

**Improved State Detection and Transport of Trapped Ion
Qubits for Scalable Quantum Computing**

by

Susanna L. Todaro

M.S. Physics, University of Colorado at Boulder

B.S. Physics, Harvey Mudd College

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

2020

This thesis entitled:
Improved State Detection and Transport of Trapped Ion Qubits for Scalable Quantum Computing
written by Susanna L. Todaro
has been approved for the Department of Physics

Andrew C. Wilson

Daniel H. Slichter

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Todaro, Susanna L. (Ph.D., Physics)

Improved State Detection and Transport of Trapped Ion Qubits for Scalable Quantum Computing

Thesis directed by Andrew C. Wilson

Despite recent progress in trapped-ion quantum computation, scaling to large numbers of qubits remain a challenge. One proposal for extending beyond tens of ions in a single string is the ‘quantum CCD’ (QCCD) architecture [1] [2]. In this approach, ion qubits are coherently transported between trapping zones dedicated to memory, readout and logical operations. In this thesis, I describe work on improving two aspects of the QCCD: trap-integrated qubit state readout and cold, fast ion transport. Trap-integrated photon detectors promise reduced complexity and improved scalability for ion qubit readout relative to external bulk collection optics, and with comparable light collection efficiency. I report the first high-fidelity readout of a trapped ion qubit in a surface electrode trap using an integrated superconducting nanowire single-photon detector (SNSPD) [3]. To maximize the clock speed of QCCD processors, fast ion transport with low motional excitation is needed. In most prior QCCD experiments, low motional excitation was achieved by slow (adiabatic) ion transport between trap zones, and ion transit times were much longer than typical laser-driven gate interactions. Faster-than-adiabatic transport between neighboring trap zones has been previously demonstrated in relatively large three-dimensional traps [4][5], but low motional excitation was only achieved with particular choices of the transport duration. I report multi-zone faster-than-adiabatic transport in a surface electrode trap with reduced dependence of the final motional excitation on the duration of the transport. Using this scheme, a single ${}^9\text{Be}^+$ ion was transported 88 μm in 25 μs with 0.03(2) quanta of heating due to transport, the lowest excitation reported to date for faster-than-adiabatic transport.

For my family, both biological and chosen.

Acknowledgements

It's a funny thing to write a dissertation and put only my name on it, when science is such a deeply collaborative process and so many people have contributed so much to this work. Over the course of my graduate career, I've been blessed (and occasionally cursed) with a large number of mentors, advisors, and supervisors. Thanks to Daniel Slichter, Andrew Wilson, Didi Leibfried and Dave Wineland for supervising the projects discussed in this thesis, sharing invaluable scientific insight, and securing funding. They didn't always agree with each other on everything (or even most things), but I learned enormously from every one of them. I'd like to especially acknowledge Daniel and Andrew, who have been incredible mentors through my graduate career and have assisted in my development not only as a physicist but as a person. An advisor that can guide you through a panic attack is a rare and valuable thing, and I've been so lucky to have worked with both of them.

I also need to acknowledge everyone down in the trenches with me. Robert Jördens began cryostat development many, many years ago and was the first post-doc I worked with on this project. Much of my foundational experience as an experimental physicist was gained working with him. Felix Knollman has only recently joined this experiment but has made great progress with investigation of integrated optics with UV light. It's tough to leave anyone in control of the apparatus that I've spent my days and nights dwelling on for years, but since I have to, I'm glad to leave it in his capable hands. David Allcock, Shaun Burd, Hannah Knaack, Elena Jordan, Katie McCormick, Justin Neidermeyer, and Raghu Srinivas also all made valuable contributions to this work, from setting up the original ${}^9\text{Be}^+$ lasers to sharing frustrations over ARTIQ code to bringing me snacks on my lowest days.

A few more people who have contributed invaluablely to this work: Varun Verma, Rich Mirin, and Sae Woo Nam from the Quantum Nanoelectronics group, whose groundbreaking work on single photon sensing made many of the results in this work possible. Varun specifically, with Daniel Slichter, fabricated the “detector trap” I used for the results in Chapter 6.

I’m also indebted to the administrative staff (both at NIST and at CU), who have saved me from an eternity’s worth of paperwork and enabled me to get paid well enough to eat through all of grad school. Also to Alejandra Collopy and Andrew Spott, who offered me a place to live with their two excellent cats my last few months of grad school.

I also need to thank the people who made me who I am. My father patiently taught me to love math in a school system that failed me; my mother, the first Dr. Todaro, bought me my first lab notebook and helped me set up my very first experiments in her lab when I was nine. I’d be no one without them. Finally, finishing this dissertation would have been impossible without the love and support of not just my biological family but my friends who have become family, who stuck with me through this colossally stressful process. To the people who sat on the floor with me while I sobbed, who answered the phone at one in the morning to listen to my rambling fears, who let me visit just to pet your dogs on my low days, who picked me up again and again every time I fell down: thank you. I have no idea where I would be without you, but I’d definitely never be able to call myself ‘doctor.’

Contents

Chapter

1	Introduction	1
1.1	Scalable quantum computing with trapped ions	2
1.1.1	The quantum CCD	3
1.1.2	Alternative scalable schemes	6
1.2	This work	8
1.2.1	Development of a cryogenic apparatus for trapped ion experiments	9
1.2.2	Scalable high-fidelity qubit state detection	10
1.2.3	Fast ion transport	12
2	Ion Traps	15
2.1	RF Paul Traps	15
2.2	Ion motion in Paul traps	18
2.2.1	Secular motion: the harmonic oscillator	19
2.2.2	Micromotion	21
2.2.3	Normal Modes	21
2.3	Surface electrode traps	23
2.3.1	Simulations of surface electrode traps	25
2.3.2	Shims and curvatures	29
2.4	Manipulation of ion crystals	29

2.4.1	Ion transport	30
2.4.2	Splitting and merging of ion crystals	30
2.4.3	Reordering	30
3	The ${}^9\text{Be}^+$ Qubit	32
3.1	Beryllium photoionization	32
3.2	Level Structure	33
3.3	State preparation and control	37
3.3.1	Doppler cooling	37
3.3.2	Optical pumping	37
3.3.3	Microwave transitions	39
3.3.4	Stimulated Raman transitions	40
3.4	Ion motional control	42
3.4.1	Sideband interactions	42
3.4.2	Raman sideband cooling	46
4	Apparatus	49
4.1	The cryostat	49
4.1.1	Vibration	50
4.1.2	Heat load estimation and mitigation	57
4.1.3	Temperature measurement	57
4.1.4	Vacuum hygiene	58
4.2	Ion trapping system	60
4.2.1	Traps	60
4.2.2	Experiment pillbox	64
4.2.3	Trap mounting	64
4.2.4	Beryllium ovens and ablation target	68
4.2.5	Filters	70

4.2.6	Generation of trapping voltages	73
4.2.7	Magnetic field coils	75
4.2.8	Laser sources	76
4.2.9	Imaging system	82
4.2.10	Experimental control system	85
4.3	Trapping ions in the new apparatus	85
4.3.1	Ion lifetime	86
4.3.2	Heating rates	86
5	Methods	89
5.1	Qubit state detection	89
5.1.1	Thresholding	91
5.1.2	Depumping and repumping	93
5.1.3	Shelving	96
5.1.4	High fidelity state detection using adaptive Bayesian statistics	96
5.2	Ion motional state measurement	98
5.3	Calibrations	100
5.3.1	Loading Ions	101
5.3.2	Stray fields	101
5.3.3	Polarization and magnetic fields	103
5.3.4	Microwave pulse times and frequencies	104
5.3.5	Secular frequency measurement	104
5.3.6	Sideband cooling calibration	105
6	Trap-integrated qubit state detection	106
6.1	Maximizing detection efficiency	106
6.1.1	Single photon detectors	108
6.1.2	Superconducting nanowire single photon detectors	108

6.2	Trap-integrated detection	111
6.2.1	Previous work	112
6.3	Hardware for SNSPD readout	113
6.3.1	Readout circuit	114
6.4	Trap-integrated SNSPD performance	116
6.4.1	RF sensitivity	116
6.4.2	Detector efficiency measurement	119
6.4.3	Spatial sensitivity	120
6.5	High fidelity state detection using a trap-integrated SNSPD	122
6.5.1	Post-selection	122
6.5.2	Calibrations for high fidelity detection measurements	123
6.5.3	SNSPD readout fidelity	125
6.6	Ion heating due to the SNSPD	126
6.7	Outlook	128
7	Cold, fast ion transport	130
7.1	Previous work	131
7.2	Theory of fast transport	131
7.2.1	Classical treatment	131
7.2.2	Quantum Treatment	133
7.2.3	Transport of ion chains	135
7.2.4	Finding transport profiles	136
7.3	Waveform Generation	137
7.3.1	Fast DACs	137
7.3.2	Finding potentials	138
7.3.3	Predistortion	139
7.4	Results	142

7.5	Summary	145
8	Conclusions and future directions	147
8.1	Trap-integrated detection	147
8.1.1	Next-generation trap-integrated SNSPDs	147
8.1.2	Integration into a larger device	148
8.2	Fast manipulation of ion chains	149
8.2.1	Separation	149
	Bibliography	150
	Appendix	
A	Heat load estimation	161
B	Trapping Potentials	166
C	Code	168
D	Derivation of transport as a displacement operator	169

Tables

Table

3.1	Lande g -factors for ${}^9\text{Be}^+$	34
4.1	Cryostat heat load	57
4.2	Thermal expansion coefficients	67
4.3	Vapor pressures and melting points of common elements used for trapped-ion quantum computing	69
4.4	RC filter components	73
4.5	Exact laser wavelengths	81
B.1	Trapping voltages	167

Figures

Figure

1.1	A cartoon illustration of the quantum CCD	3
1.2	Examples of new trap technologies	6
1.3	Double-well ion trap	8
1.4	Sample quantum CCD experiment timeline	13
2.1	Paul trap schematic	16
2.2	Electric field lines in a quadrupole trap	16
2.3	The displacement operator	20
2.4	Example surface electrode traps	24
2.5	Reordering an ion chain using a junction	31
3.1	Neutral beryllium level structure	33
3.2	${}^9\text{Be}^+$ energy levels	35
3.3	Raman carrier interactions	41
3.4	Stimulated Raman transitions	44
3.5	A red sideband on an ion in the ground state of motion	47
4.1	Low-vibration cryocooler configuration	51
4.2	Experimental configuration for vibration measurement	52
4.3	Cryostat vibration in the time domain	53
4.4	Cryostat vibration in the frequency domain	54

