multiflip, causing the readout ion to flip. If the qubit is not, the readout ion will not flip. This detection is projective, so the appropriate result will herald qubit projection into the multiflip manifold (Fig. 6.13(b)). A single multiflip measurement may not have sufficient fidelity, in which

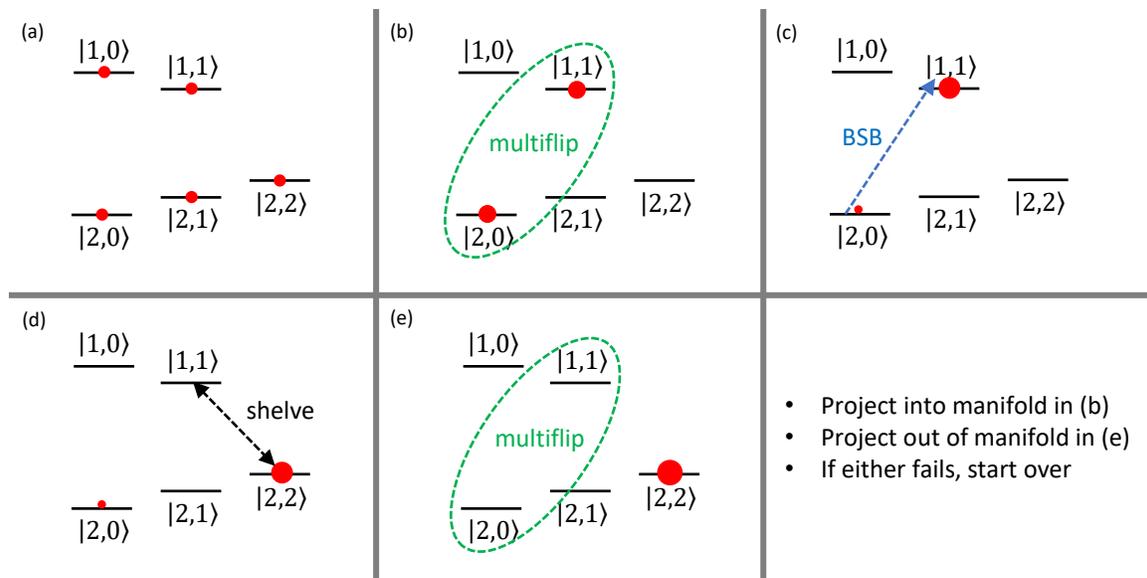


Figure 6.13: Sketch of a protocol that uses the multiflip for preparing a pure quantum state, shown on a truncated subset of the ${}^9\text{Be}^+ {}^2S_{1/2}$ manifold. (a) Assume population is distributed throughout the entire ${}^2S_{1/2}$ manifold, including the states not shown. The iterative sideband pumping procedure discussed in Fig. 6.12 can be used to accumulate population in $|1, 1\rangle$ to increase the probability of successful projection. (b) Through a separate species readout ion, use the multiflip to measure whether the qubit is in the $|2, 0\rangle \leftrightarrow |1, 1\rangle$ manifold. This can be repeated multiple times for increased confidence (see Sec. 7.2). If the measurement results indicate that the qubit is not in this manifold (indicated by the readout ion not flipping), then start over. (c) Heralded by the appropriate result in the previous frame, the qubit is projected into the chosen manifold. However, the multiflip says nothing about where population is located within that manifold. A blue sideband transition after ground state cooling will bring most of the population to $|1, 1\rangle$, limited by the ground state fidelity. (d) Shelve $|1, 1\rangle \rightarrow |2, 2\rangle$ (e) The multiflip is performed again, possibly repetitively, this time targeting measurement results that herald projection out of the chosen manifold. Given that out-of-manifold states had previously been projected out of in part (b), the only out of manifold state with population should be $|2, 2\rangle$. Since the multiflip learns nothing about where population is outside the manifold, it projects proportionally into the out-of-manifold states based on their existing populations. Hence given the right detection results, the qubit will be in $|2, 2\rangle$ with high probability.

case it can be repeated for increased confidence (see Sec. 7.2) [Hume et al., 2007]. This measurement is projective and hence probabilistic, so if the undesired measurement result is obtained the protocol must start over with the iterative pumping.

The multiflip provides no information about where in the driven manifold qubit population resides, and hence does not help distinguish between the two involved states. A motion adding sideband (Fig. 6.13(c)) can be used to collapse most of the population from $|2, 0\rangle$ to $|1, 1\rangle$, however the fidelity of the population in the target state $|1, 1\rangle$ is limited by the motional ground state fidelity. Hence additional steps must be taken. At this point we know with high confidence that the qubit is not outside of the manifold, and with reasonable confidence that it is in $|1, 1\rangle$. The population in $|1, 1\rangle$ is then transferred to $|2, 2\rangle$ with a carrier π -pulse.

Applying the multiflip again (Fig. 6.13(e)), repetitively if needed, heralds projection *out* of the driven manifold. If the result indicates projection into the $|2, 0\rangle \leftrightarrow |1, 1\rangle$ manifold, start over from the beginning. Population outside of the multiflip manifold was previously cleared out, and then injected only into $|2, 2\rangle$. The multiflip gives no information about where population is distributed outside the manifold, so it projects proportionally into the out-of-manifold states based on their existing populations. Therefore given successful projection out, the qubit should be in the $|2, 2\rangle$ pure state.

The success of this protocol depends on population not leaking out during the multiflip operations. If driven by Raman beams, this presents a significant barrier with little hope for state preparation infidelity lower than that leakage probability. For state detection, it is not necessary to end the process in a single pure state. It is sufficient to know with high confidence which state the qubit was projected into at the start of the detection process, while where it ends up is not as relevant. Motivated by this, we developed a detection protocol that can tolerate multiple spontaneous emission events that we introduce and implement in Chapter 7. It is not, however, applicable to preparing a single pure physical state of the ion. Fortunately, lasers are not the only way to couple ions to their motion. Techniques based on

near-field magnetic field gradients, rather than the electric field gradients from interfering laser beams, are particularly promising for indirect control because they are not subject to spontaneous emission.

One of the most promising approaches for entangling interactions with near-field microwaves [Sutherland et al., 2019] can be used to implement X-type or Z-type interactions. Hence an XX, XZ, or ZZ multiflip could be used for the QLS map, depending on how the readout ion is addressed. Any of these can be used in the indirect preparation protocol, with the slight modification of surrounding the multiflip with a Ramsey sequence on the readout ion if it undergoes a Z-type spin-dependent force (to convert a phase flip to a spin flip).

Chapter 7

High-Fidelity Indirect Measurement

7.1 Motivation

Resonant laser light is one of the key enabling technologies for trapped ion quantum information processing (QIP), allowing for necessary and controllable interaction with the surrounding environment. At the beginning of a computation it provides an entropy sink for state preparation through excitation to a short-lived excited state and subsequent spontaneous decay. At the end of a computation it is used for state readout by driving a state-dependent cycling transition and collecting the accompanying fluorescence [Leibfried et al., 2003a]. These two techniques together have been used to demonstrate very low state preparation and measurement (SPAM) errors in trapped ions [Myerson et al., 2008, Edmunds et al., 2020]. As the size and complexity of quantum processors increases, recooling is necessary to maintain high two-qubit gate fidelity, and mid-circuit measurements and preparations will be required for quantum error correction. Moreover, though not strictly necessary, mid-circuit measurement can be leveraged to expand capabilities in small near-term quantum processors [Kim, 2017, Foss-Feig et al., 2020, Rattew et al., 2020, Foss-Feig et al., 2021]. Finally, as discussed in Sec. 5.9.3, the resource entanglement during gate teleportation must be the primary qubit species to allow for high-fidelity and embedding it in a larger circuit. Therefore in both near-and long-term quantum processors, resonant laser light is expected to feature prominently all throughout the computation.

Trapped ions are not only an appealing QIP platform for their well controlled dis-

sipative interactions with the environment, but also for their high-fidelity coherent operations through which the processing occurs. Single-qubit rotation fidelities [Brown et al., 2011, Harty et al., 2014] already well surpass the threshold for most quantum error correction codes. Two-qubit gate fidelities [Ballance et al., 2016, Gaebler et al., 2016, Srinivas et al., 2021] are likewise above threshold for more forgiving codes [Fowler et al., 2012], though further improvement in the gates will allow for codes with much less overhead [Bermudez et al., 2019]. However, these experiments were all performed in isolated systems of one or two qubits, and as systems scale, unintended crosstalk on spectator qubits is as important as the intended impact on active qubits, and is particularly damaging to fault-tolerant quantum computation [Leibfried et al., 2004a, Sarovar et al., 2020, Parrado-Rodríguez et al., 2020].

For trapped ion qubits, laser spillover to neighboring qubits in the same potential well during gates and emitted photons being observed by the detector for another ion are two common forms of crosstalk. Naturally, careful study has gone into techniques for mitigating these effects. The impact of laser spillover can be reduced through composite pulse techniques [Parrado-Rodríguez et al., 2020, Herold et al., 2016], since these errors are coherent. Detection crosstalk can be reduced using maximum likelihood estimation [Burrell et al., 2010]. Alternatively, the QCCD architecture largely avoids either type of crosstalk by storing spectator qubits in distant potential wells [Pino et al., 2021].

Another form of crosstalk is decoherence due to scattering of resonant photons [Leibfried et al., 2004a, Bruzewicz et al., 2019], as observed in Sec. 5.5 during quantum gate teleportation, where a single errant photon scattering off a spectator qubit is enough to destroy its encoded quantum information. In that experiment stray resonant light errors were on the one percent level. However, carefully engineered experiments can and have demonstrated much lower rates [Crain et al., 2019]. However, low crosstalk per *spectator ion* can combine to high probability of crosstalk to one or more spectators per *detection*. This form of crosstalk is therefore particularly dangerous, because at least to some extent, it scales with the size and density of the register; the more qubits you have, the more that

could be hit by stray resonant photons. Hence at some point resonant light crosstalk could become a significant barrier to further benefit from quantum error correction.

Several approaches could be taken to avoid this crosstalk. Greater distances between detected and spectator qubits may help, though that is counter to the goal of scaling up the number and density of qubits in a register, and may have diminishing returns depending on the source of the stray light (e.g. it could be scattering from anywhere on the trap structure, vacuum system, etc. without a well known pattern). Even with the scattering pattern under control, the added overhead for transporting qubits that need to be together for multi-qubit gates to remote locations for readout will be unavoidable. One could shelve spectators prior to detection, however the errors in the (un)shelving process will accumulate, shelving to an excited state will suffer from spontaneous decay, and depending on the choice of shelving transition may introduce additional dephasing in the shelved period. A nice feature of the photon interconnects architecture for scalable trapped ion QIP is that it necessarily compartmentalizes this type of crosstalk, however with that comes the additional overhead and bottlenecks introduced by those interconnects. Finally, better engineering of the laser fields, trap structure, and vacuum system can mitigate this error, though it is unclear to what extent and at what cost.

Instead, this problem can be avoided entirely by trading the necessity for mixed-species operations for looser engineering requirements on the detection infrastructure. This can be accomplished by indirect measurement through a separate species readout ion whose resonances are far detuned from any relevant qubit-ion transition. A variant of this process is known as quantum logic spectroscopy (QLS) [Wineland et al., 2002, Schmidt et al., 2005]. The most basic and traditional form of QLS directly transfers population from the qubit to the motion, and then from the motion to a readout ion as shown in Fig. 6.1. While this introduces additional complexity, it is generally accepted that sympathetic cooling ions are necessary for complex algorithms, so the second species can serve a dual purpose and thus not require any additional ions than where otherwise needed. Moreover, no additional control

capabilities for the second species are necessary beyond those already used for sideband cooling, though the ability to perform geometric phase gates may be beneficial in some circumstances. In Sec. 7.4.1 we outline a protocol for indirect measurement of hyperfine qubits that ultimately enables higher detection fidelity than is otherwise feasible.

7.2 Repetitive Quantum Nondemolition Measurements

Mitigating resonant light crosstalk by indirectly reading out a qubit through a separate species introduces additional infidelity from the mapping process. However, fidelity and crosstalk are not the only figures of merit for a quantum measurement. How much the measurement affects the state of the qubit adversely after initially projecting it also matters. If it does not, the measurement is said to be quantum nondemolition (QND). Real measurements rarely perfectly satisfy this condition, so we refer to the extent to which the state is preserved as QND purity. At first the QND nature of a measurement may seem irrelevant, since quantum measurements are necessarily projective, however they allow the measurement to be repeated for increased confidence in the state of the qubit [Hume et al., 2007].

The idea of many low fidelity measurements yielding one composite high fidelity state detection may seem unfamiliar at first, however ion trappers have been exploiting this since the onset of the field. In typical fluorescence detection, each individual photon resonant with a cycling transition that encounters the ion can be considered a measurement. The first photon projects the state (the reason stray resonant light is so dangerous), and subsequent photons repeatedly measure it. A few thousand measurements occur, where if in $|\text{Bright}\rangle$ many photons will be scattered while very few are scattered from $|\text{Dark}\rangle$ (see Sec. 2.8). However, most of those scattered photons are not observed by our photomultiplier tube (PMT). The limited solid angle of the imaging system, losses through the optics along the way, and the quantum efficiency of the PMT result in very few of the scattered photons being observed. A single photon scattering off an ion is hence a *terribly* low fidelity measure-

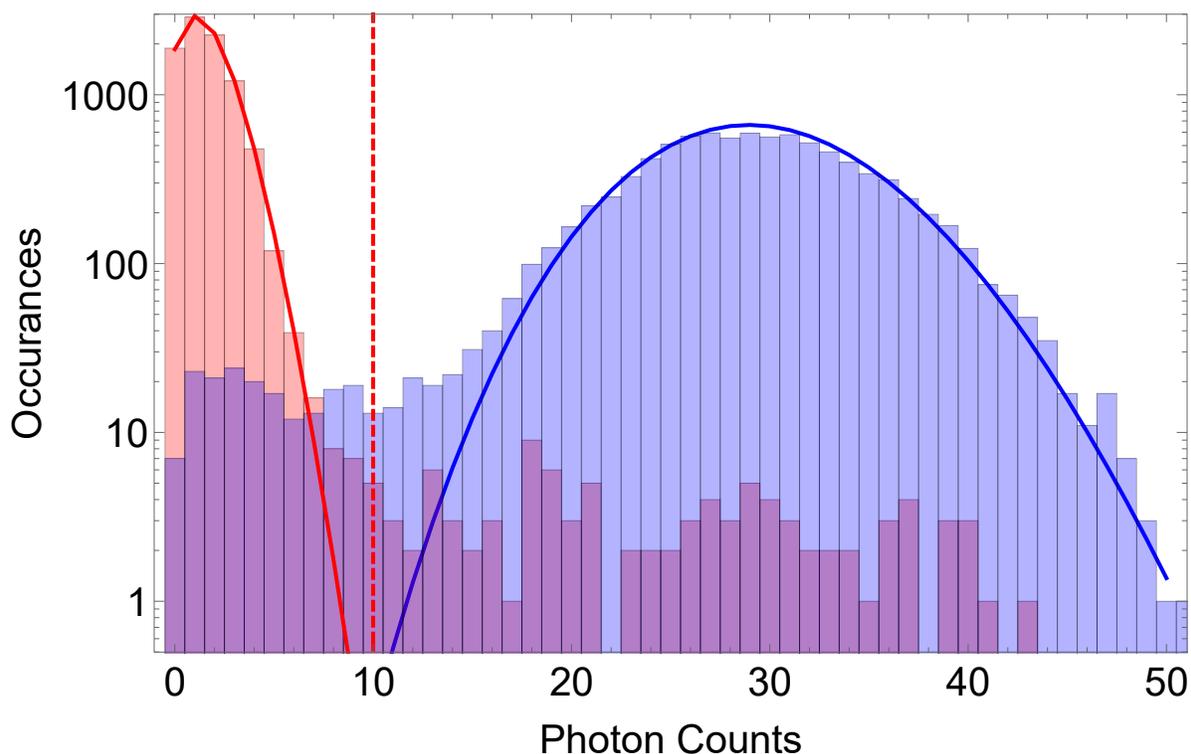


Figure 7.1: Histograms of photon counts for being in $|\text{Dark}\rangle$ (red) or $|\text{Bright}\rangle$ (blue) after preparing a superposition of the two with equal amplitudes. The two distributions are well separated, allowing for reasonably high fidelity direct qubit measurements in a single readout even by setting a threshold at 10 counts (red dashed line). However, the intermediate region around 10 counts shows more overlap than expected for ideal Poisson distributions, shown by solid line fits, due to QND impurity in the measurement.

ment, however assuming proper polarization control of the detection beam (for our choice of qubits) these measurements are *extremely* purely QND. For either state the probability of not observing a scattered photon is nearly one, however the probability of actually observing a photon is relatively larger for $|\text{Bright}\rangle$ than for $|\text{Dark}\rangle$.