4.5	Temperature measurement	59
4.6	Trap and cold head temperatures while the cryostat cools from room temperature to 4 K	59
4.7	Point trap electrodes	61
4.8	Detector trap electrodes	61
4.9	Integrated SNSPD detail	62
4.10	Microwave coupling circuit	63
4.11	Transport trap electrodes	64
4.12	The pillbox	65
4.13	Photo of an ion trap that cracked under thermal contraction stresses	66
4.14	The detector trap puck and filterboard	68
4.15	Resistive beryllium ovens	70
4.16	View over the trap boards of the ovens and the ablation target	71
4.17	Low-pass filter circuits	72
4.18	Step-up resonator	74
4.19	Laser beam configuration	76
4.20	RD/BD/BDD UV frequency shifting setup	80
4.21	Raman UV frequency shifting setup	82
4.22	Zemax ray tracing of the imaging system used for qubit state readout	83
4.23	Laser scatter on the trap scatter for beam alignment	83
4.24	A 5-ion chain	85
4.25	Heating rates as a function of frequency	86
4.26	Electric field noise vs frequency in the cryostat	88
4.27	Electric field noise vs ion-electrode distance in a large variety of experiments	88
5.1	${}^9\text{Be}^+$ level structure (detection)	90
5.2	Ideal Poissonian histograms for short and long detections	92

5.3	Depumping and repumping	94
5.4	Ideal histograms vs depumped histograms	95
5.5	Sample sideband cooling data	100
6.1	Scalable vs traditional detection	107
6.2	SNSPD principles of operation	110
6.3	Schematic showing ion position over the detector trap	114
6.4	SNSPD readout circuit	115
6.5	Sample SNSPD pulses	116
6.6	SNSPD sensitivity	117
6.7	σ^- emission pattern	119
6.8	Saturated detector count rate	121
6.9	SNSPD spatial sensitivity	121
6.10	Simultaneous ion detection using the SNSPD and the PMT	122
6.11	High fidelity state detection	126
6.12	Elevated heating rate due to detector “clicks”	127
7.1	Final excitation of a quickly transported ion; classical treatment	134
7.2	Filter circuits	139
7.3	Predistorted and target waveform, DC4	140
7.4	Filter transfer functions	142
7.5	25.0 μ s transport with 0.03(1) quanta of heating	143
7.6	Final temperature after a cold transport as a function of delay time	145
7.7	Raman scans after a 9.1 μ s and 9.2 μ s transport durations	146
A.1	Radiation Shields	165
A.2	Black-body radiation absorbed by fused silica windows	165

Chapter 1

Introduction

Quantum information processing (QIP) holds the promise of a major increase in computing power. Rather than storing binary information as a classical signal that can take either of two binary values, a quantum computer encodes information as either of two quantum states, or the “qubit.” Crucially, a qubit can take either of these two quantum state values *or* a superposition of the two; information can also be spread between multiple qubits in an entangled quantum state. By exploiting these features of quantum superposition and entanglement, proposed quantum algorithms could empower substantial speed-ups in certain complex computational problems, including but not limited to database searches [6], factoring large numbers [7] and simulation of complex quantum systems [8].

Many physical platforms have been proposed and studied for quantum information processing applications, including trapped atomic ions, superconducting Josephson junctions [9], quantum dots [10], nitrogen vacancy centers in diamond [11], single photons [12], arrays of Rydberg atoms [13], and more. I will limit my discussion to trapped ions. Qubits encoded in two different electronic states of trapped ions provide a natural platform for quantum computing. The qubit state can be manipulated with high fidelity using high-intensity laser radiation [14][15][16] or microwave fields [17], and can have quantum coherence times as long as tens of seconds [18]. Neighboring ions naturally interact via the Coulomb interaction, and two-qubit gates have been demonstrated by coupling the states of the qubits to the ions’ motion with infidelities as low as 10^{-3} per gate [15][16].

1.1 Scalable quantum computing with trapped ions

The highest fidelity two-qubit gates [15][16] were performed in systems of two ions by coupling the qubit states to the shared motion of the ions in a single harmonic well. In general, interesting quantum algorithms require more than two qubits. However, the same basic scheme can be used to perform pairwise two qubit gates on ions within a longer chain by addressing shared modes of motion [19]. Development of these techniques has allowed trapped ions to enter what John Preskill dubbed the era of the “noisy intermediate-scale quantum” (NISQ) computer [20]. In the NISQ regime, systems of 10-100 non-error-corrected qubits will be used to perform certain algorithms that may demonstrate a quantum advantage. (NISQ is not a platform-specific regime: the recent highly-publicized Google quantum supremacy result was performed in a NISQ-scale system using superconducting qubits. [21]) Quantum computers using chains of trapped ions have been demonstrated in this regime, implementing such algorithms as the quantum approximate optimization algorithm (QAOA) [22] and rudimentary simulation of molecule binding energies [23][24].

Pushing past the NISQ regime into the realm of scalable, fault-tolerant quantum computing (as is required for algorithms like Shor’s factoring algorithm to present a speed-up over classical computers) requires many more than 50 ions, with gate errors below the level of the current best two-qubit gates. For technical reasons, it appears challenging to simply increase the number of ions in these long chains to hundreds or thousands of ions. Longer chains are increasingly sensitive to decrystalization or ion loss errors [25]; they have densely packed motional spectra, resulting in cross-talk errors; and the coupling rate between ions due to the Coulomb interaction decreases in as the system grows. Instead, this thesis is focused on scaling the quantum computer that in a way that is not subject to these technical limits.

1.1.1 The quantum CCD

Various research groups (including industrial ones) are actively investigating paths to scalability in trapped ion quantum information. A few are discussed below, in Section 1.1.2, but for the bulk of this thesis, I will focus on one possible solution: the “quantum CCD” [1][2].




Figure 1.1: A cartoon illustration of the quantum CCD, adapted from [2]. In this architecture, ions can be trapped in specialized zones anywhere along given “lanes,” and can be moved in and out of these zones. Specialized zones are shown here for (A) single qubit rotations, (B) two-qubit gates, (C) long term qubit storage, and (D) qubit state detection.

The quantum CCD (QCCD), schematically illustrated in Figure 1.1, can be thought of as an scheme to build a large-scale quantum computer in a modular fashion rather than attempting to control a single system. In order to motivate this approach, I emphasize a few of characteristics of trapped-ion quantum operations:

- Any quantum algorithm can be broken down into combinations of single qubit rotations and two-qubit entangling operations [26].
- The highest fidelity quantum gates to date were performed on small systems of one [14][15][16] or two [15][16] ion qubits.

- Although two-qubit gates are typically performed using collective motion between two ions in a shared harmonic well, the quantum information is generally *stored* in the electronic (pseudo-spin) states of the ion. In the absence of a spin-dependent force, the electronic state of an ion is well decoupled from its motion.

From here, we can build to the QCCD. Rather than trying to perform all operations on a single crystal of ions, we consider an apparatus with different “zones,” optimized for different components of trapped-ion-based computation. This would include zones for single- and two-qubit gates, detection, ion loading, and long-term memory. With the application of time-varying trapping potentials, ion qubits can be moved in and out of these zones and shifted to the memory zone when are not being acted on. Because the ion’s motion is decoupled from its qubit state, it is possible to move the ion without degrading the qubit. Initial use of this architecture has enabled demonstration of complex quantum algorithms, including quantum gate teleportation [27] and creation of a four-ion Greenberger-Horne-Zeilinger State [28].

1.1.1.1 Transport, separation, and merging of ion crystals

The ability to move ions in and out of different trap zones involves four basic operations:

- **Transport:** moving a harmonic well in one dimension
- **Separation:** splitting an ion chain in a shared harmonic well into two separate, individually controllable wells
- **Merge:** joining two ion chains in separate wells into a longer chain in a shared harmonic well
- **Reorder:** arbitrarily reordering an ion chain.

With these four operations, it is possible to bring two ions into a shared two-qubit gate zone, entangle them, and separate them into individual wells. Then, two different ions can be brought into the two-qubit gate zone and entangled. These building blocks can be used to build up to larger

quantum algorithms. The implementation of these operations in ion traps is discussed in Chapter 2.

1.1.1.2 Specialized trap zones

One other benefit of the QCCD is that it enables the use of trap zones which are optimized for specific operations, rather than a single zone which functions as a “jack of all trades, master of none.” To quickly illustrate how this might be useful, consider the example of ion loading. Although the improvement of loading efficiency is an active subject of investigation [29][30], established methods of loading ions into traps involve the generation of thousands or millions more ions than are trapped. These non-trapped ions can stick to nearby surfaces and create temporally- and spatially-varying stray electric fields. To avoid this problem, it is desirable to have separated zones for ion loading and high fidelity logical operations.

Using specialized trap zones not only limits deleterious effects of certain processes on gate fidelity but also enables the use of new technologies, which may be desired for only a single operation. Trapped-ion research groups are actively investigating numerous new technologies for quantum computing applications, including:

- High power microwave antennae enabling microwave-mediated multi-qubit gates. Shown in Figure 1.2(a). [17][31][34]
- Integrated digital-analog converters for direct on-chip generation of trap potentials. [35]
- Integrated optical waveguides and grating couplers for delivery of laser beams to ions. Shown in Figure 1.2(b). [32]
- On-chip collection optics for state detection. Shown in Figure 1.2(c). [36][37][33][38][39].
- Optics for the integration of trapped ions with standing-wave optical cavities [40].
- On-chip photodetectors. Shown in Figure 1.2(d). [41][3]

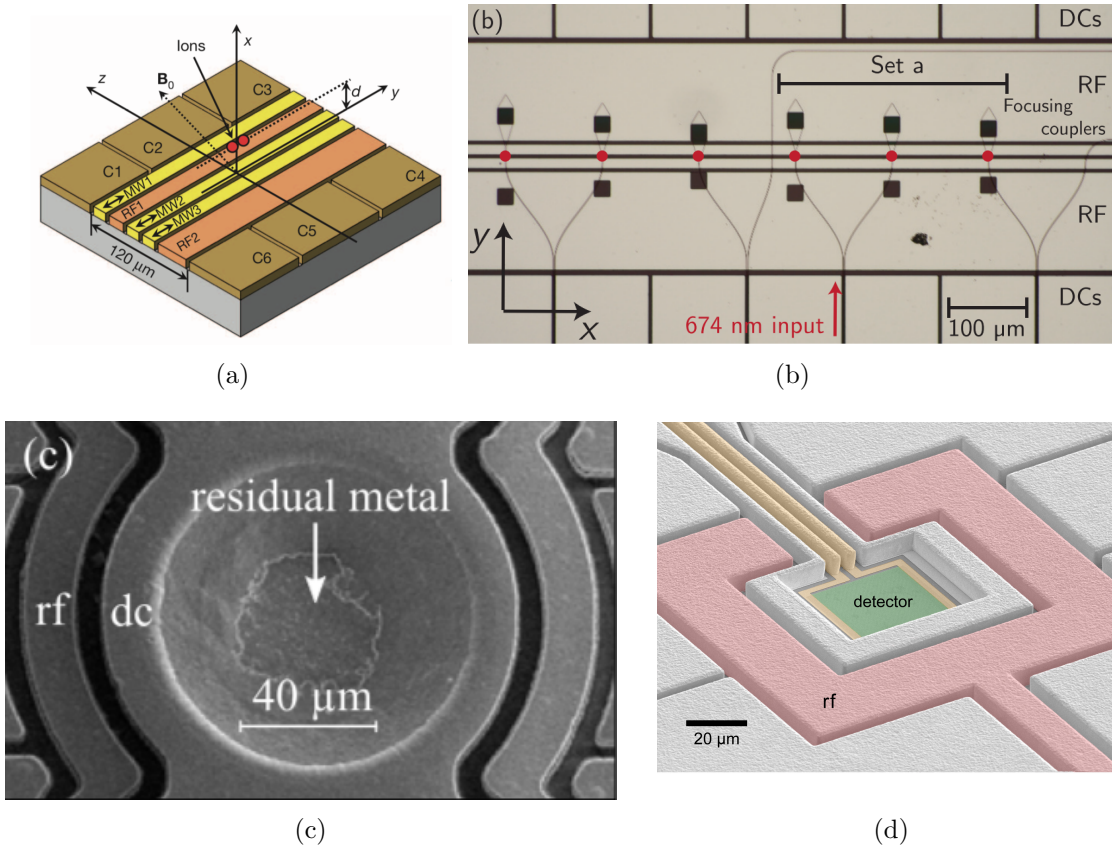


Figure 1.2: Examples of new ion trap technologies. (a) High power microwave antennae for the delivery of large microwave gradients. Microwave antennae are shown in yellow. From [31]. (b) A linear RF trap with integrated waveguides and grating couplers, enabling on-chip delivery of laser beams. From [32]. (c) An aluminum-coated micromirror integrated into a surface electrode trap geometry. From [33]. (d) A point trap with an on-chip photodetector. From [3].

It may be challenging to fabricate many of these technologies on a single chip, and operation of one technology may degrade another. However, by spatially separating trap zones and moving ions between these zones, it may be possible to incorporate multiple of these technologies onto a single chip.

1.1.2 Alternative scalable schemes

Trapped-ion quantum computing is now a broad research field, and there are investigation into several other methods for scaling the trapped ion quantum-computer in addition to the QCCD. It is likely that the truly useful scalable quantum computer will be some hybrid, incorporating

techniques developed through each of these areas of research. In the following sections, two of the other scalability schemes under investigation are described

1.1.2.1 Photonic links

Up to this point, the discussion has been limited to two-qubit entangling gates mediated by collective ion motional information. Ion motion is a natural resource for entanglement because the Coulomb force enables strong interactions between neighboring ions confined in a shared harmonic well. However, it is also possible to probabilistically entangle ions using light. Photons can travel long distances in fiber, with losses as low as < 1 dB/km at telecom wavelengths, so this scheme is a natural choice for remote entanglement of distant ions.

The primary drawback to this entanglement method is the rate of entanglement generation. Using state-of-the-art photon collection optics and high-quantum efficiency single photon counting, remote ions have been entangled using quantum interference of light in 5.5 ms with 94% fidelity [42]. In contrast, the fastest phonon-mediated two-qubit gates have gate times below a microsecond [43]. Although marginal improvements can be made (including attempts to integrate trapped ions into resonant optical cavities) geometric constraints will still limit the entanglement rate, since only a fraction of the emitted photons can be collected for entanglement. Therefore groups investigating this scheme typically envision using it as a secondary source of entanglement. In this long-term vision, there would be “nodes” of 10-50 ions, in which entanglement could be generated with high fidelity using motional modes. For algorithms involving more than 10-50 qubits, entanglement would be mediated between spatially separated nodes using photonic links [44].

1.1.2.2 Lattice traps

In the previous discussion of phonon-based entanglement, motion of ions was shared because they shared a single harmonic trap. However, since the Coulomb interaction is long-range, these ions need not be trapped in the same well to share motional information. Consider as a lowest-order example, two ions trapped in a double-well trap, as shown in Figure 1.3. If the two wells

are exactly on resonance with each other (meaning that their harmonic oscillator frequencies are the same) these separated ions can be entangled [45]. This interaction is naturally tunable, as it can be essentially turned on and off by tuning the two potential wells in and out of resonance. Investigations into this scalability scheme are still at an early stage. Ions have been entangled in a double well as described above, and preliminary attempts are underway to entangle ions at the three vertices of an equilateral triangle trap [46][47]. Such a scheme is appealing from a scalability standpoint because if ions could be controllably coupled in a two-dimensional array, the two-dimensional linkages between nodes enable better scaling for quantum algorithms. It is also an interesting platform for certain quantum simulation problems [48].




Figure 1.3: A double-well surface electrode trap, in which the interaction between two separated ions can essentially be turned on and off by bringing the ions in and out of resonance. Adapted from [45].

1.2 This work

This thesis describes addresses challenges in scaling the trapped ion quantum computer. The focus is on the QCCD. owever, many of the technical improvements discussed are applicable to other scalability schemes. The chief results of this thesis are the demonstration and characterization of a trap-integrated single photon detector for high fidelity qubit state readout and the development and demonstration of a new scheme for faster-than-adiabatic cold transport of ion qubits with reduced sensitivity to transport duration.

Chapter 2 reviews ion trapping in Paul traps and the generation of trap voltages, focusing on surface electrode traps. Chapter 3 reviews the ${}^9\text{Be}^+$ qubit and its interactions with laser light

that enables quantum control. Chapter 4 provides a comprehensive description of a new cryogenic apparatus developed for the experimental work central to this this dissertation. Chapter 5 describes calibrations in this apparatus and reviews readout methods, including high fidelity state detection. Chapter 6 describes a new trap-integrated superconducting nanowire single-photon detector (SNSPD), and measurements that characterize the performance of this device for high-fidelity quantum state readout are presented. Chapter 7 presents a new scheme for faster-than-adiabatic transport of ions, and this scheme is demonstrated with low net motional excitation and reduced sensitivity to transport duration relative to previous ion transport methods.

1.2.1 Development of a cryogenic apparatus for trapped ion experiments

All trapped-ion scalability schemes benefit from the ongoing development of microfabricated surface electrode traps (discussed in depth in Chapter 2). The high precision of microfabrication techniques enables the reliable production of larger, more complex ion trapping systems with tiny electrodes. However, users of surface traps face two major technical challenges:

- **Reduced Trap Depths:** Trapping with microfabricated chips is enabled by generating an RF quadrupole potential above the chip surface. However, away from the center of the quadrupole, there is higher anharmonicity compared to a traditional, four-rod, linear RF Paul trap [49]. This increases ion-loss due to background gas collisions.
- **Anomalous Heating:** All ion traps have electric field noise that is known to emanate from electrode surfaces [50]. This noise leads to ion heating, and because the underlying physics is not yet understood, this heating is referred to as “anomalous.” For trapped-ion quantum computing, this heating can be a significant source of errors. This is especially the case for devices that use surface-electrode traps, since ions in these traps are confined close to the surface, and the electric field noise amplitude has been shown to scale as d^{-4} , where d is the approximate distance to the nearest electrode [51][52][53].

Both of these challenges can be mitigated by operating the trap under cryogenic conditions. Since

the residual background gases common in ultra-high vacuum chambers all freeze out at temperatures below 10 K, a cryo-pumped vacuum chamber can have background gas pressures as low as 10^{-13} – 10^{-17} torr [54][25], a level at which background gas collisional loss is essentially eliminated, even in traps with low depth. Additionally, it was discovered that anomalous heating is usually reduced by cryogenically cooling the trap electrodes [55][56]. Although the temperature effect is not yet understood, several groups are actively studying this [57][58].

In addition to increased ion lifetimes and reduced anomalous heating, a cryogenic environment provides another technical benefit. Operating cryogenically enables the use of superconducting devices. Of particular interest for this thesis, this includes superconducting nanowire single photon detectors (SNSPDs) [59], which are discussed further in Chapter 6.

The new cryogenic apparatus discussed in Chapter 4, is a low-vibration closed-cycle cryogenic system. This apparatus has sufficiently good vacuum to trap a Doppler-cooled ${}^9\text{Be}^+$ ion for days and has a base temperature below 3.5 K, low enough to operate an SNSPD similar to those discussed in [3]. Additionally, the system was designed for quick and reliable exchange of traps, enabling quicker testing of new trap designs and technologies in the future. Vibration measurements from this system are also presented, showing < 1 μm peak-to-peak. This vibration may need to be reduced for experiments which require very high-fidelity gates, but is sufficient for the experiments I present in this work.

1.2.2 Scalable high-fidelity qubit state detection

Qubit state detection in trapped-ion quantum computing is performed by state-selective fluorescence detection. The quantum information is mapped onto either a “bright” state, which has a resonant cycling transition at a desired laser wavelength, or a “dark” state, which does not have a resonant transition at that wavelength. (This is discussed in the specific case of ${}^9\text{Be}^+$ in Chapter 5.) If the ion is found in the bright state, an incident photon at that wavelength will excite this cycling transition. When it falls back down to the bright state, the ion will emit a photon at the same wavelength into 4π , which can be detected. If it is in the dark state, no photons will be

counted. Detecting this re-emitted fluorescence requires sensitive single photon detection, resulting in typical count rates of 10's of photons in hundreds of microseconds. This is usually performed by imaging the ion fluorescence with high-numerical-aperture (NA) lenses onto distant high-quantum-efficiency sensors. High-NA imaging systems enable capture of photons from a large solid angle to maximize the photon counting signal. However, these imaging systems are challenging to design, with multiple conflicting design objectives.

- The objective must have high NA, so that it can collect a large fraction of the emitted photons.
- It is helpful if the objective is also high magnification and low aberrations so it can be used to image laser scatter off the trap electrodes during initial laser alignment.
- For many ions, including ${}^9\text{Be}^+$, all optics must be UV-compatible, so commonly used optical glasses like BK7 cannot be used.
- For experiments incorporating multiple species of ions [60][61], it is highly desirable that the objective operate simultaneously at the multiple wavelengths used to detect each species.
- In many setups, the objective must have a long working distance, so it doesn't interfere with vacuum hardware.
- Larger, multi-ion systems may require a wide field of view to detect multiple ions at once.

Accommodating all of these design constraints is challenging and requires trade-offs. More, these systems are relatively bulky, expensive, and represent increased experimental complexity.