After a few thousand of these single-photon measurements, this small discrepancy in probability of observing a photon will compound, resulting in approximately Poissonian count distributions that can be well resolved (Fig. 7.1). However, QND impurity results

in the distributions deviating from ideal Poissonians, making them overlap more strongly in the intermediate region which ultimately limits the fidelity of such a measurement. The main sources of impurity in our fluorescence measurements are a polarization error in the detection beam causing $|\text{Bright}\rangle$ to depump to $|\text{Dark}\rangle$, and $|\text{Dark}\rangle$ offresonantly repumping to $|\text{Bright}\rangle$. To some extent these effects can be mitigated through Bayesian techniques that incorporate photon arrival times [Langer, 2006, Myerson et al., 2008, Burrell, 2010]. Details of how we analyze a sequence of repetitive QND measurements will be discussed in the next subsection.

7.2.1 Bayesian Analysis

During standard fluorescence detection we do not know exactly how many single-photon measurements occurred. With repetitive QLS, there is no such restriction. We know exactly how many measurements we perform and get a result for every one of them. This allows for more powerful analysis techniques to get a composite detection result from the sequence of measurement rounds. Our approach is based on that of Ref. [Hume et al., 2007], where we use Bayesian inference based on reference data to determine composite detection outcomes.

Probability $P(n|i)$ of a single measurement outcome n , given qubit state i , is obtained from reference QLS data. This should ideally come from the same QLS procedure that is used repetitively for detection. The details of the protocol we implemented will be discussed in Sec. 7.5. For now we keep it general. Here n could be photon counts, but as will be discussed in Sec. 7.5 we perform two physical QLS mappings and readout ion measurements per analysis round. Each full round is then identical for analysis purposes. The qubit state i is either of two subspaces within the ground state manifold (Sec. 7.4.1).

For simplicity we only consider the two-bit outcome after thresholding each of the photon counts observed in a fixed detection duration, thus minimizing the amount of reference data needed for high statistics on each outcome. By doing this, we only need to characterize $2 \times 2 = 4$ reference parameters, compared to roughly $50 \times 50 = 2500$ if considering individual

photon counts and not truncating the bright Poisson distribution. Some benefit may be gained by grouping photon counts into $2 < n \ll 50$ bins for only a moderate increase in reference data. Detecting only a single ion at a time, it is likely that most of the benefit could be obtained from three bins, with a separate middle bin to cover the overlap between the bright and dark distributions.

From these single-round probabilities, the probability of a sequence of measurement results $\{n_j\}$ given state i is simply their product

$$P(\{n_j\}|i) = \prod_j P(n_j|i). \quad (7.1)$$

This gives $P(\{n_j\}|i)$, but we are interested in the reverse. Applying Bayes' rule yields just that:

$$P(i|\{n_j\}) = \frac{P(\{n_j\}|i)}{\sum_k P(\{n_j\}|k)} \quad (7.2)$$

where k is any possible qubit state.

The measurement result is assigned to the state $|i_{max}\rangle$ with the highest probability. The probability of measurement error can also be obtained from $1 - P(i_{max}|\{n_j\})$, which can be tracked in real time to end the sequence of measurements once a predetermined confidence threshold is reached. Doing so allows one to not only save time, but also limits the negative effects of an increasing cumulative probability of a QND impurity changing the qubit state (e.g. from spontaneous emission).

With a perfectly QND measurement, arbitrarily readout fidelity can be achieved for any nonzero single-shot fidelity with enough rounds of measurement. However, this is not to say that single-shot fidelity is not important. Higher single-shot fidelity requires fewer rounds of detection to reach a target total fidelity, thus saving time. Moreover, the fact that no measurement is perfectly QND means that a higher total fidelity can be reached before the compounding probability of changing states due to QND impurity outweighs any benefit that would otherwise be gained from additional rounds of detection.

Detection can be done through a fixed number of QLS rounds. However, as will be demonstrated in Sec. 7.7.1, it is often advantageous to actively track $P(i|\{n_j\})$ and stop once a predetermined confidence ratio threshold R_{thresh} is breached by either of two ratios

$$R_{\pm} = \frac{P(S_{\pm}|\{n_j\})}{P(S_{\mp}|\{n_j\})}. \quad (7.3)$$

Here S_+ and S_- are the two measurement subspaces for our QLS protocol, which will be defined in Sec. 7.4.1. The measurement outcome is then the more probable $P(S_{\pm}|\{n_j\})$, i.e. the one in the numerator of the breaching ratio, at the end of the sequence.

7.3 Extending From Optical to Hyperfine Qubits

Ref. [Hume et al., 2007] first demonstrated repetitive readout with traditional quantum logic spectroscopy (Fig. 6.1) and Bayesian updating of the most likely outcome in an $^{27}\text{Al}^+$ - $^9\text{Be}^+$ optical clock where the final state of the $^{27}\text{Al}^+$ clock was transferred to the $^9\text{Be}^+$ readout ion through the shared motion. Whether or not the $^{27}\text{Al}^+$ qubit was in the $^3P_0 = |\uparrow\rangle$ excited state was determined by repeatedly performing QLS on the $^1S_0 = |\downarrow\rangle$ ground state to 3P_1 auxiliary transition (Fig. 7.2). This transition is long-lived enough compared to the sideband π -time to allow for probing, but short lived enough that after transferring information to the shared motion the $^{27}\text{Al}^+$ population will decay back down to $|\downarrow\rangle$ prior to the next QLS iteration. Transferring information directly to the motion limits the fidelity of each QLS round by the motional ground state fidelity. Each individual detection cycle in this experiment therefore had approximately 15% error, however through repetitive measurements and the analysis described in 7.2.1, high QND purity enabled reaching approximately 99.94% fidelity.

In the $^{27}\text{Al}^+$ optical qubit, the main source of QND impurity is the roughly 21 second lifetime of the $|\uparrow\rangle$ excited state, while for hyperfine qubits like $^9\text{Be}^+$, the dominant source is qubit spontaneous emission during the QLS map, if driven with Raman beams. Scattering on the readout ion only impacts single round fidelity. Microwave quantum logic gates [Ospelkaus

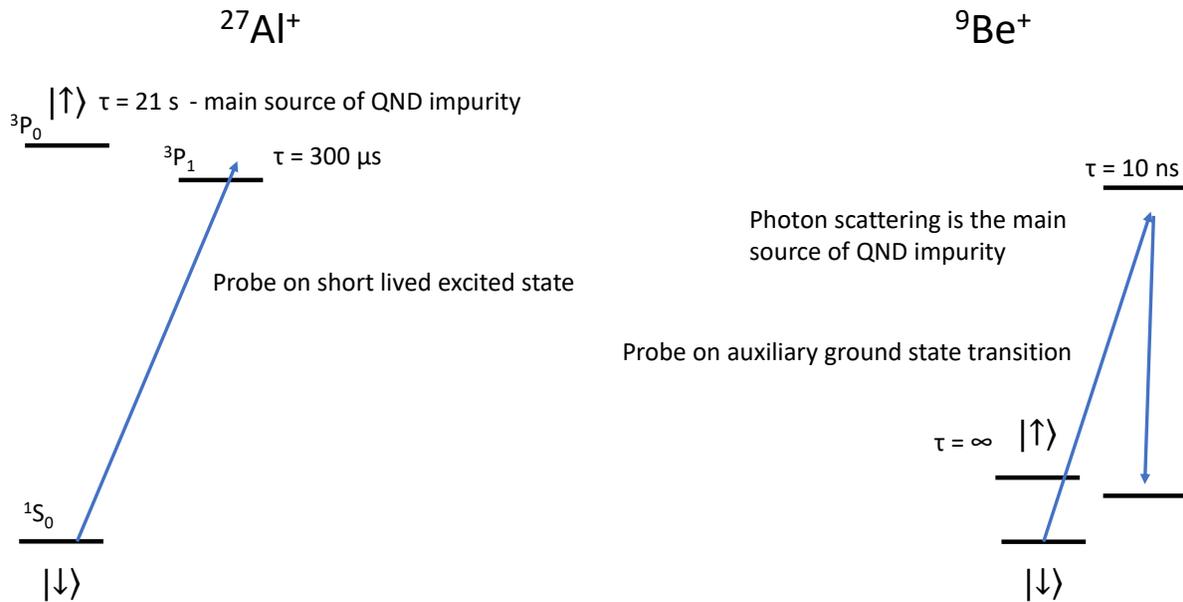


Figure 7.2: Simplified comparison of $^{27}\text{Al}^+$ and $^9\text{Be}^+$ ion energy level as they pertain to repetitive QND measurements. $^{27}\text{Al}^+$ is probed on a relatively short-lived excited state transition that will subsequently decay back to the ground state prior to the next round of QLS. Spontaneous photon scattering during the probe is of no concern. The dominant source of QND impurity is the long but finite ~ 21 s lifetime of the 3P_0 state. $^9\text{Be}^+$ ions have no comparable transitions. Instead stimulated Raman transitions can be driven to an auxiliary ground state to distinguish one qubit state from the other. The dominant source of QND impurity is photon scattering during QLS on the $^9\text{Be}^+$ ion to a state outside the transition being driven. One additional difference is that the auxiliary state in $^9\text{Be}^+$ will not decay back to one of the qubit states, so the transition has to be probed in both directions. More details on this final point will be given in Sec. 7.4.1.

et al., 2008, Harty et al., 2016, Weidt et al., 2016, Srinivas et al., 2021] avoid this source of impurity entirely, potentially allowing for extremely high readout fidelity. Next leading sources of QND impurity would likely be related to the vacuum (collisional state change or ion loss) or off-resonant coupling to spectator transitions. The former can be minimized in a cryogenic system, and the latter by decreasing the Rabi frequency and by exploiting the polarization dependence, so far higher readout fidelities should be achievable with microwave coupling.

One detail that needs to be modified for this method to be applicable to hyperfine qubits is a way to reset population driven to the auxiliary state during traditional QLS. In the $^{27}\text{Al}^+$ example, the 3P_1 auxiliary state has only a roughly $300 \mu\text{s}$ lifetime, so it decayed back to $|\downarrow\rangle$ with high certainty during subsequent sympathetic recooling on the readout ion. Hyperfine qubits do not share this feature, so the auxiliary state needs to be manually reset. Driving traditional QLS the opposite direction on the transition accomplishes this (i.e. using the opposite sideband of the first probe), and can be used to add a second bit of information about the state of the qubit by using the readout ion to measure whether the second probe injected a quantum of motion. The downside, however, is that if one QLS probe fails the next one will also most likely fail, countering to some extent the benefit of the additional bit of information. While having a chance to disrupt the subsequent probe, the action of the two direction QLS probe serves to pump population back to the initial state with high fidelity, so errors at most impact only the subsequent round, not the entire sequence of repetitive QND measurements. If this idea seems difficult to follow in the abstract, see Sec. 7.5 for details on how we implement it in our experiment.

7.4 Minimizing Spontaneous Emission

Since we work with lasers and not near-field microwaves on this apparatus, spontaneous emission is our main source of QND impurity. Minimizing it is therefore a top priority for achieving high fidelity indirect detection. Naturally, the easiest step is to use as large of

Mode	Frequency (MHz)	η_{Be}	η_{Mg}
In-phase	1.36	0.21	0.36
Out-of-phase	2.91	0.37	0.097

Table 7.1: Approximate axial mode frequencies and LD parameters for a Be-Mg crystal with weak axial confinement.

a Raman detuning Δ as feasible, since the scattering rate scales roughly as $1/\Delta^2$ [Ozeri et al., 2007]. Coupling strength to the motion scales with the Lamb-Dicke (LD) parameter $\eta = \Delta k z_0$, where $z_0 = \sqrt{\hbar/2m\omega}$ is the ground state wavefunction extent and Δk the Raman beam differential wave vector projection along the chosen motional mode (Eq. 2.12). The total scattering rate does not depend on η , therefore a large value will allow for less scatter per sideband interaction. The first component, Δk , can be increased by having the two Raman beams counterpropagate along the chosen normal mode. The second, z_0 , can be increased by reducing the strength of the harmonic confinement.

Table 7.1 shows the axial mode frequencies and LD parameters for both species in a Be-Mg crystal with weak axial confinement. We choose to work with the out-of-phase mode at approximately 2.91 MHz because of its large $\eta_{Be} \approx 0.37$, at the expense of weak $^{25}\text{Mg}^+$ coupling with $\eta_{Mg} \approx 0.097$. This makes efficient use of our $^9\text{Be}^+$ Raman beams, at the expense of a high spontaneous emission rate on $^{25}\text{Mg}^+$. Our choice of imbalanced mode makes high-fidelity Mølmer-Sørensen (MS) based QLS maps more difficult, and with lower single-round fidelity due to increased $^{25}\text{Mg}^+$ spontaneous emission [Tan et al., 2015]. Moreover, MS based protocols come with roughly $\sim 2.5 - 3$ times more $^9\text{Be}^+$ scatter per bit of mapped information compared to traditional QLS. This is due to driving two sidebands instead of one and either including additional wrapper pulses to mitigate fluctuating interferometric phases if constructing a CNOT gate (Sec. 4.7) or having to drive the interaction for $\sqrt{2}$ times as long if using the multiflip (Sec. 4.5). Finally, Raman beam polarizations should be optimized to only include components that drive the chosen transition, and not any other

components that contribute to spontaneous emission but not the QLS.