Any scalable scheme for trapped-ion quantum computing that seeks to detect many ions in parallel will require an alternative to this traditional method of fluorescence detection. There are many possible solutions. This dissertation presents one such solution: the integration of a photon sensor directly into an ion trap chip, eliminating the need for collection optics. Traps incorporating photon detectors have been previously demonstrated, including a commercial photodiode positioned beneath a transparent, electrically conductive electrode [41] and on-chip avalanche

photodiodes (APDs) [32]. Another option for high-efficiency single photon counting is a superconducting nanowire single-photon detector (SNSPD). These detectors (discussed in detail in Section 6) have high quantum efficiency, low dark counts, and appropriate count rates for trapped-ion experiments [62]. Previous work has shown that SNSPDs operate well at the UV wavelengths that are used for state detection in many ions [3][63] (up to 76% quantum efficiency at 315 nm) and under conditions that suggested that such a detector could be integrated into a trap chip [3]. Demonstrating this trap integration was one of the primary goals of the work described in this dissertation, and Chapter 6 describes the successful implementation. Results include the first demonstration of trapping ions in a device with an on-chip SNSPD and using this SNSPD for state detection with $> 99.9\%$ detection fidelity.

1.2.3 Fast ion transport

In order to demonstrate a practical trapped-ion quantum computer, improvements to the speed of each operation are still needed. Others are working on faster quantum gates, and the current record is $1.6 \mu\text{s}$ with 99.8% fidelity [43]. The work described in this thesis focuses on one of the major issues slowing down algorithms in QCCD architectures: the speed of ion transport.

For context, consider the example quantum algorithm illustrated in Figure 1.4, which was performed using a QCCD-style approach, from [60]. The experiment is a modified Bell test experiment, providing a careful measurement of the extent to which the state of two entangled ions cannot be expressed as a classical product state. This experiment illustrates the issue of transport as a major time-sink in a QCCD-type experiment. The timeline of this experiment is shown in Figure 1.4. Transport steps (shown in orange and light blue) took $250 \mu\text{s}$ each, while separation and recombination (in yellow and light green, respectively) steps took $300 \mu\text{s}$ each. All transport, separation, and recombination steps combined took 1.9 ms. Neglecting initial ground state cooling and final state detection, the full experiment took 2.3 ms.

In this example, transport and separation and separation represent a large fraction of the experimental timeline. This is generally true of quantum-CCD-style algorithms [64]. This is a chal-

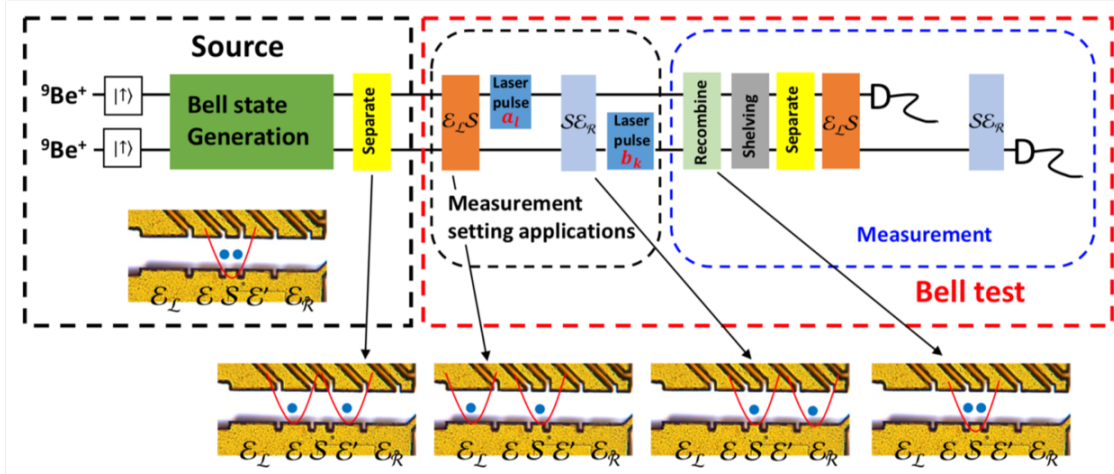


Figure 1.4: Timeline of a sample experiment performed using a quantum-CCD architecture, with adiabatic transport. Transport, separation, and recombination operations take up 1.9 ms, 83% of the time not devoted to initial ground-state cooling and final state detection. From [60].

lenge, not just because we wish to be able to perform operations quickly but because as algorithms grow longer, the ions heat more due to both technical noise and anomalous heating. The highest fidelity gates are performed on ions cooled to the motional ground state, so long algorithms need additional time-intensive re-cooling steps before high-fidelity steps, *further* extending the runtime of the algorithm. One might naïvely simply perform the transport more quickly; however, if not performed carefully, speeding transport up introduces further heating and therefore gate errors. On timescales of tens to hundreds of trap motional periods (typically on the order of $0.1 - 10 \mu\text{s}$), the ion is largely insensitive to the motion of the harmonic well. This is generally referred to as “adiabatic” transport. However, as the transport speeds up to timescales approaching the timescale of a few tens of trap motional periods, the ion’s motion is excited by the motion of the potential well. Under these conditions, the ion can end up in a highly excited motional state at the end of the transport, requiring re-cooling.

Transport can be performed faster than adiabatically if the ion is transported in an integer number of trap periods. This can be thought of classically as the ion receiving a “kick” as the well starts moving and an equal and opposite “kick” at the end of the transport. (This is formalized more in Chapter 7.) Ions have previously been transported a single zone in relatively large, mi-

crostructured 3D traps with transport times of $3.6 \mu\text{s}$ [5] and $8 \mu\text{s}$ [4] with ~ 0.1 quanta of motional excitation due to transport. However, in practice, this method is sensitive to both fluctuations in the trap secular frequency and miscalibration of the transport duration and thus requires frequent careful calibration. Further, there is a single free parameter (the time of the transport) on which to optimize; this does not scale well to systems of more degrees of freedom. In Chapter 7, a new scheme for fast transport of ions which has more degrees of freedom and is thus generalizable to more complex systems and which is less sensitive to the duration of the transport, is presented and demonstrated.

Chapter 2

Ion Traps

We require a method of trapping and storing the ions for long times; ideally, many days. Although there are other methods of trapping ions, I only discuss here linear RF Paul traps.

2.1 RF Paul Traps

Ions interact strongly with electric fields, which are generated by applying voltages to nearby electrodes. However, it is impossible to trap an ion with only static electric fields, since we are bound by Maxwell's equations. Specifically, a 3D confining electric potential would violate Gauss's Law, which necessitates that electric fields be divergence-free in the absence of sources:

$$\nabla \cdot \vec{E} = 0 \tag{2.1}$$

[65]. It is, however, possible to generate a trapping potential using purely electric fields if we consider the trap in a *time-averaged* way. Gauss's Law permits confinement in two dimension but de-confines in the third dimension [65]. However, if we quickly alternate between confining and de-confining, the ion will see, on average, a confining potential, commonly called the "pseudopotential" [1].

First define the system geometry. A schematic of a simple linear Paul trap is shown in Figure 2.1(a). A radio frequency tone $V_{RF}(t) = V_{RF} \cos(\Omega_{RF}t)$ at amplitude V_{RF} and angular frequency Ω_{RF} is applied to electrodes A and C (the RF electrodes). Electrodes B2 and D2 are grounded. A DC voltage V_{DC} is applied to the endcap electrodes B1, B3, D1, and D3. In a transverse slice to

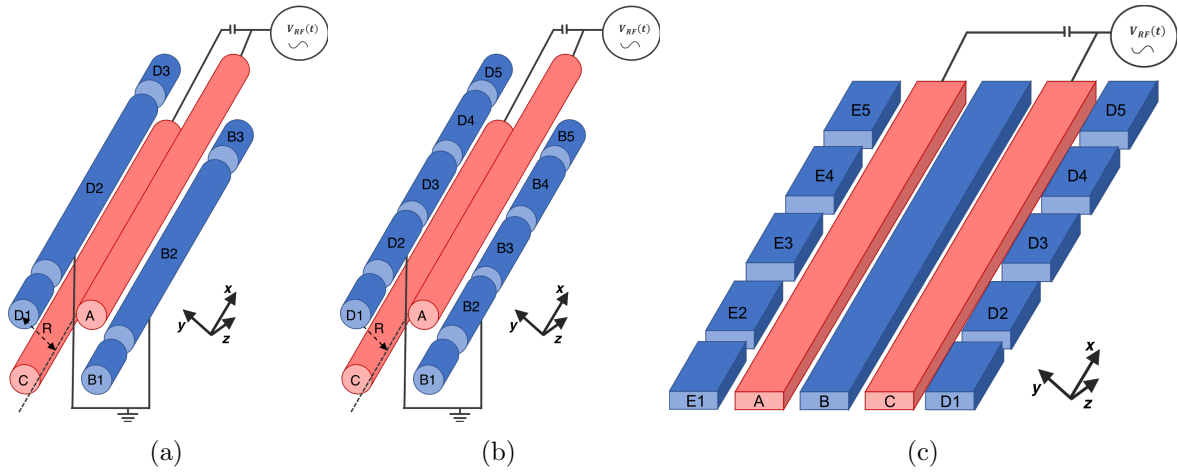


Figure 2.1: Schematics of three linear Paul traps. (a) A classic linear Paul trap. RF applied to the red electrodes A and C will produce a quadrupole transverse to the central axis, a distance R from the electrodes, as described by Eq. 2.2, confining the ion in the yz plane. Confinement along the trap axis is provided by DC voltages applied to electrodes B1, B3, D1, and D3. (b) A segmented Paul trap. The RF confinement is identical to the case in (a), but DC confinement can be provided by any combination of the DC electrodes, enabling multiple trapping zones. (c) A surface electrode trap, functionally equivalent to the case in (b) but that has been “unfolded” onto a single plane so that it can be made using microfabrication techniques.

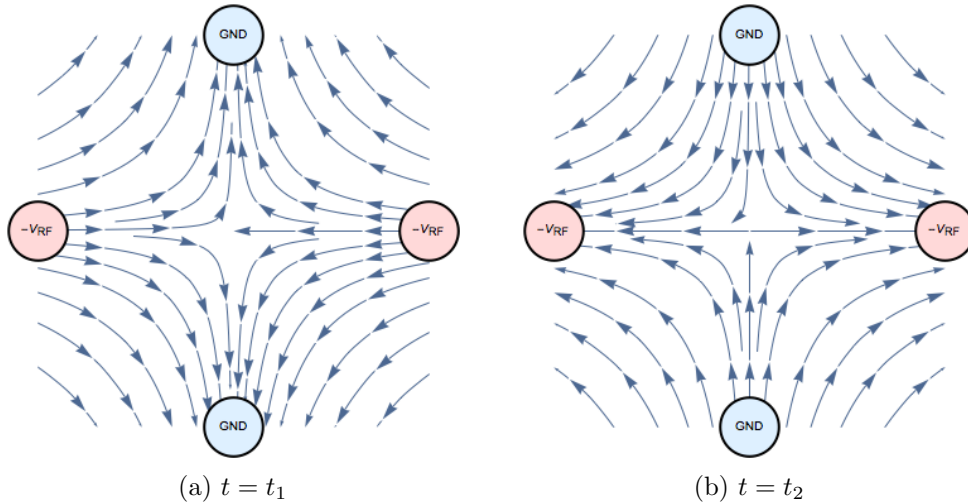


Figure 2.2: Electric field lines in a quadrupole ion trap, at times t_1, t_2 corresponding to two opposite points of the RF phase, i.e. $t_2 = t_1 + 2\pi/\Omega_{RF}$.

the axis of this trap, the electric field at time $t = 0$ will be as shown in Figure 2.2(a), with field lines directed from electrodes A and C towards electrodes B and D. There is zero electric field at the very center of the trap, but this is a saddle point of the potential not a minimum. At this time, an ion will travel along the field lines towards GND.

At a later time corresponding to a 180° phase change, the electric field lines transverse to the axis will appear as shown in Figure 2.2(b), with field directions reversed from those in Figure 2.2(a), directed from electrodes B and D towards electrodes A and C. At this time, the force on an ion will be reversed, and it will instead travel towards the RF electrodes.

For an ion near the central axis, the electric potential is

$$\Phi(x, y, z) \approx \frac{qV_{RF}}{2} \cos(\Omega_{RF}t) \left(1 + \frac{y^2 - z^2}{R^2} \right) + \kappa V_{DC} \left(x^2 - \frac{1}{2}(y^2 + z^2) \right) \quad (2.2)$$

for ion-electrode distance R , where κ is a geometric factor that depends on the shape and positioning of the electrodes [1]. This has a time dependent and static components. The static component (provided by the DC voltages on the endcap electrodes) is deconfining in y and z . For this to function as a three-dimensional confining potential, the values of V_{RF} , Ω_{RF} , and V_{DC} must be such that the confinement provided by the time-dependent piece of this potential overcomes the deconfinement of the static piece. Eq. 2.2 gives ion equations of motion in the radial direction described by the Mathieu equation

$$\frac{d^2y}{d\zeta^2} + (a_y + 2q_y \cos(2\zeta)) y = 0 \quad (2.3)$$

$$\frac{d^2z}{d\zeta^2} + (a_z + 2q_z \cos(2\zeta)) z = 0 \quad (2.4)$$

where

$$\zeta \equiv \Omega_{RF}t/2 \quad (2.5)$$

$$a_y = a_z = -\frac{4\kappa q V_{DC}}{m x_0^2 \Omega_{RF}^2} \quad (2.6)$$

$$q_y = -q_z = \frac{2q V_{RF}}{m R^2 \Omega_{RF}^2} \quad (2.7)$$

[1]. Ion traps are generally operated in the limit

$$a_x < q_x^2 \ll 1 \quad (2.8)$$

[1]. In this limit, the ion's position $x_i \in \{y, z\}$ as a function of time is described by

$$x_i(t) \propto \cos(\omega_i t + \phi_i) \left(1 + \frac{q_i}{2} \cos(\Omega_{RF} t) + \frac{q_i^2}{32} \cos(2\Omega_{RF} t) \right) + \beta_i \frac{q_i}{2} \sin(\omega_i + \phi_i) \sin(\Omega_{RF} t) \quad (2.9)$$

where $\beta_i = \sqrt{a_i + q_i^2/2}$ [1]. The quickly varying terms in this are expression generally referred to as the micromotion terms and the motion at frequency ω_i is the trap secular motion. In practice, we operate in conditions where the micromotion is minimized, so it is a reasonable approximation to consider only the secular motion. This acts as a ‘‘pseudopotential’’, described by

$$\Phi_{pp}(y, z) = \frac{m\omega_r^2}{2q}(y^2 + z^2) \quad (2.10)$$

where

$$\omega_r \approx \frac{qV_{RF}}{\sqrt{2}\Omega_{RF}mR^2}. \quad (2.11)$$

[1]. I have briefly neglected motion in the axial direction x in this discussion. The x -dependent piece of Eq. 2.2 is

$$\Phi_{ax}(x) = \kappa V_{DC} x^2 = \frac{1}{2q} m \omega_{ax}^2 x^2 \quad (2.12)$$

describing harmonic confinement with secular frequency

$$\omega_{ax} = \sqrt{\frac{2\kappa q V_{DC}}{m}} \quad (2.13)$$

I will refer to the three secular frequencies as the axial (ω_{ax}) and radial (ω_{r1}, ω_{r2}) frequencies, and motion along the corresponding directions as the axial and radial modes, respectively.

The limit for which we made the pseudopotential approximation, Eq 2.8, is important to consider because it is a good guide for trap design and selection of trapping parameters. RF and DC voltage amplitudes and RF frequency should all be chosen so that this limit holds.

2.2 Ion motion in Paul traps

There are two types of ion motion to consider in ion traps: *secular motion*, the harmonic oscillator motion of the ion; and *micromotion*, fast oscillations at the RF frequency.

2.2.1 Secular motion: the harmonic oscillator

The potential $\Phi_{pp}(\vec{x}) + \Phi_z(\vec{x})$ can be described by the second-order curvature tensor

$$C_{ij} = \frac{\partial^2}{\partial x_i \partial x_j} (\Phi_{pp}(\vec{x}) + \Phi_z(\vec{x})) . \quad (2.14)$$

In the harmonic limit, we treat that all higher-order curvatures $C_{ijk} = 0$; this is generally true in the limit of small amplitude oscillations. An ion in such a potential experiences the Hamiltonian

$$\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m \left(\sum_{i,j \in (1,2,3)} C_{ij} \hat{x}_i \hat{x}_j \right) \quad (2.15)$$

We generally operate in the curvature eigenvector basis, in which case this can be rewritten as

$$\hat{H} = \sum_{i \in (1,2,3)} \left(\frac{\hat{p}_i^2}{2m} + \frac{1}{2}m\omega_i^2 \hat{q}_i^2 \right) \quad (2.16)$$

where ω_i is the eigenvalue of the curvature along the i th direction. and we can treat this as three decoupled harmonic oscillators. I will refer to the eigenvectors \vec{q}_i as the “secular modes” and the eigenvalues as the “secular frequencies,” as they correspond directly to the direction and frequency of the ion’s harmonic motion.

In the eigenvector basis, the Hamiltonian decouples into a sum of three independent harmonic oscillator Hamiltonians. For simplicity of notation, I will for the rest of this treatment discuss the one-dimensional Hamiltonian of a single secular mode. As is derived in most undergraduate quantum mechanics textbooks, including [66], the position-space solutions to the harmonic oscillator Hamiltonian have equally spaced eigenenergies

$$E_n = \hbar\omega \left(n + \frac{1}{2} \right) \quad (2.17)$$

This is often referred to as the harmonic oscillator “ladder,” since $E_n - E_{n-1} = \hbar\omega$ for all n . These energy eigenstates are generally what we work with in trapped ion. It is thus useful to move from the position/momentum basis described by \hat{q}, \hat{p} into a basis described by creation/annihilation operators \hat{a}^\dagger, \hat{a} , where

$$\begin{aligned} \hat{a} &= \sqrt{\frac{m\omega}{2\hbar}} \left(\hat{q} + \frac{i}{m\omega} \hat{p} \right) \\ \hat{a}^\dagger &= \sqrt{\frac{m\omega}{2\hbar}} \left(\hat{q} - \frac{i}{m\omega} \hat{p} \right) \end{aligned} \quad (2.18)$$




Figure 2.3: A phase space illustration of the displacement operator with parameter $\alpha = |\alpha|e^{i\phi}$ acting on the ground state of the harmonic operator $|0\rangle$.

This basis is called the Fock basis. These operators are so-called “annihilation” and “creation” operators, respectively because they transfer the ion from $|n\rangle \rightarrow |n \pm 1\rangle$ by

$$\begin{aligned}\hat{a}|n\rangle &= \sqrt{n}|n-1\rangle \\ \hat{a}^\dagger|n\rangle &= \sqrt{n+1}|n+1\rangle\end{aligned}\tag{2.19}$$

With these operators, we can rewrite the Hamiltonian for a harmonic oscillator as

$$H = \hbar\omega \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right)\tag{2.20}$$

Although I will not work through full derivations of the above relations, these can be found in [67] or any standard quantum optics text.

Considering a harmonic oscillator’s motion in the Fock basis is a powerful tool that has enabled the creation of interesting nonclassical states of motion in trapped ions [68][69][70]. I will, however, limit myself to discussion of a classical operation in this basis: the displacement operator.

$$\hat{D}(\alpha) = \exp\left(\alpha\hat{a}^\dagger - \alpha^*\hat{a}\right)\tag{2.21}$$

where α is a complex number $\alpha = |\alpha|e^{i\phi}$. The effect of the displacement operator on the ground state of the harmonic oscillator is shown in (q, p) phase space in Figure 2.3. As the name suggests, it displaces the state in phase space. A displaced ground state gives the motion classically associated with a harmonic oscillator: the ion oscillating in its well with amplitude $|\alpha|$ and phase ϕ . As I will discuss in Chapter 7, moving a harmonic well relative to an ion is equivalent to acting on the ion with the displacement operator [71].

2.2.2 Micromotion

In approximating the confining potential as Equation 2.10, I have neglected terms which oscillate at Ω_{RF} , $2\Omega_{RF}$, $3\Omega_{RF}$, etc. This is a good approximation for an ion located on a null of the RF potential. However, these terms - the ion's "micromotion" - grow with distance from the axis. Because an ion which is oscillating at Ω_{RF} can have transition frequencies which are substantially Doppler shifted, we wish to operate in a regime in which this micromotion is minimized, i.e. with the ion located on the central trap axis. For a single trapped ion, this is a matter of minimizing stray electric field, which would push the ion off the RF null. For systems with larger numbers of ions, micromotion minimization requires both cancelling stray electric field and operating in conditions in which all of the ions are located on the RF null. This can be ensured by operating with $\omega_{ax} \ll \omega_{r1, r2}$, in which case the ions crystalize in a 1D chain along the pseudopotential null [1].