7.4.1 Containing Scattering within a Recoverable Subspace

Reducing the scattering rate goes far in minimizing the QND impurity, however a single scattering event can still change the qubit state or leak out of the transition being driven. The right choice of transition can reduce this probability by restricting the number of other physical states that could be scattered to. For example, if QLS is driven on the $|2, 1\rangle \leftrightarrow |1, 1\rangle$ transition with two balanced σ^+ -polarized Raman beams, the only external state that can be scattered to is $|2, 2\rangle$ (see Fig. 7.3). We call this the $\Delta m_F = 0$ configuration. Alternatively, with a strong lower energy σ^+ -polarized beam and a weak π -polarized beam, what we call the $\Delta m_F = 1$ configuration can drive QLS on the $|2, 2\rangle \leftrightarrow |1, 1\rangle$ transition. It is convenient to use the same set of Raman beams for qubit manipulations and readout, so a corresponding set of qubit states should be chosen. The $\Delta m_F = 0$ configuration is compatible with the $|2, 1\rangle \leftrightarrow |1, 1\rangle$ qubit transition that is field insensitive at $B \approx 22.307$ mT, and the $\Delta m_F = 1$ configuration is compatible with the $|2, 1\rangle \leftrightarrow |1, 0\rangle$ qubit transition at $B \approx 11.964$ mT (Sec. 2.1).

Both configurations ideally contain scatter within the subspaces $S_+ \equiv \{|F, m_F \geq 1\rangle\}$ and $S_- \equiv \{|F, m_F \leq 0\rangle\}$. The $\Delta m_F = 0$ configuration requires a σ^+ beam polarization impurity to scatter out of the S_+ subspace, while the $\Delta m_F = 1$ configuration opens up an additional pathway to S_- by scattering with the weak π beam. On the other hand, if shelved to $|2, -2\rangle$, S_- requires multiple scattering events to change to S_+ . If an entire subspace can somehow be incorporated into the QLS probe itself, then QND purity between the two subspaces can be maintained even after otherwise destructive Raman scattering events, thereby enabling extremely high fidelity. Turning a subspace measurement into a qubit measurement is then simply a matter of appropriate shelving. The $\Delta m_F = 0$ configuration uses the qubit transition itself for readout, so prior to measurement one should shelve $|\uparrow\rangle = |1, 1\rangle \rightarrow |2, -2\rangle$. The $\Delta m_F = 1$ configuration uses a different transition for

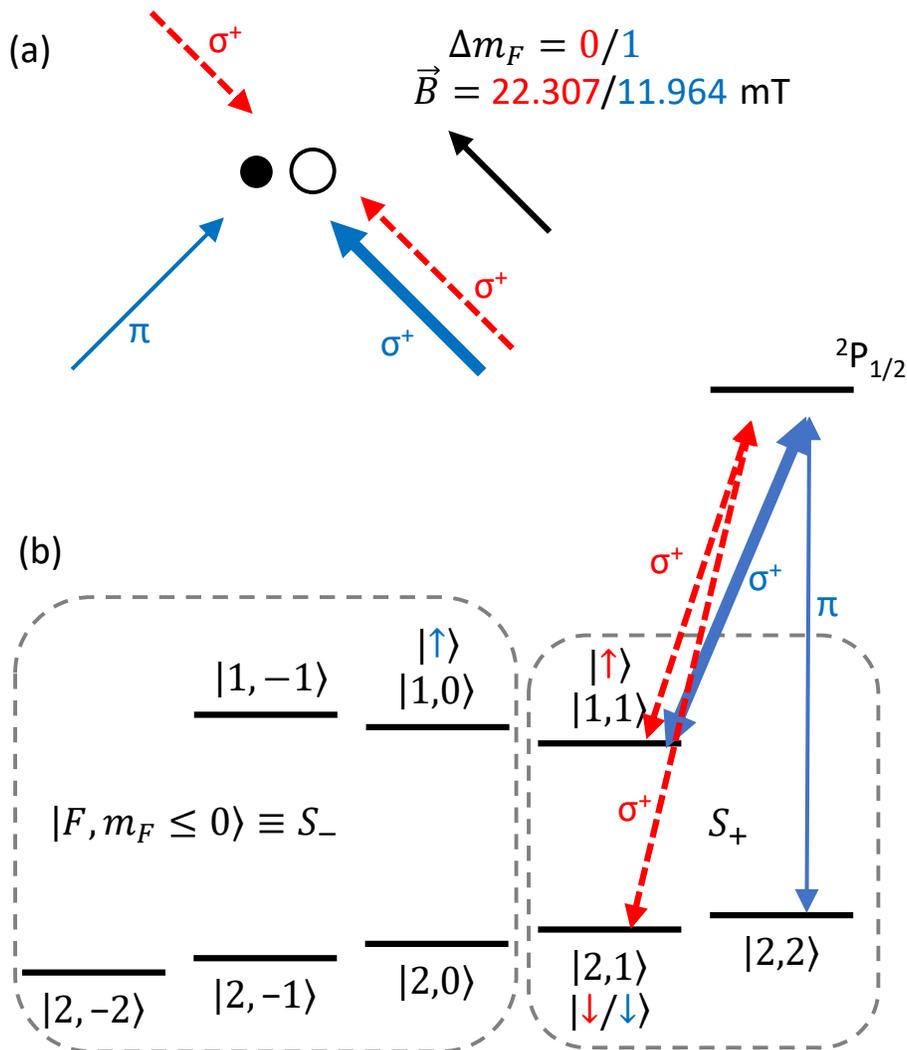


Figure 7.3: (a) Raman beam configurations for containing spontaneous emission within a subspace of the $^2S_{1/2}$ manifold. Two possible configurations are shown, color coded and labeled by the change in magnetic quantum number m_F that they can drive. Dashed red arrows represent counterpropagating σ^+ Raman beams for driving QLS on the $\Delta m_F = 0$ $|2, 1\rangle \leftrightarrow |1, 1\rangle$ transition. Solid blue arrows show the $\Delta m_F = 1$ configuration, with their thicknesses qualitatively representing the difference in intensities. (b) $^9\text{Be}^+ \ ^2S_{1/2}$ manifold divided into subspaces, where S_+ is the one where sideband transitions occur. Associated field-insensitive qubit transitions (Sec. 2.1) that require the same Raman polarizations are labeled appropriately.

readout, so one would shelve $|\downarrow\rangle = |2, 1\rangle \rightarrow |2, 2\rangle$ and $|\uparrow\rangle = |1, 1\rangle \rightarrow |2, -2\rangle$.

7.4.2 Experimental Limitations

Due to fewer pathways for QND impurity and more efficient use of Raman beam power, we consider the $\Delta m_F = 0$ variant superior, however limitations of our apparatus restrict us from easily implementing it. As shown in Fig. 3.1, one of the entry ports to our trap that is parallel to the magnetic field is occupied by all of our resonant and photoionization beams. Though possible, this comes with the additional Raman beam inefficiency (and hence higher QND impurity) from projecting the Raman beam Δk vector onto the trap axis. We considered injecting a Raman beam into this port to not be worth the additional work that would be required for our proof-of-principle demonstration when the alternate $\Delta m_F = 1$ variant is still available. In principle the magnetic field could be reoriented along the trap axis, allowing for counterpropagating pure σ^+ Raman beams along it, however not in a way where our system has the necessary optical access.

Implementing the $\Delta m_F = 1$ variant still required changes to our Raman beams compared to those discussed in Sec. 3.6.2. Fig. 2.1 depicts a higher frequency σ^+ -polarized and lower frequency π -polarized Raman beam. This is usually the case in our apparatus, however the $\Delta m_F = 1$ transition shown in Fig. 7.3(b) requires opposite polarizations. We found it easier to swap Raman beam entry ports than to redo the AOM setup that generates our two Raman tones (Fig. 3.8). Other experiments we were working on had different polarization constraints and had to be put on hold, so this also made reversing the change straightforward when we were ready to resume those experiments.

7.5 QLS Protocol

The circuit for one round of the QLS protocol is shown in Fig. 7.4. Each individual round begins with a fluorescence check on the readout ion to ensure crystalization. If the observed photon counts are below a set threshold, additional $^{25}\text{Mg}^+$ cooling is applied to

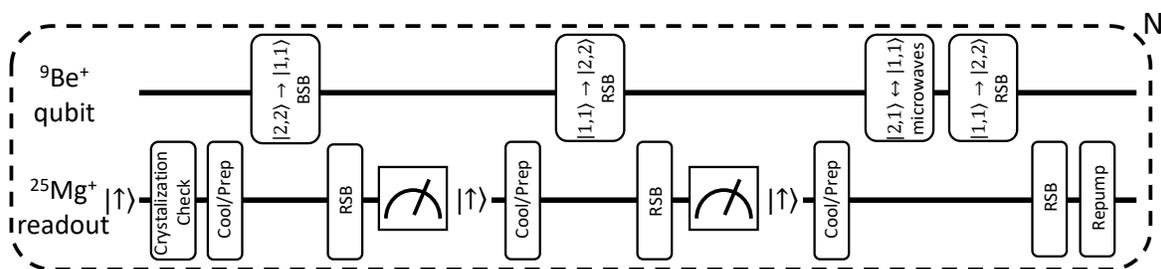


Figure 7.4: Circuit for one round of the QLS protocol, driving various transitions in S_+ . First, crystalization is verified through $^{25}\text{Mg}^+$ readout ion fluorescence, followed by ground state cooling. Most likely, the $^9\text{Be}^+$ qubit is initially in $|2,2\rangle$ if in S_+ . If not, the protocol serves to pump it there from the remaining states in S_+ for the next iteration. Two quanta of motion are likely injected by the qubit and measured by the readout ion if the qubit is in S_+ . If in S_- , it is unlikely that any quanta are observed. After the second measurement on the readout ion, population scattered to the inactive state in S_+ ($|2,1\rangle$) is brought back to $|2,2\rangle$.

attempt to recrystallize the ions prior to a second fluorescence detection. Counts again below threshold on this second detection indicate a failure. We next perform two probes of S_+ (yielding a two-bit outcome), followed by a recovery procedure for population that may have scattered to $|2, 1\rangle$. Grouping the probes this way is done for analysis purposes to ensure that each observation in our Bayesian analysis (Sec. 7.2.1) comes from an identical process, and only the entire protocol on maintains QND purity on S_+ . Rounds can be repeated a predetermined number of times or until a target likelihood ratio threshold is reached.

If in S_+ , the qubit ideally begins in $|2, 2\rangle$. As will subsequently be described, the action of the protocol serves to pump population within S_+ towards $|2, 2\rangle$. After a crystalization check, the readout ion cools the collective motion to the nominal ground state. The first probe then begins with a qubit $|2, 2\rangle \rightarrow |1, 1\rangle$ BSB, which injects a quantum of motion into the OOPH mode if the qubit was in $|2, 2\rangle$. This quantum is then detected with a readout ion motion-subtracting sideband followed by fluorescence measurement, where the presence of a motional quantum is correlated with a spin flip.

At this point of the protocol, the qubit nominally is in $|1, 1\rangle$. If it had begun there, the first qubit sideband leaves it there if the motion was successfully cooled to the ground state. Failure to ground state cool allows the first sideband to act backwards (i.e. motion-subtracting) and transfer population that started in $|1, 1\rangle$ to $|2, 2\rangle$. In this case it would impact the rest of this round of the protocol, but be in the proper starting place for the next round. Any population that may have scattered to $|2, 1\rangle$ is untouched by this operation.

The second probe is therefore performed under the assumption that the qubit population is likely in $|1, 1\rangle$ if it is within S_+ . After recooling and repreparing the readout ion, injecting a second quantum of motion is attempted with a qubit $|1, 1\rangle \rightarrow |2, 2\rangle$ RSB (motion adding if in $|1, 1\rangle$). This quantum is then detected with the same motion-subtracting sideband and fluorescence detection on the readout ion. If the qubit population instead was in $|2, 2\rangle$, this second sideband will leave it there for the start of the next round (again limited by motional ground state fidelity).

Both probes leave population in $|2, 1\rangle$ untouched, though have the chance of scattering to there. To maintain the QND purity of this protocol we therefore need a way to reincorporate population in $|2, 1\rangle$. The second probe nominally clears population out of $|1, 1\rangle$, so we transfer population from $|2, 1\rangle$ to $|1, 1\rangle$ with a microwave π -pulse. After cooling with the readout ion, this population then gets recombined with the bulk population in $|2, 2\rangle$ by a $|1, 1\rangle \rightarrow |2, 2\rangle$ RSB that is unidirectional if starting in the motional ground state (going the other way would require subtracting a ideally nonexistent quantum of motion). Since this is expected to be a rare occurrence, this potential quantum of motion is simply cooled away rather than measured. This recovery operation allows population to leak back out of $|2, 2\rangle$ if the motion is not in the ground state, which would impact the next round of measurement, but would be subsequently pumped back to $|2, 2\rangle$ with high probability for later rounds.

This ends one round of the protocol, where if in S_+ there is a high probability of detecting two quanta of motion, while if in S_- there is low probability of detecting any since all the qubit operations are far off resonance. More importantly, there is a very low probability of changing subspaces during this probe, since most scatter is contained within the subspaces, allowing for it to be repeated many times.

7.6 Data Filtering

Our primary focus was the capabilities of the protocol itself, not how well we had engineered our apparatus to reliably implement it on demand. We therefore apply a number of independent filters for experimental defects that would hinder the QLS, and hence the fidelity of the detection, but could be avoided in better engineered systems. Any experiments that fail these filters are not included. For example, if any of the fluorescence checks at the beginning of each QLS round fail, for example due to decrystallization from background gas collisions, that experiment is discarded. Such collisions are comparatively common in our room temperature vacuum system, but not in cryogenic systems where those background gases freeze out. We therefore do not consider this a fundamental limitation to the detection

protocol.

Similarly, we check the fluorescence of both species at the beginning and end of an experiment (each consisting of many QLS rounds), where if either fails for one or both species that experiment is discarded. All of these fluorescence checks combined serve to validate that both ions were there at the beginning of the experiment, remained crystalized, and were not lost. This covers defects related to the ions themselves.

We also want to protect against defects in our control. As discussed in Fig. 3.5, our Raman and detection beam intensities are all actively stabilized (and Raman beam pulse amplitudes are shaped). Servos for pulse control failing to lock will impact the Raman transitions central to the QLS, but is a technical rather than fundamental problem. Our servos have digital outputs that indicate whether they are locked, so we use those signals to filter out experiments. Likewise, we are interested in whether other resonant beams are present in the experiment. Since all resonant repump and BDD beams are not actively stabilized, we have nothing monitoring whether they are present and sufficiently strong. Instead, we read in the status of the resonant doubling cavity locks, trusting the passive stability of the beamlines after that. For these beams, precise pulse timing is not necessary, so small drifts are tolerable. Finally, we also filter out experiments where our magnetic field servo was unlocked.