2.2.3 Normal Modes

In long chains of ions, we must consider not only the externally applied potential but for the Coulomb interaction between neighboring ions. So, for the k th ion in an ion chain, the Hamiltonian is

$$\hat{H} = \frac{\hat{p}^2}{2m_k} + \Phi_{ext}(\hat{x}_k) + \frac{1}{4\pi\epsilon_0} \sum_{l \neq k} \frac{e^2}{|\hat{x}_l - \hat{x}_k|} \quad (2.22)$$

where I allow for different ions in the chain to have different masses, but have forced that all ions in the chain have charge $+e$.

In general, fully solving for the Hamiltonian accounting for the Coulomb interaction is a challenging problem. However, we can simplify things substantially working in the limit of small oscillations around equilibrium $(\vec{x}_1^{(0)} \dots \vec{x}_N^{(0)})$, where

$$\left. \frac{\partial \Phi}{\partial x_k} \right|_{x_k^{(0)}} = 0 \quad (2.23)$$

Solving for the equilibrium positions $x_k^{(0)}$ analytically becomes more challenging for longer, mixed-species ion crystals, but it is a straightforward numerical problem.

Rather than the 3×3 curvature matrix of Equation 2.14, we must now consider $3N$ coordinates, as the three-dimensional position of each ion will define the shape of the potential. So we actually have

$$C_{ij} = \left. \frac{\partial^2 \Phi}{\partial x_i \partial x_j} \right|_{x_i^{(0)}} \quad (2.24)$$

where Φ includes both the externally applied field Φ_{app} and the Coulomb potential Φ_{Coul} . This is a larger matrix than Equation 2.14, but it is still a two dimensional tensor. Then the equations of motion for dimension j (one of the 3 spatial parameters of ion k) can be written

$$m_k \delta \ddot{x}_j = - \sum_{i=1}^{3N} C_{ij} \delta x_i + O(\delta x^3) \quad (2.25)$$

where $\delta x_i = x_i - x_i^{(0)}$ is the perturbation from equilibrium. In the limit of small perturbations, we can drop the higher order terms to obtain a linear equation. We can then derive a new set of coordinates - the “normal modes” - which diagonalize this linear equation.

In the simplest case, consider a two-ion system in one dimension with ions of identical mass, so

$$\hat{\Phi} = \frac{\hat{p}_1^2 + \hat{p}_2^2}{2m} + \frac{1}{2} m \omega^2 (\hat{x}_1^2 + \hat{x}_2^2) + \frac{1}{4\pi\epsilon_0} \frac{e^2}{\hat{x}_1 - \hat{x}_2}, \quad (2.26)$$

where $x_1 > x_2$ was chosen with no loss of generality. The equilibrium positions of such a crystal are

$$x_{1,2}^{(0)} = \pm \left(\frac{1}{16\pi\epsilon_0} \frac{e^2}{2m\omega^2} \right)^{1/3} \quad (2.27)$$

relative to the single-ion equilibrium position. The curvature matrix around these equilibrium positions is given by

$$C_{ij} = \begin{pmatrix} \left. \frac{\partial^2 V}{\partial x_1^2} \right|_{x_1^{(0)}} & \left. \frac{\partial^2 V}{\partial x_2 \partial x_1} \right|_{x_2^{(0)}} \\ \left. \frac{\partial^2 V}{\partial x_1 \partial x_2} \right|_{x_1^{(0)}} & \left. \frac{\partial^2 V}{\partial x_2^2} \right|_{x_2^{(0)}} \end{pmatrix} = \begin{pmatrix} C & -C/2 \\ -C/2 & C \end{pmatrix} \quad (2.28)$$

where

$$C = m\omega^2 \quad (2.29)$$

This can be diagonalized with eigenvectors

$$\vec{x}_{com} = \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad \vec{x}_{str} = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \quad (2.30)$$

commonly notated as the “center of mass” and “stretch” modes, respectively. These have eigenenergies $m\omega^2/2$ and $3m\omega^2/2$, corresponding to eigenfrequencies ω and $\sqrt{3}\omega$.

2.3 Surface electrode traps

For the hundreds or thousands of qubits which will be necessary for larger scale algorithms, hundreds or thousands of electrodes will be likely required. As the number of electrodes grows, so too do the tolerance requirements on the fabrication of these electrodes. If the fabrication scheme has a 1% fatal error rate per electrode; this is acceptable for a system like Figure 2.1(b) with just five electrodes but unacceptable for a large-scale system with a thousand electrodes.

However, it is not necessary to reinvent the wheel. Classical computers are built daily with tolerances on the order of 10’s of nanometers using a whole host of powerful microfabrication techniques. But in order to take advantage of these technologies, we need a trap which lies on a single plane. Fortunately, the four-rod trap can be “unwrapped” into a single plane [49]. This is illustrated in Figure 2.1c. For small perturbations from the potential minimum, the surface trap produces a RF quadrupole as described by Eq. 2.2.

Such an ion trap - the “surface electrode trap” - introduces new challenges. The approximation that the electric field is a pure RF quadrupole is valid for much smaller excursions from the RF




Figure 2.4: Examples of surface electrode traps. (a) The “high optical access 2.0” trap, fabricated by Sandia National Labs and distributed to many groups [72]. (b) The “triangle trap,” which traps ions in three independent wells separated by $30\ \mu\text{m}$ [47]. (c) A linear trap fabricated in a commercial CMOS facility [32]

null [49]. In a shallow trap, the ions are more susceptible to background gas collisions, which are typically the dominant cause of ion loss in trapped ion experiments. However, with laser cooling and an apparatus with adequately low vacuum pressure, like the one I will describe in Chapter 4, it is possible to hold ions in a surface electrode trap for days at a time.

Surface electrode traps can have relatively small ion-electrode distances (as low as $30\ \mu\text{m}$ has been demonstrated [34]), which enables the application of stronger electric fields and curvatures (and potentially faster gates) using more moderate electrode voltages. However, with these small electrode-ion distances, it becomes increasingly challenging to focus lasers tightly onto an ion while limiting scattering on nearby trap surfaces. This scattering is particularly of concern when working with ions like ${}^9\text{Be}^+$ or ${}^{25}\text{Mg}^+$, which have cooling wavelengths near $313\ \text{nm}$ and $280\ \text{nm}$, respectively, since high intensity UV laser scatter on a dielectric will introduce stray charge through to the photoelectric effect [66].

Due to the reliability of fabrication and the complexity electrode geometries achievable, surface electrode traps are increasingly becoming the workhorse of the trapped ion quantum computing community. Surface electrode traps are beginning to be produced and distributed more widely so that every research group need not make their own. The “high optical access 2.0” trap from Sandia

[72] has been distributed to many trapped ion research groups, and MIT Lincoln Lab has successfully trapped ions in traps fabricated in a commercial CMOS foundry [32]. Some representative or interesting surface electrode traps are shown in Figure 2.4. In this thesis I used surface electrode traps (discussed in more detail in Section 4.2.1) fabricated in the Boulder Microfabrication Facility at NIST by Daniel Slichter and Varun Verma.

2.3.1 Simulations of surface electrode traps

In large 3-dimensional traps, it is relatively straightforward to find DC electrode voltages that produce a harmonic potential well. In a trap with 5 segmented electrodes, as illustrated in Figure 2.1(b), a well can be produced in zone C by applying a high voltage $V_B = V_D > 0$ to DC electrodes B and D, and a low voltage $V_C < 0$ to DC electrode C. Since the electrodes are far from the ion, the potentials on the end electrodes V_A, V_E will apply relatively minor corrections to the potential at the ion's location. To tune the trap frequency, we can adjust the amplitudes of the electrode voltages V_B, V_C, V_D . The same general principle can be used to produce trapping potentials in a surface electrode trap. However, since the ion is an order of magnitude closer to the electrodes, we must be much more careful to choose voltages that create a deeper and more harmonic trapping potential.

To find these voltages, we need the electric potential produced at the ion location due each electrode. The exact geometry of the trap electrode substantially impacts the shape of the potential. This can be simulated precisely using COMSOL or boundary element methods. However, for a surface electrode trap, there is a quicker and simpler alternative. For an electrode in a single plane with a voltage V , there exists an analog to the Biot-Savart law for electrostatics to give the electric field $\vec{E}(\vec{x})$ at position \vec{x} :

$$\vec{E}(\vec{x}) = \frac{V}{2\pi} \oint \frac{(\vec{x}' - \vec{x}) \times d\vec{s}'}{|\vec{x}' - \vec{x}|^3} \quad (2.31)$$

where the integration limits are the edges of the electrode [73]. This allows for relatively quick (i.e. calculable on a laptop) calculation of the electric fields generated by each electrode.

Electric potentials sum linearly, so we can write a linear equation to find the voltages that will give our target potential. I will identify a few characteristics of the desired DC electrical potential:

- (1) There should be no electric field at the well center:

$$E_i \propto \frac{\partial}{\partial x_i} \Phi(\vec{x}_0) = 0 \quad (2.32)$$

for $x_i \in (x, y, z)$.

- (2) The second-order curvature tensor (Eq. 2.14) should generate the desired mode frequencies and orientations. Specifically,

$$C_{xx} = \frac{\partial^2}{\partial x^2} \Phi(\vec{x}_0) = m\omega_x^2 \quad (2.33)$$

where I have, without loss of generality, defined x to lie along the trap axis. The radial confinement in y, z is provided by the RF potential, not by the externally applied DC potential. However, the radial curvatures will shift upon application of C_{xx} , since by Laplace

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \right) \Phi(\vec{x}) = 0 \quad (2.34)$$

[65]. So only $\partial^2 \Phi / \partial x^2$ and $\partial^2 \Phi / \partial z^2$ need to be constrained. From these, $\partial^2 \Phi / \partial y^2$ is fully set. If we wish to shift the frequency of one of the radial modes by a given amount, we can apply a DC curvature. Typically, to simplify things, one of these shifts will be set to zero

$$\frac{\partial^2}{\partial x^2} \Phi(\vec{x}_0) = 0 \quad (2.35)$$

We must also consider the off-axis components of the curvature matrix

$$\frac{\partial^2}{\partial x_i \partial x_j} \Phi(\vec{x}_0) \quad (2.36)$$

for $i \neq j$. The curvature tensor is symmetric, so $C_{ij} = C_{ji}$. These off-axis components rotate mode orientations, so a non-zero curvature C_{ij} will rotate the modes around the $x_i x_j$ axis. For efficient Doppler laser cooling, the modes \vec{q}_i should be oriented so that all modes

have some overlap with the cooling beam. We can work in the coordinate space defined by the modes C'_{ij} that have the desired mode orientations and rotate into the xyz frame by

$$C_{ij} = R^T C'_{ij} R \quad (2.37)$$

for rotation matrix R . In this frame, all of the off-diagonal terms

$$C'_{xy} = C'_{yz} = C'_{xz} = 0 . \quad (2.38)$$

- (3) The well should be harmonic. This can be forced to lowest order in the axial dimension by requiring the third axial derivative to be zero,

$$C_{xxx} = \frac{\partial^3}{\partial x^3} V(\vec{x}_0) = 0 \quad (2.39)$$

It is useful to define a vector containing our target derivatives:

$$\begin{aligned} \vec{\varphi}^{(0)} &\equiv [E_x, E_y, E_z, C_{xx}, C_{zz}, C_{xy}, C_{yz}, C_{xz}, C_{xxx}]^T \\ &= [0, 0, 0, m\omega^2, 0, 0, 0, 0, 0]^T . \end{aligned} \quad (2.40)$$