Due to the limited numbers of TTL inputs in our ARTIQ control system (Sec. 3.5), we cannot read in each servo status individually. We therefore combine all servo statuses into a single signal by taking the logical OR of all of them in a custom circuit, then read in that combined signal. This gives a shot-by-shot record on whether one or more of these servos failed, but not which servo. While the limited information is sufficient for filtering out erroneous experiments, future control system upgrades with more TTL channels could provide better diagnostics by reading in each servo individually.

All of the filters discussed so far have single-experiment resolution. When executing an algorithm on a quantum computer flagging failures of the classical control can be used for

error mitigation through post-selection. However, in characterizing the protocol we are also interested in protecting against defects that may not be identifiable in a single experiment, but still impact the QLS in a technical but not fundamental manner. For example, if a $^{25}\text{Mg}^+$ repumper changes frequency by jumping iodine lock fringes, ground state cooling will not work properly and the sideband transitions in the QLS protocol will have reduced, but non-zero, fidelity.

To guard against this, and other effects that cannot be resolved shot-by-shot, we test a single round of the QLS protocol prior to every experiment after preparing each subspace ($|2, 2\rangle$ for S_+ and $|2, -2\rangle$ for S_-). This also ensures that measuring S_- is not successful simply because the QLS stopped working. A running average over 100 experiments is kept for the fraction of time these QLS tests give ideal result (assuming perfect cooling and operations). If at any point that average falls below a preset threshold the entire window is discarded. The thresholds were set low enough that statistical noise would be unlikely to breach them often (thereby wasting otherwise good data), but high enough to be triggered should only a small fraction of the window be compromised. An example of the running averages for the 490 GHz Raman detuning data run are shown in Fig. 7.5.

7.7 Measurement Results

Before presenting results on this detection protocol, we must first define how we identify a detection error. Our state preparation error is expected to be much larger than the detection error, so simply measuring a combined state preparation and measurement (SPAM) error does not tell us anything useful about the detection fidelity. Though non-deterministic, the detection protocol provides excellent heralded state preparation into one of the subspaces. We therefore estimate the detection fidelity by comparing two back-to-back detections, without doing anything to the qubit in between. The first detection serves to prepare the qubit to either subspace with very high fidelity, and the second to test the detection protocol. A readout ion fluorescence result corresponding to the qubit being in S_+ is assigned a value of

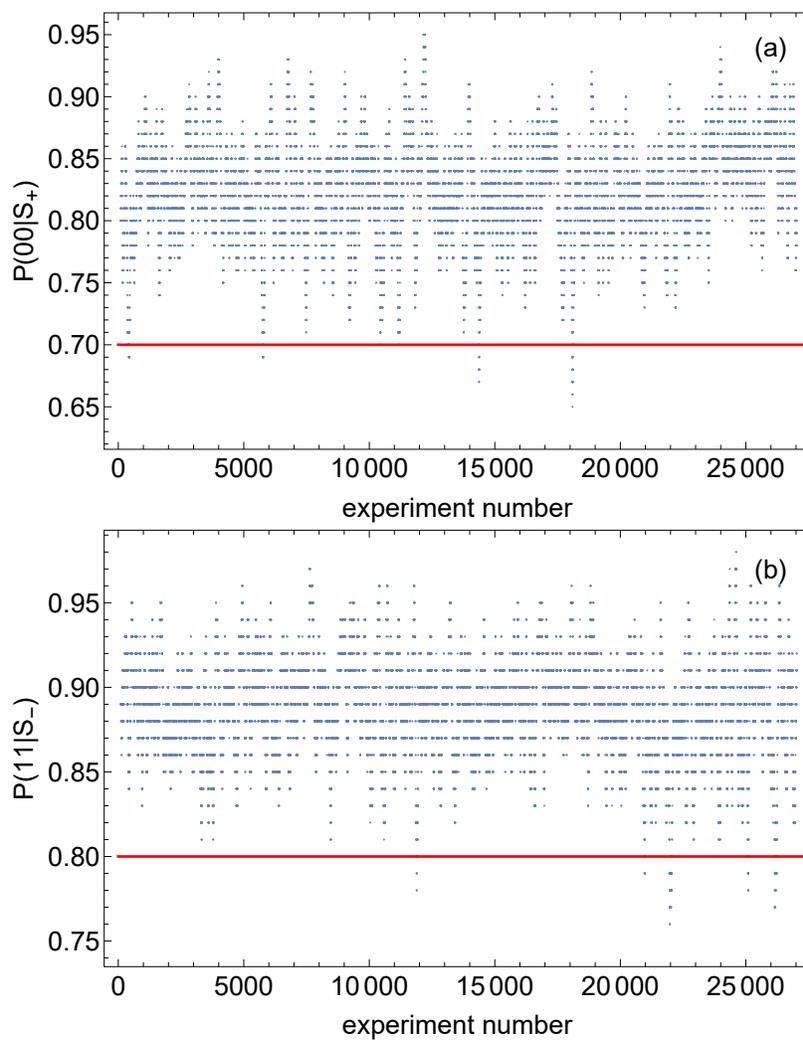


Figure 7.5: Running averages over 100 experiments of the fraction of time QLS tests performed before each experiment return the ideal expected result after preparing (a) S_+ and (b) S_- . Red lines show the thresholds for which if the running averages fall below the entire window is discarded. The two-bit results 00 and 11 relate to $^{25}\text{Mg}^+$ fluorescence results during QLS that both correspond to S_+ or S_- , respectively.

zero, and a result corresponding to being in S_- a value of one. We refer to the two results r_1 and r_2 in binary form r_1r_2 , which yields four possible outcomes per round of detection. These four outcome results are treated as a single unit in our Bayesian analysis to ensure that each result comes from an identical process (Sec. 7.2.1). The binary outcomes of the two detection sequences in each experiment are discussed similarly.

In this approach, the QND impurity of the first detection sequence cannot be distinguished from a detection error from the second. They two detections are identical, so we assume that the QND impurity and detection error rates are the same for both. Therefore, we report our infidelity as the sum of the QND impurity plus detection error, which we obtain from that rate at which the second measurement disagrees with the first. This sum then serves as an upper bound for both of its constituents. Explicitly, the errors for measuring the two subspaces and the average measurement fidelity are

$$\begin{aligned}\epsilon_+ &= P(01|r_1 = 0) \\ \epsilon_- &= P(10|r_1 = 1) \\ \bar{\epsilon} &= \frac{\epsilon_+ + \epsilon_-}{2}\end{aligned}\tag{7.4}$$

7.7.1 Test Data at 45 GHz Raman Detuning

Given the anticipated very low error rates, we first characterize the protocol at low Raman detuning where the main source of QND impurity (spontaneous emission) is greatly amplified. We choose to first operate at 45 GHz detuning. This is a factor of 20 less in detuning, and hence roughly a factor of 400 in scattering rate, than was used for the high-fidelity entangling gate experiment performed on this same apparatus [Gaebler et al., 2016]. The available Raman power in principle allows for an extremely large power imbalance in the two beams, but other constraints limited it in practice. The π -beam requires a minimum power to reliably shape and servo it, and too strong of a σ^+ -beam increases off-resonant carrier excitation. Ultimately we chose a ratio of 35:1 for this test.

Collecting a statistical sample with enough trials to certify errors of order 10^{-3} takes

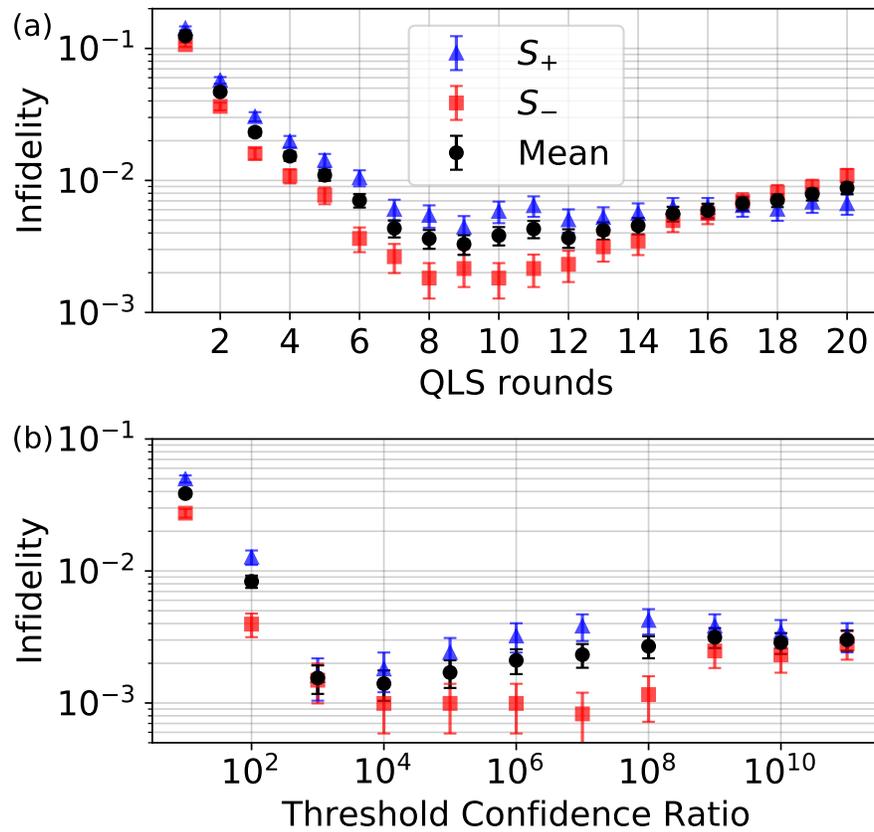


Figure 7.6: Test data taken at 45 GHz Raman detuning. (a) Infidelity for measuring S_+ , S_- , and their average are shown by blue triangles, red squares, and black circles, respectively, as a function of the fixed number of QLS rounds. (b) Infidelities as a function of the threshold confidence ratio used for adaptive detection (Sec. 7.2.1).

several hours, so implementing and testing many variations of the protocol is cumbersome. To alleviate this problem, we take a single large dataset, and then analyze it various ways in post-processing. We perform experiments with 40 rounds of QLS, and then analyze it as though we had decided to stop after $1 \leq n \leq 20$ rounds for each of two back-to-back detections. Any additional rounds after the second detection were ignored. Fig. 7.6(a) shows the infidelity for measuring either subspaces and the mean infidelity as a function of number of QLS iterations per detection. Infidelity after a single iteration is dominated by the single-round QLS fidelity, however we see that with additional iterations infidelity rapidly decreases. Eventually the increasing cumulative probability of changing subspaces due to the QND impurity outweighs any further benefit that would otherwise be derived from additional QLS rounds, causing the infidelity to plateau and then slowly rise. The average error reached a minimum of $3.3(6) \times 10^{-3}$ after nine rounds, biased towards measuring S_+ .

Measuring for a fixed number of rounds is a useful diagnostic, but is rather wasteful in most instances. Not all equal-length sequences of measurement results reach the same level of confidence in their result (Eq. 7.3). For example, the sequences $\{00, 00, 00, 00, 00\}$ and $\{00, 01, 10, 00, 00\}$ both correspond to S_+ , but the former will do so with higher confidence. Similarly, different detection sequences will reach a given target confidence in the result after varying numbers of QLS rounds. The sequences $\{00, 00, 00, 00, 00\}$ and $\{00, 11, 00, 01, 10, 00, 00, 00, 00\}$ both ultimately reach a similar level of confidence in measuring S_+ , but the latter takes much longer to do so. With an increasing cumulative probability of changing subspaces during the measurement, it is therefore highly beneficial to stop once a desired confidence level is reached.

We analyzed the same test data adaptively in post-processing. Fig. 7.6(b) shows the resulting infidelities plotted against the threshold confidence ratio for being in one subspace over the other. Not only is a lower error of $1.4(4) \times 10^{-3}$ reached with a 10^4 confidence ratio, it requires an average of only 3.47 QLS rounds. Higher single-round QLS fidelity would result in even fewer rounds needed for a given confidence level, Histograms of the number of rounds

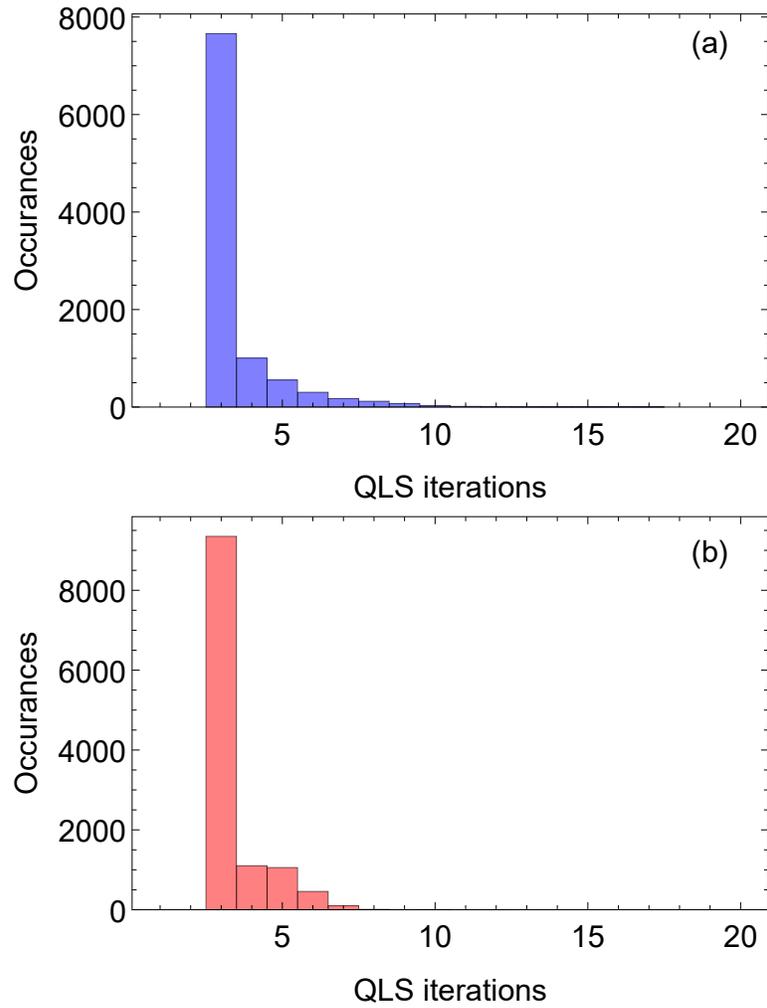


Figure 7.7: Histograms of number of rounds needed to reach a confidence ratio (Eq. 7.3) of 10^4 in the 45 GHz Raman detuning test data for (a) measuring S_+ , and (b) measuring S_- .

taken to reach this 10^4 ratio are shown in Fig. 7.7. Very rarely are large numbers of rounds necessary, and most of the time three suffice with our single-round fidelity. The adaptive method therefore clearly outperforms the fixed method, providing both an improvement in fidelity and time. We therefore programmed our FPGA to perform the adaptive analysis on the fly and to terminate each run in real time once the threshold is breached, allowing for much faster characterization of the high-fidelity measurements discussed in the next section.

7.7.2 High-Fidelity Measurement

With the adaptive measurement programmed to be executed in real-time, we first test it in similar conditions to the test data in Sec. 7.7.1 with the same 45 GHz Raman detuning and a roughly 35:1 Raman beam power ratio. Using the optimal 10^4 threshold confidence ratio obtained from Fig. 7.6(b), we measure a mean infidelity $\bar{\epsilon} = 1.7(6) \times 10^{-3}$ after an average of 3.55 QLS cycles (Table 7.2). Both the fidelity and number of rounds are consistent with the post-processed test data.

Increasing the detuning to 90 GHz allowed for a larger Raman beam power imbalance. Again limited by the constraints of keeping the weak π -beam strong enough to reliably servo and the total Rabi rate low enough to avoid off-resonant carrier coupling, we chose a ratio of 120:1. Expecting an increase in fidelity, we also raised the target confidence threshold to 10^7 to ensure that no further benefit would be derived from additional QLS rounds (i.e. that we are in the plateau region exemplified in Fig. 7.6(b)). The fidelity for measuring S_+ was $\epsilon_+ = 8(3) \times 10^{-4}$, which is a factor of four better than the 45 GHz data. This is roughly what one would expect considering just the detuning, since the scattering rate scales inversely with the detuning squared. Reaching a higher fidelity requires more rounds of QLS, resulting in a greater cumulative probability of changing subspaces during the measurement that suppresses this benefit slightly.

However, we also changed the power ratio of the Raman beams between these two detunings, which is also expected to have an effect. It should suppress the QND impurity from

Table 7.2: Results for adaptive detection performed in real-time on the FPGA.

Detuning (GHz)	Power Ratio	Threshold Ratio	Mean Cycles	S_+ Infidelity	S_- Infidelity	Mean Infidelity
45	35:1	10^4	3.55	$3.2(5) \times 10^{-3}$	$3.0(1.7) \times 10^{-4}$	$1.7(6) \times 10^{-3}$
90	120:1	10^7	5.13	$8(3) \times 10^{-4}$	$2.6(1.5) \times 10^{-4}$	$5(3) \times 10^{-4}$
210	35:1	10^7	5.92	7/25,000	0/100,000	$\sim 1.4 \times 10^{-4}$
490	15:1	10^9	8.55	3/25,000	0/100,000	$\sim 6 \times 10^{-5}$

weak π -beam scattering, but amplify the impurity from a σ^+ -beam polarization impurity scattering out of S_+ . We therefore suspect that the dominant source of impurity changed from the weak π -beam to the σ^+ -beam polarization impurity at this higher power ratio. Likewise, the stronger relative power of the σ -beam increases the rate that S_- pumps to S_+ , so we observed a comparable error for measuring S_- .

When increasing the Raman detuning to 210 GHz, we could no longer operate with such a large power imbalance due to the same constraints discussed above, so we return to a 35:1 ratio. Infidelity for measuring both subspaces decreased, and in the case of measuring S_- we were unable to observe any disagreements after 100,000 experiments. In measuring S_+ , we observed seven error events in 25,000. The difference in improvement of the two errors over those measured at lower detuning can be understood by the underlying sources of QND impurity. Scattering from S_+ to S_- requires only a single scattering event, so all else equal we expect the associated QND impurity to scale roughly linearly with the photon scattering rate. On the other hand, given successful shelving to $|2, -2\rangle$, multiple scattering events are required to change subspaces, making S_- to S_+ subspace changes extremely rare. As a rough point estimate, we observe $\sim 1.4 \times 10^{-4}$ average infidelity across the two subspaces.

We tried going up to 1.2 THz detuning, the maximum supported by our set of seed lasers. At this large of a detuning, we ran into several problems. The Raman detuning was changed by replacing the ~ 1550 nm seed laser, however the fiber amplifier that boosts this seed laser was less efficient at the higher wavelength. Instead of 1.2 W of 313 nm light out of our doubler, we were only able to muster 600 mW in this configuration. Moreover,

${}^9\text{Be}^+$ sideband π -times on the OOPH mode were 300-400 μs , even with smaller power imbalances of around 10:1, thus making the transitions extremely narrow and hence susceptible to drifts that affected the fidelity of the QLS.

This highlights a weakness of the $\Delta m_F = 1$ variant of the protocol. Inefficient use of Raman beam power requires more of it to maintain fast Rabi rates, but that leads to increased Stark shifts. Given a certain Rabi rate, the relative strength of the Stark shift increases with detuning compared to the transition linewidth, making the resonance increasingly sensitive to beam alignment onto the ions or intensity fluctuations. Small misalignments are then sufficient to upset the QLS, particularly with how narrow our transitions were.

To reduce this instability and keep transitions sufficiently fast, we reduced the Raman detuning to the highest that was compatible with our next lowest wavelength seed laser near 1550 nm. This allowed for 490 GHz Raman detuning while maintaining our typical level of 313 nm power. However, to maintain a sideband π -time of 35-40 μs , we had to lower the Raman beam power ratio to 15:1, exacerbating the QND impurity due to scattering from the π -polarized beam. Weary of beam pointing instability impacting the measurement, we set a high threshold confidence ratio of 10^9 for this measurement. At only a minor cost in total fidelity, though an appreciable cost in time, such a threshold can compensate for instability in the QLS. We performed 25,000 experiments measuring S_+ , observing three errors. Once again we did not observe any errors in measuring S_- after 100,000 experiments. Combined this gives a point estimate of $\sim 6 \times 10^{-5}$ for the average infidelity, which is an order of magnitude improvement over the previous best indirect measurement with ions [Hume et al., 2007] and competitive with highest demonstrated measurement fidelities of any qubit system, direct or indirect.

For comparison, after removing the procedure to recover scattered population from $|2, 1\rangle$, we measured an S_+ readout error of $4(2) \times 10^{-4}$. Another useful point of comparison is the Raman scattering rate, which we measured to average $5(1) \times 10^{-4}$ per round of QLS. After the average number of rounds, this yields a 4.3×10^{-3} chance of scatter per full

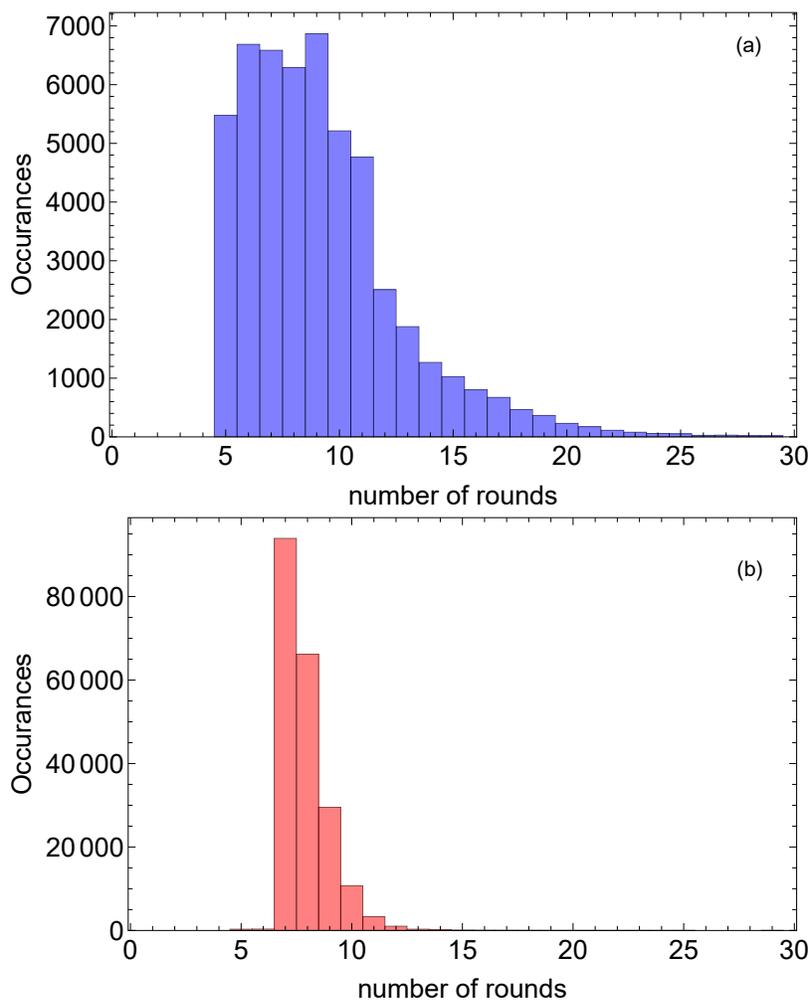


Figure 7.8: Histogram of QLS rounds at 490 GHz Raman detuning and a 10^9 threshold confidence ratio for measuring (a) S_+ and (b) S_- . The mean across both subspaces is 8.55.

detection sequence. This implies that the probability of a measurement error on S_+ is a factor of roughly 70 less than the probability of having one or more Raman scattering events on the level of the individual energy levels.

7.8 Improvements to the Indirect Measurement Protocol

The downside of these measurements is that they were quite slow compared to direct detection through state-dependent fluorescence. Being adaptive, detection time varied, but averaged around 100 ms for the largest detuning. This is prohibitively slow to be practical in a large-scale quantum information processor, but our demonstration chose to focus on maximizing fidelity. Cooling time dominated this experiment, which can be drastically reduced with more advanced techniques like EIT cooling [Roos et al., 2000, Lin et al., 2013b]. Further improvement could be made by optimizing optical pumping and Doppler cooling times, which in this experiment were chosen very conservatively. We wanted to ensure sufficient pumping/cooling, even if laser power sagged, so appreciable time savings could likely be found. After these more significant measures, a few hundred μs could be removed from each cycle on average by using adaptive techniques for the $^{25}\text{Mg}^+$ fluorescence detections [Myerson et al., 2008].

This QLS protocol applies directly to ions with nuclear spin greater than or equal to $3/2$, though it can be adapted to others. Another popular choice of trapped ion qubit is $^{171}\text{Yb}^+$ with nuclear spin $1/2$. Though lacking sufficient ground-state levels for the protocol, it can still be applied by shelving one of the qubit states to an extremely long-lived excited state, for example the $^2F_{7/2}$ levels [Edmunds et al., 2020, Ransford, 2020].

To improve fidelity, the $\Delta m_F = 0$ variant has obvious benefits. It removes the source of QND impurity from the weak π -polarized beam. Moreover, it makes more efficient use of available Raman beam power, allowing for greater detuning for a given laser power and target sideband coupling rate. This will also make it less sensitive to fluctuating Stark shifts due to beam pointing or laser intensity instabilities. Finally, as discussed in Sec. 7.3, replacing

Raman beams with near-field microwave gradients would eliminate our dominant source of QND impurity, potentially allowing for readout errors that are negligible compared to other imperfections.

Chapter 8

Conclusions and Outlook

When working simultaneously with two species of ion, the whole is more than the sum of its parts, allowing for division of labor based on the ions' relative strengths, and opening up capabilities that become available due to a synergy of the species. For example, sympathetic cooling of the collective motion through one species becomes possible, where the necessary dissipative link to the environment involves laser beams resonant only with one of the species. This protects quantum information encoded in the other species throughout cooling.

However, cooling the motion of the ions is not the only task in quantum experiments that requires dissipation to the environment through resonant laser beams. Their internal states must also be initialized and measured. While this may seem innocuous in small quantum computing experiments, mid-circuit measurements and reinitializations will require this dissipation throughout the computation as systems scale to more qubits, for example in quantum error correction. A single resonant photon scattering off of a spectator qubit is enough to destroy its encoded information, making direct mid-circuit qubit preparation and measurement a dangerous endeavor and a barrier to increasing qubit density and potentially to fault tolerance. Such errors became apparent in our demonstration of quantum gate teleportation described in Chapter 5, and have also appeared as a prominent obstacle in the experiments of others recently.

The motion of ions in close proximity is strongly coupled through Coulomb repulsion,

but their internal states are not. Outsourcing the dissipative tasks of state preparation and measurement is complicated by that disconnect, and requires actively induced coupling of the internal states through the collective motion. In Chapter 7 we proposed and demonstrated a protocol for qubit state detection that is robust to spontaneous photon scattering, a typically dominating error source in laser-based quantum logic with trapped ions. This protocol divides the $^2S_{1/2}$ ground state manifold into two subspaces, allowing us to achieve high-fidelity compatible with stringent fault-tolerance thresholds and with clear pathways for further improvement.

While the measurement is projective like all quantum measurements, the projection is into one of the two subspaces, not individual states. It is therefore less useful for indirect state preparation. In Chapter 6 we demonstrated the basic features of a “mutually-controlled multiflip” operation and proposed a protocol for using it to indirectly prepare a pure fiducial qubit state.

In the process of performing the experiments presented in this thesis, a number of areas for further improvement became apparent. While the indirect detection protocol is robust to spontaneous photon scattering, spin-motion coupling through near-field microwave gradients would remove this error source entirely and potentially increase the usefulness of this protocol for state preparation. Ultimately, such coupling could reduce laser requirements to only low power resonant beams for the secondary species and photoionization lasers for both, where all manipulations on the primary qubit are done with microwaves. Such a system would perform all qubit preparation and measurement indirectly, avoiding decoherence from stray resonant light entirely. Moreover, the absence of high power Raman beams, and accompanying drifting stray fields from charging of dielectric surfaces near the ions, will result in more stable motional frequencies and ion positions.

In the QCCD architecture, significant amounts of recoiling are necessary after separation, transportation, or mid-circuit measurements. As observed in the gate teleportation and high fidelity indirect detection experiments, cooling time can dominate the duty cycle.

One approach to reducing this time is to use more advanced cooling techniques than the slow but simple and effective sideband cooling used in this thesis. Fellow graduate student Jenny Wu is currently tackling this problem, and has produced preliminary results on electromagnetically-induced transparency cooling of $^{25}\text{Mg}^+$. The approximate tripod structure of the involved states will allow us to hopefully push the final temperature even lower by exploiting a double-dark resonance. With properly tuned parameters, one of those dark resonances can be placed on the carrier transition and the other on the blue sideband, leaving ideally just the red sideband.

Another consequence of working with mixed-species crystals is the imbalanced normal mode participation. This feature was exploited to give strong sideband coupling to $^9\text{Be}^+$ in the high-fidelity indirect detection experiment, but in general it can be a hindrance. In particular, the problem is worse for radial modes, which are typically the natural choice for spin-motion coupling through near-field microwave gradients. It also hinders cooling of modes in which the cooling species does not strongly participate. Postdoc Panyu Hou has made great progress in this area through fast motional mode-mode coupling of multi-ion strings by modulating the trapping potential, allowing for iterative cooling while swapping motional state populations between modes with strong and weak participation. Furthermore, this allows cooling of modes that the cooling lasers might not have any access to at all, for example modes that are perpendicular to a Raman beam Δk or cooling the stretch mode on $^{25}\text{Mg}^+$ in a BMB crystal that has no participation for the middle ion, potentially reducing the number of ions needed in a processor.

He is also using these tools to explore alternative approaches to quantum information processing through coherent control of the motion itself as the information carrier, rather than a bus through which ions are coupled. Another alternative approach that we have been investigating in the lab, led by postdoc Dan Cole, is generating entanglement through engineered dissipation that is robust to fluctuations in certain experimental parameters like pulse duration. The improved scheme that he is working on removes the need for sympathetic

cooling ions, building the cooling into the scheme itself, and has the potential to reach higher fidelity than initial demonstrations in a BMMB crystal [Lin et al., 2013a].