Using the Biot-Savart analogue for electrostatics (Eq. 2.31), the exact fields and curvatures that will be generated at a desired trap location by applying 1 V to a given electrode \mathcal{E}_i can be simulated:

$$\vec{\varphi}^{(i)} \equiv [E_x^{(i)}, E_y^{(i)}, E_z^{(i)}, C_{xx}^{(i)}, C_{zz}^{(i)}, C_{xy}^{(i)}, C_{yz}^{(i)}, C_{xz}^{(i)}, C_{xxx}^{(i)}]^T \quad (2.41)$$

The relationship between $\vec{\varphi}^{(0)}$, $\vec{\varphi}^{(i)}$ is a linear equation:

$$\vec{\varphi}^{(0)} = [\vec{\varphi}^{(1)} \ \vec{\varphi}^{(2)} \ \dots \ \vec{\varphi}^{(n)}] \vec{V} \quad (2.42)$$

where \vec{V} is an n -dimensional vector describing the voltage that must be applied to each of the n electrodes to generate the fields and curvatures described by $\vec{\varphi}^{(0)}$. If there are exactly 9 electrodes then there will be at most a single set of voltages that will produce a potential that fulfilling all nine constraints. Typically surface traps have many more than 9 electrodes, and there is substantial space to explore to find an optimal solution.

Optimization is a well-known field in computer science, and there are multiple built in computer packages which use well-established optimization algorithms to solve linear equations of this form. I have used `cvxopt` with Python, but this is largely a choice of personal coding preference. Similar packages exist which are compatible with MATLAB, C++, etc., and they generally provide the same algorithms.

More interesting is the choice of cost function. An optimization algorithm requires a cost function, which ranks different solutions \vec{V} and returns the best. A solution that requires a voltage of 100 V on each electrode to trap an ion, while it may provide a nice solution on a computer is impractical, and should be rejected from the solver.

A few cost functions we might imagine:

- *Well depth*: The well should be sufficiently deep to trap ions for long periods of time.
- *Harmonicity*: The well should be as harmonic as possible. This could be achieved by either a cost function that penalizes anharmonicities in the potential or by introducing more higher order terms into our linear equation.
- *Electrode contribution strength*: every voltage should be significant, i.e. a large voltage that has little effect on the potential seen at the ion voltage should be penalized.
- *Voltage smoothness*: neighboring electrodes should have similar voltages, to minimize capacitive coupling between electrodes.
- *Transport smoothness*: for operations like transport or separation that involve time-varying voltages, the voltage vectors $V(\vec{t})$ should change as slowly as possible.

In practice, some of these constraints might serve redundant purposes. For example, voltage smoothness and transport smoothness cost functions will often yield similar solutions. Others are at odds: a deep, highly harmonic well will likely have higher voltages and less smooth voltages than we might like. Further, the weighting of different priorities will vary substantially between apparatuses and even between experiments. In a room temperature trap, where the ion is more

susceptible to background gas collisions, the well depth requirement would need to be more stringent, whereas a cryogenic apparatus can trap ions for long times in surface traps with quite low well depths. For the fast transport experiments discussed in Chapter 7, I heavily weighted the transport smoothness requirement.

2.3.2 Shims and curvatures

All ion traps will have stray fields that must be compensated to controllably create desired mode frequencies and orientations. This is especially challenging in the system described in Chapter 4 due to the combination of UV laser frequencies, small ion-electrode distance, and a cryogenic apparatus. With a short ion-electrode distance, the lasers scatter more on the trap and produce more stray charge because photoelectrons are removed by UV light. At cryogenic temperatures, stray charges dissipate more slowly than they would at room temperature.

We thus require a set of compensation voltages to correct for these stray fields. We refer to these compensation voltages as the shims, as we can consider this as if we are “shimming up” the potential and making it “level” at the trap. The goal is a set of shims that affects a single derivative of the potential but leaves all others unchanged. We can use the same basic scheme as is described above, but with a different target state. For example, to find the y shim, we would use the target vector

$$\begin{aligned}\vec{\varphi}^{(0)} &\equiv [E_x, E_y, E_z, C_{xx}, C_{zz}, C_{xy}, C_{yz}, C_{xz}, C_{xxx}]^T \\ &= [0, 1, 0, 0, 0, 0, 0, 0, 0]^T\end{aligned}\tag{2.43}$$

to find a set of voltages that produces an electric field in y and otherwise leaves the derivatives unchanged.

2.4 Manipulation of ion crystals

In order to implement the quantum CCD architecture described in Section 1.1.1, we require the ability to move, split, merge, and reorder ion crystals.

2.4.1 Ion transport

Ions can be transported from location \vec{x}_a to \vec{x}_b by linearly changing the DC voltages from those needed to produce a well at \vec{x}_a to those needed to produce a well at \vec{x}_b . For sufficiently small $|\vec{x}_b - \vec{x}_a|$, the ion will follow the well minimum. More complicated position and velocity profiles like the ones discussed in Chapter 7 can be produced by breaking this profile up into small pieces that can be approximated linearly.

2.4.2 Splitting and merging of ion crystals

The goal in splitting an ion crystal is to shift from a single harmonic well to a double well. This necessarily has higher-than-harmonic terms. Assuming symmetry (i.e. no third-order term) the axial potential can be described to lowest order as

$$V(x) = \alpha x^2 + \beta x^4 \quad (2.44)$$

A purely harmonic well has $\alpha > 0, \beta = 0$, in which case the ions will be separated by a distance $\delta x = (e^2/\pi\epsilon_0\alpha)^{1/3}$. Separation is performed by reducing the harmonic term and applying a strong quartic, which provides confinement even in the absence of a harmonic term. Under a purely quartic potential $\alpha = 0, \beta > 0$, the ions will separate under the Coulomb interaction to a distance of $\delta x = (4e^2/\pi\epsilon_0\beta)^{1/5}$. We can then apply a “separation wedge” by applying a negative harmonic term $\alpha < 0$. Near the ion equilibrium positions, this produces to lowest order two harmonic wells centered at $x_0 = \pm(-\alpha/2\beta)^{1/2}$ with curvature $\alpha' = -4\alpha$. To merge wells, this process is reversed.

2.4.3 Reordering

I will not discuss reordering in this thesis but note that it is a required component of the scalable quantum CCD. We need to be able to reorder ion crystals in order to bring together arbitrary pairs of ions for two-qubit gates. There are two ways in which we could consider reordering ions: rotating the ion crystal or using a junction. The crystal can be rotated by applying an xy

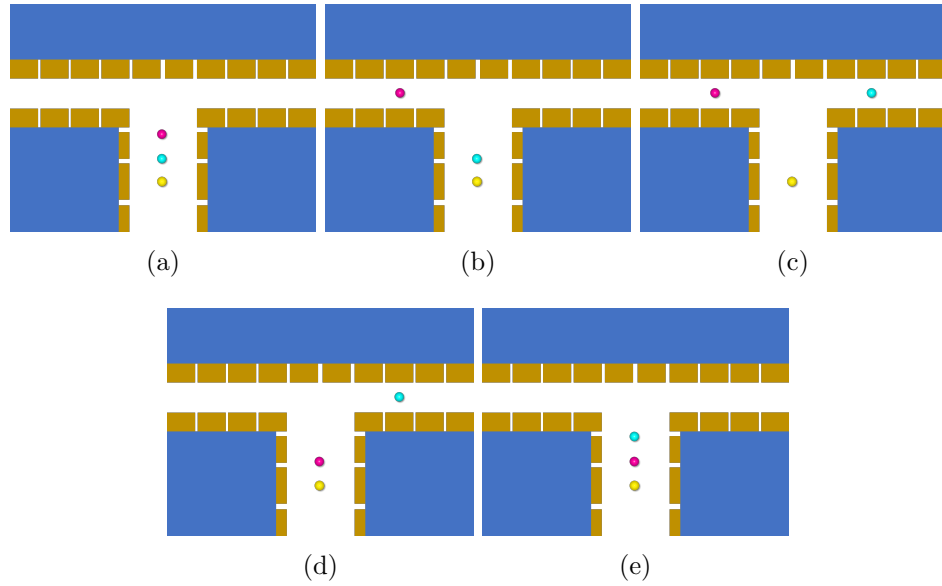


Figure 2.5: Reordering an ion chain using a junction. The colors identify different ions. We can rearrange the orders of an ion in a chain using a junction as shown; however, the junction introduces a “micromotion hump” that can cause ion heating.

shim (or, equivalently, an xz shim):

$$V_{xy} = \frac{\partial^2 V}{\partial x \partial y} \quad (2.45)$$

generated as I have described above in Section 2.3.2. However, this introduces excess micromotion, since the ions will come off of the line of the RF null as they are rotated. Using a junction, as drawn in Figure 2.5 makes some intuitive sense. However, junctions also introduce excess micromotion. Two intersecting RF null lines cannot be created without violating Laplace, so passing through a junction necessarily means the ion moves off of the RF null. Finding a fast, reliable, cold way to reorder ion crystals is a critical piece of the puzzle to building a scalable quantum computer. However, it is beyond the scope of this thesis.

Chapter 3

The ${}^9\text{Be}^+$ Qubit

I have up till now referred to the qubit states in the abstract as $|0\rangle$ and $|1\rangle$ or as spin states $|\uparrow\rangle$ and $|\downarrow\rangle$. Any real atomic system is substantially more complex than this simple two-level approximation and has many more states to consider. However, with careful state preparation and manipulation, we can operate in a limited subspace where most of the additional energy levels can be neglected.

Singly-ionized alkaline earth metals make good choices of qubit ion, since they have a single valence electron and can be understood in a relatively straightforward analysis. In ions with hyperfine structure, the qubit can be encoded in hyperfine ground states, which have very long radiative lifetimes. There are multiple options, including ${}^{25}\text{Mg}^+$, ${}^{43}\text{Ca}^+$, ${}^{133}\text{Ba}^+$, and ${}^{171}\text{Yb}^+$, all of which have been investigated for quantum computing applications [60][74][75][76]. I will discuss ${}^9\text{Be}^+$, which has the lowest mass of any alkaline earth metal. This low mass increases coupling strength between neighboring ions and allows trapping with reduced RF amplitudes.