Quantum experiments are difficult because we want the qubits to not interact with each other or the environment, except for when we need them to. Being able to turn laser beams on and off has provided a high level of *temporal* isolation for these operations on small systems. However, as these systems scale, they additionally require *spatial* isolation between addressed qubits and idle qubits. This is particularly important for dissipative tasks like qubit state initialization or readout, where only a single stray resonant photon is enough to destroy a quantum superposition. Although not yet the dominant error source, if spatial barriers prove to have insufficient isolation, we expect this form of crosstalk to be a barrier to further improvement in fault tolerant quantum information processing. We therefore advocate adding an additional *spectral* layer of isolation for all dissipative qubit operations.

Bibliography

- [Ahmadi et al., 2018] Ahmadi, M., Alves, B., Baker, C., Bertsche, W., Capra, A., Carruth, C., Cesar, C., Charlton, M., Cohen, S., Collister, R., et al. (2018). Characterization of the 1s–2s transition in antihydrogen. Nature, 557(7703):71–75.
- [Akerman and Ozeri, 2018] Akerman, N. and Ozeri, R. (2018). Atomic combination clocks. New Journal of Physics, 20(12):123026.
- [Bacon, 2006] Bacon, D. (2006). Operator quantum error-correcting subsystems for self-correcting quantum memories. Physical Review A, 73(1):012340.
- [Baldwin et al., 2021] Baldwin, C., Bjork, B., Foss-Feig, M., Gaebler, J., Hayes, D., Kokish, M., Langer, C., Sedlacek, J., Stack, D., and Vittorini, G. (2021). High-fidelity light-shift gate for clock-state qubits. Physical Review A, 103(1):012603.
- [Ballance et al., 2016] Ballance, C., Harty, T., Linke, N., Sepiol, M., and Lucas, D. (2016). High-fidelity quantum logic gates using trapped-ion hyperfine qubits. Physical Review Letters, 117(6):060504.
- [Ballance, 2014] Ballance, C. J. (2014). High-Fidelity Quantum Logic in Ca. Ph. D. thesis, Oxford University.
- [Baron et al., 2014] Baron, J., Campbell, W. C., DeMille, D., Doyle, J. M., Gabrielse, G., Gurevich, Y. V., Hess, P. W., Hutzler, N. R., Kirilov, E., Kozyryev, I., et al. (2014). Order of magnitude smaller limit on the electric dipole moment of the electron. Science, 343(6168):269–272.
- [Barrett et al., 2004] Barrett, M., Chiaverini, J., Schaetz, T., Britton, J., Itano, W., Jost, J., Knill, E., Langer, C., Leibfried, D., Ozeri, R., et al. (2004). Deterministic quantum teleportation of atomic qubits. Nature, 429(6993):737–739.
- [Beloy et al., 2020] Beloy, K., Bodine, M. I., Bothwell, T., Brewer, S. M., Bromley, S. L., Chen, J.-S., Deschênes, J.-D., Diddams, S. A., Fasano, R. J., et al. (2020). Frequency ratio measurements with 18-digit accuracy using a network of optical clocks. arXiv preprint arXiv:2005.14694.

- [Benioff, 1980] Benioff, P. (1980). The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. Journal of Statistical Physics, 22(5):563–591.
- [Bennett and Brassard, 1984] Bennett, C. H. and Brassard, G. (1984). Quantum cryptography: Public key distribution and coin tossing. Proceedings of IEEE International Conference on Computers, Systems and Singal Processing, 175:8.
- [Bennett et al., 1993] Bennett, C. H., Brassard, G., Crépeau, C., Jozsa, R., Peres, A., and Wootters, W. K. (1993). Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. Physical Review Letters, 70(13):1895.
- [Bermudez et al., 2019] Bermudez, A., Xu, X., Gutiérrez, M., Benjamin, S., and Müller, M. (2019). Fault-tolerant protection of near-term trapped-ion topological qubits under realistic noise sources. Physical Review A, 100(6):062307.
- [Blakestad et al., 2009] Blakestad, R., Ospelkaus, C., VanDevender, A., Amini, J., Britton, J., Leibfried, D., and Wineland, D. J. (2009). High-fidelity transport of trapped-ion qubits through an x-junction trap array. Physical Review Letters, 102(15):153002.
- [Blakestad et al., 2011] Blakestad, R., Ospelkaus, C., VanDevender, A., Wesenberg, J., Biercuk, M., Leibfried, D., and Wineland, D. (2011). Near-ground-state transport of trapped-ion qubits through a multidimensional array. Physical Review A, 84(3):032314.
- [Blakestad, 2010] Blakestad, R. B. (2010). Transport of trapped-ion qubits within a scalable quantum processor. Ph. D. thesis, University of Colorado.
- [Bourdeauducq et al., 2017] Bourdeauducq, S., Jördens, R., Zotov, P., Britton, J., Slichter, D., Leibbrandt, D., Allcock, D., Hankin, A., Kermarrec, F., Sionneau, Y., et al. (2017). m-labs/artiq: 3.6. Zenodo, December.
- [Bowler, 2015] Bowler, R. (2015). Coherent ion transport in a multi-electrode trap array. PhD thesis, University of Colorado.
- [Bowler et al., 2012] Bowler, R., Gaebler, J., Lin, Y., Tan, T. R., Hanneke, D., Jost, J. D., Home, J., Leibfried, D., and Wineland, D. J. (2012). Coherent diabatic ion transport and separation in a multizone trap array. Physical Review Letters, 109(8):080502.
- [Bowler et al., 2013] Bowler, R., Warring, U., Britton, J. W., Sawyer, B., and Amini, J. (2013). Arbitrary waveform generator for quantum information processing with trapped ions. Review of Scientific Instruments, 84(3):033108.
- [Breit and Rabi, 1931] Breit, G. and Rabi, I. (1931). Measurement of nuclear spin. Physical Review, 38(11):2082.
- [Brewer et al., 2019] Brewer, S., Chen, J.-S., Hankin, A., Clements, E., Chou, C.-W., Wineland, D., Hume, D., and Leibbrandt, D. (2019). Al+ 27 quantum-logic clock with a systematic uncertainty below 10⁻¹⁸. Physical Review Letters, 123(3):033201.

- [Brown et al., 2011] Brown, K. R., Wilson, A. C., Colombe, Y., Ospelkaus, C., Meier, A. M., Knill, E., Leibfried, D., and Wineland, D. J. (2011). Single-qubit-gate error below 10^{-4} in a trapped ion. Physical Review A, 84(3):030303.
- [Brown and Brown, 2019] Brown, N. C. and Brown, K. R. (2019). Leakage mitigation for quantum error correction using a mixed qubit scheme. Physical Review A, 100(3):032325.
- [Bruzewicz et al., 2019] Bruzewicz, C. D., Chiaverini, J., McConnell, R., and Sage, J. M. (2019). Trapped-ion quantum computing: Progress and challenges. Applied Physics Reviews, 6(2):021314.
- [Burd et al., 2020] Burd, S., Penttinen, J.-P., Hou, P.-Y., Knaack, H., Ranta, S., Mäki, M., Kantola, E., Guina, M., Slichter, D., Leibfried, D., et al. (2020). Vecsel systems for quantum information processing with trapped beryllium ions. arXiv preprint arXiv:2003.09060.
- [Burd et al., 2019] Burd, S., Srinivas, R., Bollinger, J., Wilson, A., Wineland, D., Leibfried, D., Slichter, D., and Allcock, D. (2019). Quantum amplification of mechanical oscillator motion. Science, 364(6446):1163–1165.
- [Burd et al., 2016] Burd, S. C., Allcock, D. T., Leinonen, T., Penttinen, J.-P., Slichter, D. H., Srinivas, R., Wilson, A. C., Jördens, R., Guina, M., Leibfried, D., et al. (2016). Vecsel systems for the generation and manipulation of trapped magnesium ions. Optica, 3(12):1294–1299.
- [Burrell et al., 2010] Burrell, A., Szwer, D., Webster, S., and Lucas, D. (2010). Scalable simultaneous multiqubit readout with 99.99% single-shot fidelity. Physical Review A, 81(4):040302.
- [Burrell, 2010] Burrell, A. H. (2010). High fidelity readout of trapped ion qubits. PhD thesis, Oxford University, UK.
- [Chao and Reichardt, 2018] Chao, R. and Reichardt, B. W. (2018). Quantum error correction with only two extra qubits. Physical Review Letters, 121(5):050502.
- [Chen et al., 2017] Chen, J.-S., Brewer, S. M., Chou, C., Wineland, D., Leibbrandt, D., and Hume, D. (2017). Sympathetic ground state cooling and time-dilation shifts in an $^{27}\text{Al}^+$ optical clock. Physical Review Letters, 118(5):053002.
- [Chen et al., 2016] Chen, L., Chen, L., Jordan, S., Liu, Y.-K., Moody, D., Peralta, R., Perlmutter, R., and Smith-Tone, D. (2016). Report on post-quantum cryptography, volume 12. US Department of Commerce, National Institute of Standards and Technology.
- [Chou et al., 2020] Chou, C., Collopy, A. L., Kurz, C., Lin, Y., Harding, M. E., Plessow, P. N., Fortier, T., Diddams, S., Leibfried, D., and Leibbrandt, D. R. (2020). Frequency-comb spectroscopy on pure quantum states of a single molecular ion. Science, 367(6485):1458–1461.

- [Chou et al., 2010] Chou, C.-W., Hume, D. B., Rosenband, T., and Wineland, D. J. (2010). Optical clocks and relativity. Science, 329(5999):1630–1633.
- [Chou et al., 2017] Chou, C.-w., Kurz, C., Hume, D. B., Plessow, P. N., Leibbrandt, D. R., and Leibfried, D. (2017). Preparation and coherent manipulation of pure quantum states of a single molecular ion. Nature, 545(7653):203–207.
- [Chou et al., 2018] Chou, K. S., Blumoff, J. Z., Wang, C. S., Reinhold, P. C., Axline, C. J., Gao, Y. Y., Frunzio, L., Devoret, M., Jiang, L., and Schoelkopf, R. (2018). Deterministic teleportation of a quantum gate between two logical qubits. Nature, 561(7723):368–373.
- [Chuang and Nielsen, 1997] Chuang, I. L. and Nielsen, M. A. (1997). Prescription for experimental determination of the dynamics of a quantum black box. Journal of Modern Optics, 44(11-12):2455–2467.
- [Cirac and Zoller, 1995] Cirac, J. I. and Zoller, P. (1995). Quantum computations with cold trapped ions. Physical Review Letters, 74(20):4091.
- [Clark et al., 2021] Clark, C. R., Tinkey, H. N., Sawyer, B. C., Meier, A. M., Burkhardt, K. A., Seck, C. M., Shappert, C. M., Guise, N. D., Volin, C. E., Fallek, S. D., et al. (2021). High-fidelity bell-state preparation with 40ca+ optical qubits. arXiv preprint arXiv:2105.05828.
- [Clos et al., 2014] Clos, G., Enderlein, M., Warring, U., Schaetz, T., and Leibfried, D. (2014). Decoherence-assisted spectroscopy of a single mg+ ion. Physical Review Letters, 112(11):113003.
- [Cole et al., 2021] Cole, D. C., Wu, J. J., Erickson, S. D., Hou, P.-Y., Wilson, A. C., Leibfried, D., and Reiter, F. (2021). Dissipative preparation of w states in trapped ion systems. arXiv preprint arXiv:2103.02088.
- [Colombe et al., 2014] Colombe, Y., Slichter, D. H., Wilson, A. C., Leibfried, D., and Wineland, D. J. (2014). Single-mode optical fiber for high-power, low-loss uv transmission. Optics Express, 22(16):19783–19793.
- [Crain et al., 2019] Crain, S., Cahall, C., Vrijsen, G., Wollman, E. E., Shaw, M. D., Verma, V. B., Nam, S. W., and Kim, J. (2019). High-speed low-crosstalk detection of a 171 yb+ qubit using superconducting nanowire single photon detectors. Communications Physics, 2(1):1–6.
- [Deutsch and Jozsa, 1992] Deutsch, D. and Jozsa, R. (1992). Rapid solution of problems by quantum computation. Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences, 439(1907):553–558.
- [Devitt et al., 2013] Devitt, S. J., Munro, W. J., and Nemoto, K. (2013). Quantum error correction for beginners. Reports on Progress in Physics, 76(7):076001.

- [Dick, 1987] Dick, J. (1987). Local oscillator induced instabilities in trapped ion frequency standards. Proceedings of the 19th Precise Time and Time Interval, pages 133–147.
- [DiVincenzo, 2000] DiVincenzo, D. P. (2000). The physical implementation of quantum computation. Fortschritte der Physik: Progress of Physics, 48(9-11):771–783.
- [Edmunds et al., 2020] Edmunds, C. L., Tan, T. R., Milne, A. R., Singh, A., Biercuk, M. J., and Hempel, C. (2020). Scalable hyperfine qubit state detection via electron shelving in the $^2D_{5/2}$ and $^2F_{7/2}$ manifolds in $^{171}\text{Yb}^+$. arXiv preprint arXiv:2012.14606.
- [Eisert et al., 2000] Eisert, J., Jacobs, K., Papadopoulos, P., and Plenio, M. B. (2000). Optimal local implementation of nonlocal quantum gates. Physical Review A, 62(5):052317.
- [Ejtemaee and Haljan, 2017] Ejtemaee, S. and Haljan, P. (2017). 3d sisyphus cooling of trapped ions. Physical Review Letters, 119(4):043001.
- [Elder et al., 2020] Elder, S. S., Wang, C. S., Reinhold, P., Hann, C. T., Chou, K. S., Lester, B. J., Rosenblum, S., Frunzio, L., Jiang, L., and Schoelkopf, R. J. (2020). High-fidelity measurement of qubits encoded in multilevel superconducting circuits. Physical Review X, 10(1):011001.
- [Feng et al., 2020] Feng, L., Tan, W., De, A., Menon, A., Chu, A., Pagano, G., and Monroe, C. (2020). Efficient ground-state cooling of large trapped-ion chains with an electromagnetically-induced-transparency tripod scheme. Physical Review Letters, 125(5):053001.
- [Feynman, 1982] Feynman, R. P. (1982). Simulating physics with computers. International Journal of Theoretical Physics, 21(6/7).
- [Fiurášek, 2001] Fiurášek, J. (2001). Maximum-likelihood estimation of quantum measurement. Physical Review A, 64(2):024102.
- [Foss-Feig et al., 2020] Foss-Feig, M., Hayes, D., Dreiling, J. M., Figgatt, C., Gaebler, J. P., Moses, S. A., Pino, J. M., and Potter, A. C. (2020). Holographic quantum algorithms for simulating correlated spin systems. arXiv preprint arXiv:2005.03023.
- [Foss-Feig et al., 2021] Foss-Feig, M., Ragole, S., Potter, A., Dreiling, J., Figgatt, C., Gaebler, J., Hall, A., Moses, S., Pino, J., Spaun, B., et al. (2021). Entanglement from tensor networks on a trapped-ion qccd quantum computer. arXiv preprint arXiv:2104.11235.
- [Fowler, 2013] Fowler, A. G. (2013). Coping with qubit leakage in topological codes. Physical Review A, 88(4):042308.
- [Fowler et al., 2012] Fowler, A. G., Mariantoni, M., Martinis, J. M., and Cleland, A. N. (2012). Surface codes: Towards practical large-scale quantum computation. Physical Review A, 86(3):032324.

- [Gaebler et al., 2016] Gaebler, J. P., Tan, T. R., Lin, Y., Wan, Y., Bowler, R., Keith, A. C., Glancy, S., Coakley, K., Knill, E., Leibfried, D., et al. (2016). High-fidelity universal gate set for 9+ ion qubits. Physical Review Letters, 117(6):060505.
- [Gao et al., 2010] Gao, W.-B., Goebel, A. M., Lu, C.-Y., Dai, H.-N., Wagenknecht, C., Zhang, Q., Zhao, B., Peng, C.-Z., Chen, Z.-B., Chen, Y.-A., et al. (2010). Teleportation-based realization of an optical quantum two-qubit entangling gate. Proceedings of the National Academy of Sciences, 107(49):20869–20874.
- [Georgescu et al., 2014] Georgescu, I. M., Ashhab, S., and Nori, F. (2014). Quantum simulation. Reviews of Modern Physics, 86(1):153.
- [Gorman et al., 2014] Gorman, D. J., Schindler, P., Selvarajan, S., Daniilidis, N., and Häffner, H. (2014). Two-mode coupling in a single-ion oscillator via parametric resonance. Physical Review A, 89(6):062332.
- [Gottesman, 2010] Gottesman, D. (2010). An introduction to quantum error correction and fault-tolerant quantum computation. In Quantum information science and its contributions to mathematics, Proceedings of Symposia in Applied Mathematics, volume 68, pages 13–58.
- [Gottesman and Chuang, 1999] Gottesman, D. and Chuang, I. L. (1999). Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. Nature, 402(6760):390–393.
- [Greenberger et al., 1989] Greenberger, D. M., Horne, M. A., and Zeilinger, A. (1989). Going beyond bell’s theorem. In Bell’s theorem, quantum theory and conceptions of the universe, pages 69–72. Springer.
- [Grover, 1996] Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In Proceedings of the twenty-eighth annual ACM symposium on Theory of computing, pages 212–219.
- [Harty et al., 2014] Harty, T., Allcock, D., Ballance, C. J., Guidoni, L., Janacek, H., Linke, N., Stacey, D., and Lucas, D. (2014). High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. Physical Review Letters, 113(22):220501.
- [Harty et al., 2016] Harty, T., Sepiol, M., Allcock, D., Ballance, C., Tarlton, J., and Lucas, D. (2016). High-fidelity trapped-ion quantum logic using near-field microwaves. Physical Review Letters, 117(14):140501.
- [Hayes et al., 2020] Hayes, D., Stack, D., Bjork, B., Potter, A., Baldwin, C., and Stutz, R. (2020). Eliminating leakage errors in hyperfine qubits. Physical Review Letters, 124(17):170501.
- [Herold et al., 2016] Herold, C. D., Fallek, S. D., Merrill, J., Meier, A. M., Brown, K. R., Volin, C., and Amini, J. M. (2016). Universal control of ion qubits in a scalable microfabricated planar trap. New Journal of Physics, 18(2):023048.

- [Hite et al., 2013] Hite, D., Colombe, Y., Wilson, A. C., Allcock, D., Leibfried, D., Wineland, D., and Pappas, D. (2013). Surface science for improved ion traps. MRS Bulletin, 38(10):826–833.
- [Home and Steane, 2006] Home, J. and Steane, A. (2006). Electrode configurations for fast separation of trapped ions. Quantum Information and Computation, 6(4-5).
- [Home, 2013] Home, J. P. (2013). Quantum science and metrology with mixed-species ion chains. In Advances in Atomic, Molecular, and Optical Physics, volume 62, pages 231–277. Elsevier.
- [Hsiao et al., 2008] Hsiao, J.-T., Wang, L.-R., Sun, H.-L., Lin, S.-F., Lu, C.-L., and Huang, K.-N. (2008). Photoionization processes of the beryllium atom. Physical Review A, 78(1):013411.
- [Huang and Leibfried, 2004] Huang, P. and Leibfried, D. (2004). Achromatic catadioptric microscope objective in deep ultraviolet with long working distance. In Novel Optical Systems Design and Optimization VII, volume 5524, pages 125–133. International Society for Optics and Photonics.
- [Huang et al., 2004] Huang, Y.-F., Ren, X.-F., Zhang, Y.-S., Duan, L.-M., and Guo, G.-C. (2004). Experimental teleportation of a quantum controlled-not gate. Physical Review Letters, 93(24):240501.
- [Huelga et al., 1997] Huelga, S. F., Macchiavello, C., Pellizzari, T., Ekert, A. K., Plenio, M. B., and Cirac, J. I. (1997). Improvement of frequency standards with quantum entanglement. Physical Review Letters, 79(20):3865.
- [Hughes et al., 2020] Hughes, A., Schäfer, V., Thirumalai, K., Nadlinger, D., Woodrow, S., Lucas, D., and Ballance, C. (2020). Benchmarking a high-fidelity mixed-species entangling gate. Physical Review Letters, 125(8):080504.
- [Hume et al., 2007] Hume, D., Rosenband, T., and Wineland, D. J. (2007). High-fidelity adaptive qubit detection through repetitive quantum nondemolition measurements. Physical Review Letters, 99(12):120502.
- [Huntemann et al., 2012] Huntemann, N., Lipphardt, B., Okhapkin, M., Tamm, C., Peik, E., Taichenachev, A., and Yudin, V. (2012). Generalized ramsey excitation scheme with suppressed light shift. Physical Review Letters, 109(21):213002.
- [Itano and Wineland, 1982] Itano, W. M. and Wineland, D. (1982). Laser cooling of ions stored in harmonic and penning traps. Physical Review A, 25(1):35.
- [Jefferts et al., 1995] Jefferts, S. R., Monroe, C., Bell, E., and Wineland, D. J. (1995). Coaxial-resonator-driven rf (paul) trap for strong confinement. Physical Review A, 51(4):3112.

- [Jordan, 2020] Jordan, S. (2020). Quantum algorithm zoo. <https://quantumalgorithmzoo.org/>. Accessed: 2020-04-22.
- [Keith et al., 2018] Keith, A. C., Baldwin, C. H., Glancy, S., and E, K. (2018). Joint quantum-state and measurement tomography with incomplete measurements. Physical Review A, 98(4):042318.
- [Keller et al., 2020] Keller, J., Hou, P.-Y., McCormick, K. C., Cole, D. C., Erickson, S. D., Wu, J. J., Wilson, A. C., and Leibfried, D. (2020). Quantum harmonic oscillator spectrum analyzers. arXiv preprint arXiv:2010.10438.
- [Kielpinski et al., 2000] Kielpinski, D., King, B., Myatt, C., Sackett, C., Turchette, Q., Itano, W. M., Monroe, C., Wineland, D. J., and Zurek, W. (2000). Sympathetic cooling of trapped ions for quantum logic. Physical Review A, 61(3):032310.
- [Kielpinski et al., 2002] Kielpinski, D., Monroe, C., and Wineland, D. J. (2002). Architecture for a large-scale ion-trap quantum computer. Nature, 417(6890):709–711.
- [Kienzler et al., 2020] Kienzler, D., Wan, Y., Erickson, S., Wu, J., Wilson, A., Wineland, D., and Leibfried, D. (2020). Quantum logic spectroscopy with ions in thermal motion. Physical Review X, 10(2):021012.
- [Kim, 2017] Kim, I. H. (2017). Holographic quantum simulation. arXiv preprint arXiv:1702.02093.
- [Ladd et al., 2010] Ladd, T. D., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C., and O’Brien, J. L. (2010). Quantum computers. Nature, 464(7285):45–53.
- [Langer et al., 2005] Langer, C., Ozeri, R., Jost, J. D., Chiaverini, J., DeMarco, B., Ben-Kish, A., Blakestad, R., Britton, J., Hume, D., Itano, W. M., et al. (2005). Long-lived qubit memory using atomic ions. Physical Review Letters, 95(6):060502.
- [Langer, 2006] Langer, C. E. (2006). High fidelity quantum information processing with trapped ions. PhD thesis, University of Colorado at Boulder.
- [Lee et al., 2005] Lee, P. J., Brickman, K.-A., Deslauriers, L., Haljan, P. C., Duan, L.-M., and Monroe, C. (2005). Phase control of trapped ion quantum gates. Journal of Optics B: Quantum and Semiclassical Optics, 7(10):S371.
- [Leibfried, 2017] Leibfried, D. (2017). Playing tricks to ions. Applied Physics B, 123(1):10.
- [Leibfried et al., 2004a] Leibfried, D., Barrett, M. D., Kish, A. B., Britton, J., Chiaverini, J., DeMarco, B., Itano, W. M., Jelenković, B., Jost, J. D., Langer, C., et al. (2004a). Building blocks for a scalable quantum information processor based on trapped ions. In Laser Spectroscopy, pages 295–303. World Scientific.
- [Leibfried et al., 2004b] Leibfried, D., Barrett, M. D., Schaetz, T., Britton, J., Chiaverini, J., Itano, W. M., Jost, J. D., Langer, C., and Wineland, D. J. (2004b). Toward heisenberg-limited spectroscopy with multiparticle entangled states. Science, 304(5676):1476–1478.

- [Leibfried et al., 2003a] Leibfried, D., Blatt, R., Monroe, C., and Wineland, D. (2003a). Quantum dynamics of single trapped ions. Reviews of Modern Physics, 75(1):281.
- [Leibfried et al., 2003b] Leibfried, D., DeMarco, B., Meyer, V., Lucas, D., Barrett, M., Britton, J., Itano, W. M., Jelenković, B., Langer, C., Rosenband, T., et al. (2003b). Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. Nature, 422(6930):412–415.
- [Leibrandt and Heidecker, 2015] Leibrandt, D. R. and Heidecker, J. (2015). An open source digital servo for atomic, molecular, and optical physics experiments. Review of Scientific Instruments, 86(12):123115.
- [Levitt, 1982] Levitt, M. H. (1982). Symmetrical composite pulse sequences for nmr population inversion. i. compensation of radiofrequency field inhomogeneity. Journal of Magnetic Resonance (1969), 48(2):234–264.
- [Lin et al., 2013a] Lin, Y., Gaebler, J., Reiter, F., Tan, T. R., Bowler, R., Sørensen, A., Leibfried, D., and Wineland, D. J. (2013a). Dissipative production of a maximally entangled steady state of two quantum bits. Nature, 504(7480):415–418.
- [Lin et al., 2016] Lin, Y., Gaebler, J. P., Reiter, F., Tan, T. R., Bowler, R., Wan, Y., Keith, A., Knill, E., Glancy, S., Coakley, K., et al. (2016). Preparation of entangled states through hilbert space engineering. Physical Review Letters, 117(14):140502.
- [Lin et al., 2013b] Lin, Y., Gaebler, J. P., Tan, T. R., Bowler, R., Jost, J. D., Leibfried, D., and Wineland, D. J. (2013b). Sympathetic electromagnetically-induced-transparency laser cooling of motional modes in an ion chain. Physical Review Letters, 110(15):153002.
- [Ludlow et al., 2015] Ludlow, A. D., Boyd, M. M., Ye, J., Peik, E., and Schmidt, P. O. (2015). Optical atomic clocks. Reviews of Modern Physics, 87(2):637.
- [Marangos, 1998] Marangos, J. P. (1998). Electromagnetically induced transparency. Journal of Modern Optics, 45(3):471–503.
- [Maslov, 2017] Maslov, D. (2017). Basic circuit compilation techniques for an ion-trap quantum machine. New Journal of Physics, 19(2):023035.
- [McCormick et al., 2019] McCormick, K. C., Keller, J., Burd, S. C., Wineland, D. J., Wilson, A. C., and Leibfried, D. (2019). Quantum-enhanced sensing of a single-ion mechanical oscillator. Nature, 572(7767):86–90.
- [McEwen et al., 2021] McEwen, M., Kafri, D., Chen, Z., Atalaya, J., Satzinger, K., Quintana, C., Klimov, P. V., Sank, D., Gidney, C., Fowler, A., et al. (2021). Removing leakage-induced correlated errors in superconducting quantum error correction. Nature Communications, 12(1):1–7.
- [McGrew et al., 2018] McGrew, W., Zhang, X., Fasano, R., Schäffer, S., Beloy, K., Nicolodi, D., Brown, R., Hinkley, N., Milani, G., Schioppo, M., et al. (2018). Atomic clock performance enabling geodesy below the centimetre level. Nature, 564(7734):87–90.

- [Merkel et al., 2019] Merkel, B., Thirumalai, K., Tarlton, J., Schäfer, V., Ballance, C., Harty, T., and Lucas, D. (2019). Magnetic field stabilization system for atomic physics experiments. Review of Scientific Instruments, 90(4):044702.
- [Micke et al., 2020] Micke, P., Leopold, T., King, S., Benkler, E., Spieß, L., Schmoeger, L., Schwarz, M., López-Urrutia, J. C., and Schmidt, P. (2020). Coherent laser spectroscopy of highly charged ions using quantum logic. Nature, 578(7793):60–65.
- [Moehring et al., 2007] Moehring, D. L., Maunz, P., Olmschenk, S., Younge, K. C., Matsukevich, D. N., Duan, L.-M., and Monroe, C. (2007). Entanglement of single-atom quantum bits at a distance. Nature, 449(7158):68–71.
- [Monroe and Kim, 2013] Monroe, C. and Kim, J. (2013). Scaling the ion trap quantum processor. Science, 339(6124):1164–1169.
- [Monroe et al., 1995a] Monroe, C., Meekhof, D., King, B., Itano, W. M., and Wineland, D. J. (1995a). Demonstration of a fundamental quantum logic gate. Physical Review Letters, 75(25):4714.
- [Monroe et al., 1995b] Monroe, C., Meekhof, D., King, B., Jefferts, S. R., Itano, W. M., Wineland, D. J., and Gould, P. (1995b). Resolved-sideband raman cooling of a bound atom to the 3d zero-point energy. Physical Review Letters, 75(22):4011.
- [Monroe et al., 2014] Monroe, C., Raussendorf, R., Ruthven, A., Brown, K., Maunz, P., Duan, L.-M., and Kim, J. (2014). Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. Physical Review A, 89(2):022317.
- [Morigi et al., 2000] Morigi, G., Eschner, J., and Keitel, C. H. (2000). Ground state laser cooling using electromagnetically induced transparency. Physical Review Letters, 85(21):4458.
- [Morinaga et al., 1989] Morinaga, A., Riehle, F., Ishikawa, J., and Helmcke, J. (1989). A ca optical frequency standard: Frequency stabilization by means of nonlinear ramsey resonances. Applied Physics B, 48(2):165–171.
- [Mur-Petit et al., 2012] Mur-Petit, J., García-Ripoll, J. J., Pérez-Ríos, J., Campos-Martínez, J., Hernández, M. I., and Willitsch, S. (2012). Temperature-independent quantum logic for molecular spectroscopy. Physical Review A, 85(2):022308.
- [Myerson et al., 2008] Myerson, A., Szwer, D., Webster, S., Allcock, D., Curtis, M., Imreh, G., Sherman, J., Stacey, D., Steane, A., and Lucas, D. (2008). High-fidelity readout of trapped-ion qubits. Physical Review Letters, 100(20):200502.
- [Negnevitsky et al., 2018] Negnevitsky, V., Marinelli, M., Mehta, K. K., Lo, H.-Y., Flühmann, C., and Home, J. P. (2018). Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register. Nature, 563(7732):527–531.

- [Nielsen, 2002] Nielsen, M. A. (2002). A simple formula for the average gate fidelity of a quantum dynamical operation. Physics Letters A, 303(4):249–252.
- [Nielsen and Chuang, 2002] Nielsen, M. A. and Chuang, I. (2002). Quantum computation and quantum information.
- [Ospelkaus et al., 2008] Ospelkaus, C., Langer, C. E., Amini, J. M., Brown, K. R., Leibfried, D., and Wineland, D. J. (2008). Trapped-ion quantum logic gates based on oscillating magnetic fields. Physical Review Letters, 101(9):090502.
- [Ozeri et al., 2007] Ozeri, R., Itano, W. M., Blakestad, R., Britton, J., Chiaverini, J., Jost, J. D., Langer, C., Leibfried, D., Reichle, R., Seidelin, S., et al. (2007). Errors in trapped-ion quantum gates due to spontaneous photon scattering. Physical Review A, 75(4):042329.
- [Parrado-Rodríguez et al., 2020] Parrado-Rodríguez, P., Ryan-Anderson, C., Bermudez, A., and Müller, M. (2020). Crosstalk suppression for fault-tolerant quantum error correction with trapped ions. arXiv preprint arXiv:2012.11366.
- [Paul, 1990] Paul, W. (1990). Electromagnetic traps for charged and neutral particles. Reviews of Modern Physics, 62(3):531.
- [Pino et al., 2021] Pino, J., Dreiling, J., Figgatt, C., Gaebler, J., Moses, S., Allman, M., Baldwin, C., Foss-Feig, M., Hayes, D., Mayer, K., et al. (2021). Demonstration of the trapped-ion quantum ccd computer architecture. Nature, 592(7853):209–213.
- [Poulsen et al., 1975] Poulsen, O., Andersen, T., and Skouboe, N. (1975). Fast-beam, zero-field level-crossing measurements of radiative lifetimes, fine and hyperfine structures in excited states of ionic and neutral beryllium. Journal of Physics B: Atomic and Molecular Physics, 8(9):1393.
- [Preskill, 1998] Preskill, J. (1998). Lecture notes for physics 229: Quantum information and computation. California Institute of Technology, 16.
- [Qiao et al., 2021] Qiao, M., Wang, Y., Cai, Z., Du, B., Wang, P., Luan, C., Chen, W., Noh, H.-R., and Kim, K. (2021). Double-electromagnetically-induced-transparency ground-state cooling of stationary two-dimensional ion crystals. Physical Review Letters, 126(2):023604.
- [Rafiq et al., 2007] Rafiq, M., Hussain, S., Saleem, M., Kalyar, M., and Baig, M. (2007). Measurement of photoionization cross section from the 3s3p 1p1 excited state of magnesium. Journal of Physics B: Atomic, Molecular and Optical Physics, 40(12):2291.
- [Ransford, 2020] Ransford, A. M. (2020). Old Dog, New Trick: High Fidelity, Background-free State Detection of an Ytterbium Ion Qubit. PhD thesis, UCLA.
- [Rattew et al., 2020] Rattew, A. G., Sun, Y., Minssen, P., and Pistoia, M. (2020). Quantum simulation of galton machines using mid-circuit measurement and reuse. arXiv preprint arXiv:2009.06601.

- [Reiter et al., 2020] Reiter, F., Nguyen, T. L., Home, J. P., and Yelin, S. F. (2020). Cooperative breakdown of the oscillator blockade in the dicke model. Physical Review Letters, 125(23):233602.
- [Riebe et al., 2004] Riebe, M., Häffner, H., Roos, C., Hänsel, W., Benhelm, J., Lancaster, G., Körber, T., Becher, C., Schmidt-Kaler, F., James, D., et al. (2004). Deterministic quantum teleportation with atoms. Nature, 429(6993):734–737.
- [Roos et al., 2000] Roos, C., Leibfried, D., Mundt, A., Schmidt-Kaler, F., Eschner, J., and Blatt, R. (2000). Experimental demonstration of ground state laser cooling with electromagnetically induced transparency. Physical Review Letters, 85(26):5547.
- [Roos et al., 2006] Roos, C. F., Chwalla, M., Kim, K., Riebe, M., and Blatt, R. (2006). ‘designer atoms’ for quantum metrology. Nature, 443(7109):316–319.
- [Rosenband et al., 2008] Rosenband, T., Hume, D., Schmidt, P., Chou, C.-W., Brusch, A., Lorini, L., Oskay, W., Drullinger, R. E., Fortier, T. M., Stalnaker, J. E., et al. (2008). Frequency ratio of al^+ and hg^+ single-ion optical clocks; metrology at the 17th decimal place. Science, 319(5871):1808–1812.
- [Rosenband et al., 2007] Rosenband, T., Schmidt, P. O., Hume, D. B., Itano, W. M., Fortier, T. M., Stalnaker, J. E., Kim, K., Diddams, S. A., Koelemeij, J., Bergquist, J., et al. (2007). Observation of the $s\ 0\ 1\ -j\ p\ 0\ 3$ clock transition in al^+ 27. Physical Review Letters, 98(22):220801.
- [Sackett et al., 2000] Sackett, C. A., Kielpinski, D., King, B. E., Langer, C., Meyer, V., Myatt, C. J., Rowe, M., Turchette, Q., Itano, W. M., Wineland, D. J., et al. (2000). Experimental entanglement of four particles. Nature, 404(6775):256–259.
- [Sanner et al., 2018] Sanner, C., Huntemann, N., Lange, R., Tamm, C., and Peik, E. (2018). Autobalanced ramsey spectroscopy. Physical Review Letters, 120(5):053602.
- [Santarelli et al., 1998] Santarelli, G., Audoin, C., Makdissi, A., Laurent, P., Dick, G. J., and Clairon, A. (1998). Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator. IEEE Transactions On Ultrasonics, Ferroelectrics, and Frequency Control, 45(4):887–894.
- [Sarovar et al., 2020] Sarovar, M., Proctor, T., Rudinger, K., Young, K., Nielsen, E., and Blume-Kohout, R. (2020). Detecting crosstalk errors in quantum information processors. Quantum, 4:321.
- [Scarani et al., 2009] Scarani, V., Bechmann-Pasquinucci, H., Cerf, N. J., Dušek, M., Lütkenhaus, N., and Peev, M. (2009). The security of practical quantum key distribution. Reviews of Modern Physics, 81(3):1301.
- [Schmidt et al., 2005] Schmidt, P. O., Rosenband, T., Langer, C., Itano, W. M., Bergquist, J. C., and Wineland, D. J. (2005). Spectroscopy using quantum logic. Science, 309(5735):749–752.

- [Shaniv et al., 2018] Shaniv, R., Manovitz, T., Shapira, Y., Akerman, N., and Ozeri, R. (2018). Toward heisenberg-limited rabi spectroscopy. Physical Review Letters, 120(24):243603.
- [Shor, 1994] Shor, P. W. (1994). Algorithms for quantum computation: discrete logarithms and factoring. In Proceedings 35th annual symposium on foundations of computer science, pages 124–134. Ieee.
- [Shor, 1995] Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. Physical Review A, 52(4):R2493.
- [Shor, 1996] Shor, P. W. (1996). Fault-tolerant quantum computation. In Proceedings of 37th Conference on Foundations of Computer Science, pages 56–65. IEEE.
- [Sørensen and Mølmer, 1999] Sørensen, A. and Mølmer, K. (1999). Quantum computation with ions in thermal motion. Physical Review Letters, 82(9):1971.
- [Sørensen and Mølmer, 2000] Sørensen, A. and Mølmer, K. (2000). Entanglement and quantum computation with ions in thermal motion. Physical Review A, 62(2):022311.
- [Srinivas et al., 2021] Srinivas, R., Burd, S., Knaack, H., Sutherland, R., Kwiatkowski, A., Glancy, S., Knill, E., Wineland, D., Leibfried, D., Wilson, A., et al. (2021). High-fidelity laser-free universal control of two trapped ion qubits. arXiv preprint arXiv:2102.12533.
- [Stricker et al., 2020] Stricker, R., Vodola, D., Erhard, A., Postler, L., Meth, M., Ringbauer, M., Schindler, P., Monz, T., Müller, M., and Blatt, R. (2020). Experimental deterministic correction of qubit loss. Nature, 585(7824):207–210.
- [Suchara et al., 2015] Suchara, M., Cross, A. W., and Gambetta, J. M. (2015). Leakage suppression in the toric code. In 2015 IEEE International Symposium on Information Theory (ISIT), pages 1119–1123. IEEE.
- [Sutherland, 2019] Sutherland, R. (2019). Analog quantum simulation of superradiance and subradiance in trapped ions. Physical Review A, 100(6):061405.
- [Sutherland et al., 2019] Sutherland, R., Srinivas, R., Burd, S. C., Leibfried, D., Wilson, A. C., Wineland, D. J., Allcock, D., Slichter, D. H., and Libby, S. (2019). Versatile laser-free trapped-ion entangling gates. New Journal of Physics, 21(3):033033.
- [Tan, 2016] Tan, T. R. (2016). High-Fidelity Entangling Gates with Trapped-Ions. PhD thesis, University of Colorado.
- [Tan et al., 2015] Tan, T. R., Gaebler, J. P., Lin, Y., Wan, Y., Bowler, R., Leibfried, D., and Wineland, D. J. (2015). Multi-element logic gates for trapped-ion qubits. Nature, 528(7582):380–383.
- [Tan et al., 2017] Tan, T. R., Wan, Y., Erickson, S., Bierhorst, P., Kienzler, D., Glancy, S., Knill, E., Leibfried, D., and Wineland, D. (2017). Chained bell inequality experiment with high-efficiency measurements. Physical Review Letters, 118(13):130403.

- [Todaro, 2020] Todaro, S. L. (2020). Scalable State Detection and Fast Transport of Trapped-Ion Qubits for Quantum Computing. PhD thesis, University of Colorado at Boulder.
- [Wan et al., 2020] Wan, Y., Jördens, R., Erickson, S. D., Wu, J. J., Bowler, R., Tan, T. R., Hou, P.-Y., Wineland, D. J., Wilson, A. C., and Leibfried, D. (2020). Ion transport and reordering in a 2d trap array. Advanced Quantum Technologies, 3(11):2000028.
- [Wan et al., 2019] Wan, Y., Kienzler, D., Erickson, S. D., Mayer, K. H., Tan, T. R., Wu, J. J., Vasconcelos, H. M., Glancy, S., Knill, E., Wineland, D. J., et al. (2019). Quantum gate teleportation between separated qubits in a trapped-ion processor. Science, 364(6443):875–878.
- [Wang et al., 2021] Wang, P., Luan, C.-Y., Qiao, M., Um, M., Zhang, J., Wang, Y., Yuan, X., Gu, M., Zhang, J., and Kim, K. (2021). Single ion qubit with estimated coherence time exceeding one hour. Nature Communications, 12(1):1–8.
- [Wehlitz et al., 2003] Wehlitz, R., Lukić, D., and Bluett, J. (2003). Resonance parameters of autoionizing $be\ 2\ p\ n\ l$ states. Physical Review A, 68(5):052708.
- [Weidt et al., 2016] Weidt, S., Randall, J., Webster, S., Lake, K., Webb, A., Cohen, I., Navickas, T., Lekitsch, B., Retzker, A., and Hensinger, W. (2016). Trapped-ion quantum logic with global radiation fields. Physical Review Letters, 117(22):220501.
- [Wilson et al., 2011] Wilson, A. C., Ospelkaus, C., VanDevender, A., Mlynek, J. A., Brown, K. R., Leibfried, D., and Wineland, D. J. (2011). A 750-mw, continuous-wave, solid-state laser source at 313 nm for cooling and manipulating trapped $9\ be+$ ions. Applied Physics B, 105(4):741–748.
- [Wineland et al., 2002] Wineland, D., Bergquist, J., Bollinger, J., Drullinger, R., and Itano, W. (2002). Quantum computers and atomic clocks. In Frequency Standards And Metrology, pages 361–368. World Scientific.
- [Wineland et al., 1983] Wineland, D., Bollinger, J., and Itano, W. M. (1983). Laser-fluorescence mass spectroscopy. Physical Review Letters, 50(9):628.
- [Wineland et al., 1992] Wineland, D. J., Dalibard, J., and Cohen-Tannoudji, C. (1992). Sisyphus cooling of a bound atom. Journal of the Optical Society of America B, 9(1):32–42.
- [Wineland and Itano, 1979] Wineland, D. J. and Itano, W. M. (1979). Laser cooling of atoms. Physical Review A, 20(4):1521.
- [Wineland et al., 1998] Wineland, D. J., Monroe, C., Itano, W. M., Leibfried, D., King, B. E., and Meekhof, D. M. (1998). Experimental issues in coherent quantum-state manipulation of trapped atomic ions. Journal of Research of the National Institute of Standards and Technology, 103(3):259.

- [Wolf et al., 2016] Wolf, F., Wan, Y., Heip, J. C., Gebert, F., Shi, C., and Schmidt, P. O. (2016). Non-destructive state detection for quantum logic spectroscopy of molecular ions. Nature, 530(7591):457–460.
- [Wootters and Zurek, 1982] Wootters, W. K. and Zurek, W. H. (1982). A single quantum cannot be cloned. Nature, 299(5886):802–803.
- [Xu et al., 2017] Xu, Z., Deng, K., Che, H., Yuan, W., Zhang, J., and Lu, Z. (2017). Precision measurement of the $\text{mg}+ 25$ ground-state hyperfine constant. Physical Review A, 96(5):052507.
- [Yudin et al., 2010] Yudin, V., Taichenachev, A., Oates, C., Barber, Z., Lemke, N. D., Ludlow, A. D., Sterr, U., Lisdat, C., and Riehle, F. (2010). Hyper-ramsey spectroscopy of optical clock transitions. Physical Review A, 82(1):011804.
- [Zarantonello et al., 2019] Zarantonello, G., Hahn, H., Morgner, J., Schulte, M., Bautista-Salvador, A., Werner, R., Hammerer, K., and Ospelkaus, C. (2019). Robust and resource-efficient microwave near-field entangling $\text{be}+ 9$ gate. Physical Review Letters, 123(26):260503.
- [Zhou and Lin, 1995] Zhou, B. and Lin, C. (1995). Photoionization of the beryllium atom. Physical Review A, 51(2):1286.