3.1 Beryllium photoionization

We produce ${}^9\text{Be}^+$ ions by photoionization of neutral beryllium using a 235nm laser. The relevant level structure for neutral beryllium is shown in Figure 3.1. We photoionize in a two-photon process, first exciting the valence electron from $2s^2\ ^1S_0$ to $2s2p\ ^1P_1$ and then exciting it to the continuum. The transition wavelength from $2s^2\ ^1S_0 \leftrightarrow 2s2p\ ^1P_1$ is 234.9329 nm, and the transition wavelength from $2s2p\ ^1P_1$ to the continuum is 306.492 nm [77]. Since the second step is

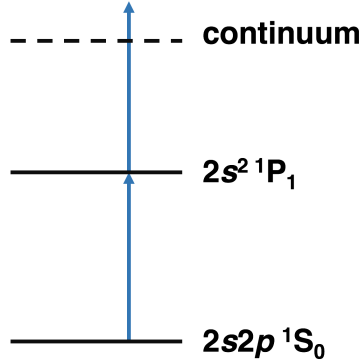


Figure 3.1: Relevant level structure for beryllium photoionization. Photoionization is a two-step process, first exciting the neutral atom in a resonant step from 1S_0 to 1P_1 , then exciting the ion to the continuum.

lower energy, the electron can be excited first to $2s2p$ and then to the continuum by two photons resonant with the first transition. Beryllium can also be ionized by electron beam bombardment. However, photoionization is species-selective and does not generate as much stray charge.

3.2 Level Structure

In order to reasonably approximate the energy levels of an ion as a two-level system, both states must be long lived and initializable with high fidelity. A good choice in $^9\text{Be}^+$ is to encode the qubit between the ground state hyperfine sublevels, as shown in Figure 3.2. For this treatment, we need only consider ions in three energy levels (described by the standard term symbols $^{2S+1}L_J$): $1s^22s\ ^2S_{1/2}$, $1s^22s\ ^2P_{1/2}$, and $1s^22s\ ^2P_{3/2}$. All of these states have orbital configuration $1s^22s$, so I will drop it from my notation hereafter. The quantum numbers S , L , and J describe the electron spin, electron orbital angular momentum, and total electron angular momentum, respectively. All other energy levels are split from the ground state $1s^22s$ by inaccessible laser wavelengths. The nearest orbital transition, $1s^22s \leftrightarrow 1s^23s$ is at 113 nm [78].

All of the relevant energy structure for $^9\text{Be}^+$ is the fine and hyperfine structure and Zeeman splitting. $^9\text{Be}^+$ interacts with a magnetic field by

$$H = -\vec{\mu} \cdot \mathbf{B} \quad (3.1)$$

State	F	g_F
$^2S_{1/2}$	1	-1/2
	2	1/2
$^2P_{1/2}$	1	-1/6
	2	1/6
$^2P_{3/2}$	0	N/A
	1	2/3
	2	2/3
	3	2/3

Table 3.1: Lande g -factors for the ground and first excited states of $^9\text{Be}^+$

[66]. There are three sources of magnetic field we must consider: (1) spin-orbit coupling between the valence electron's spin and orbital angular momentum, (2) spin-spin coupling between the valence electron and the nucleus, and (3) coupling between the valence electron and an externally applied magnetic field \mathbf{B}_{app} . The spin-orbit coupling gives the fine structure splitting between the two P -state levels, measured as 197.2 GHz [79].

The spin-spin interaction gives the hyperfine structure, splitting the two qubit states. Since $^9\text{Be}^+$ has a non-zero nuclear spin $I = 3/2$, the nuclear spin \mathbf{I} will interact with the electronic angular momentum \mathbf{J} . In the low field limit, this spin-spin interaction can be straightforwardly diagonalized by considering eigenstates of the *total* angular momentum

$$\mathbf{F} = \mathbf{J} + \mathbf{I} . \quad (3.2)$$

This spin-spin interaction introduces a hyperfine shift, dependent on the total angular momentum quantum number F . In the ground state, the hyperfine shift splits the $F = 1$ and $F = 2$ levels by 1250 MHz [81]. In the $2P_{1/2}$ state, the hyperfine shift splits $F = 1$ and $F = 2$ by 237 MHz [79], and in the $2P_{3/2}$ state, the hyperfine shift splits the $F = 0..3$ levels by less than 1 MHz [82].

In the low field, the ion will experience a Zeeman shift dependent on F and m_F , the projection of the total angular momentum onto a quantization axis set by the external magnetic field \mathbf{B}_{app} . This energy shift is given by

$$E_Z = -g_F \mu_B m_F B \quad (3.3)$$




Figure 3.2: Ground and first excited state energy levels of the ${}^9\text{Be}^+$ ion in a low magnetic field, adapted from [80]. We use the ${}^2S_{1/2}|F=2, m_F=-2\rangle \leftrightarrow {}^2P_{3/2}|F=3, m_F=-3\rangle$ transition for optical pumping, Doppler cooling, and state detection, plus an additional beam far detuned from this transition for Doppler cooling of very excited states; these are drawn in blue and labeled “blue doppler” (BD) and “blue doppler detuned” (BDD), respectively. Assuming that these beams have pure σ^- polarization, this is a cycling transition. To limit state preparation errors, we also repump on the ${}^2S_{1/2}|F=2, m_F=-1\rangle \leftrightarrow {}^2P_{1/2}|F=2, m_F=-2\rangle$ and ${}^2S_{1/2}|F=1, m_F=-1\rangle \leftrightarrow {}^2P_{1/2}|F=2, m_F=-2\rangle$ transitions, drawn in red.

[66]. The Lande g -factor g_F is dependent on the specific F state as

$$g_F = g_J \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} \quad (3.4)$$

where

$$g_J = \frac{3}{2} + \frac{S(S+1) - L(L+1)}{2J(J+1)} \quad (3.5)$$

[83]. Values of g_F for the relevant states are given in Table 3.1.

The above is true in the low-field limit. It is possible to substantially extend the coherence time between our two qubit states by choosing a magnetic field at which the energy splitting between the two qubit states is first-order independent of magnetic field. This requires a magnetic field high enough that the Zeeman effect depends quadratically on the external magnetic field B_{app} . Coherence times as long as 10 seconds are possible in ${}^9\text{Be}^+$ by operating at such a field [18]. The lowest magnetic field at which this first-order magnetic field insensitivity can be achieved in ${}^9\text{Be}^+$ is 11.94 mT. Details of the high-field energy level configuration of ${}^9\text{Be}^+$ can be found in [80]. Producing such a magnetic field in the lab requires currents on the order of 50–100 A, depending on the magnetic field coil configuration, and presents a substantial increase in experimental complexity. Therefore, I will limit my discussion to the low-field limit, in which we apply only a small magnetic field to break the Zeeman sublevel degeneracy.

States in the ${}^2S_{1/2}$ manifold make good qubit states, since the $F = 1$ states have very long radiative lifetimes, estimated as 5×10^{14} seconds, or about 16 million years [1]. We could choose any pair of states within the ground state manifold as qubit states. For all the work discussed in this thesis, I used the qubit states

$$\begin{aligned} |\downarrow\rangle &= {}^2S_{1/2}|F = 2, m_F = -2\rangle \\ |\uparrow\rangle &= {}^2S_{1/2}|F = 1, m_F = -1\rangle \end{aligned} \quad (3.6)$$

primarily chosen for ease of state preparation. We apply an external DC magnetic field of 5.7 G, under which these qubit levels are separated by 1262 MHz.

3.3 State preparation and control

3.3.1 Doppler cooling

In the experiments discussed in this thesis, we Doppler cool on the ${}^2S_{1/2}|F = 2, m_F = -2\rangle \leftrightarrow {}^2P_{3/2}|F = 3, m_F = -3\rangle$ cycling transition, nicknamed the “stretch transition.” This transition is drawn in Figure 3.2, labelled as the “blue Doppler,” or BD. (The “blue” in “blue Doppler” is to distinguish it from the repumping beam, nicknamed the “red Doppler” for historical reasons. The BD is still detuned red of resonance for Doppler cooling.) During cooling cycles, this beam is detuned 6 MHz to the red of resonance. We also apply a far-detuned Doppler beam tuned ~ 620 MHz to the red of resonance with approximately 8x higher intensity, which serves to cool very hot ions. This laser is nicknamed “blue Doppler detuned,” or BDD. With this combination of far-detuned and near-resonant Doppler cooling, we are able to cool ions in the i th motional mode to near the Doppler limit [84]

$$k_B T |\hat{x}_i \cdot \hat{k}_{Dop}|^2 = \frac{\hbar \gamma}{2} \quad (3.7)$$

where I have modified the usual equation to account for overlap between the mode orientation \hat{x}_i and the cooling laser \hat{k}_{Dop} . The transition linewidth $\gamma = 2\pi \times 19.4$ MHz in angular frequency units [82]. This cools the ion to $\bar{n} \approx 9$ in a 2 MHz mode oriented 45° from the Doppler cooling laser. The ion can be cooled below the Doppler limit by sideband cooling, discussed in Section 3.4.2.

3.3.2 Optical pumping

In order to use the ion as a qubit, it is necessary to controllably produce ions in a single qubit state. The excited state lifetime of the 2P state is 8.2 ns [82], so in a very short time the ion begins with all population in the ${}^2S_{1/2}$ manifold. This population can be shifted into the $|\downarrow\rangle$ qubit state by optically pumping on the ${}^2S_{1/2} \leftrightarrow {}^2P_{3/2}$ transition using a σ^- polarized laser at 313.13 nm. We use the same laser as that used for Doppler cooling, the “BD,” for this process. This laser is tuned to near-resonance with the ${}^2S_{1/2}|F = 2, m_F = -2\rangle \leftrightarrow {}^2P_{3/2}$ transition. $|F = 2, m_F = -2\rangle$ is approximately 1.25 GHz detuned from the $F = 1$ manifold, so we also use the far-detuned laser

to clear out the remainder of the $^2S_{1/2}$ manifold. The BDD is detuned ~ 620 MHz detuned to the red of the $^2S_{1/2}|F = 2, m_F = -2\rangle \leftrightarrow ^2P_{3/2}$ transition, making it approximately equally detuned from the $F = 2$ and $F = 1$ manifolds. Although this high detuning gives a lower probability-per-photon of making a transition than the near-resonant optical pumping laser, this beam has a higher intensity, and effectively optically pumps the full $^2S_{1/2}$ manifold. In an ideal case, this would produce 100% population in the goal qubit state of $|F = 2, m_F = -2\rangle$, since the population will eventually be transferred to the $^2P_{3/2}|F = 3, m_F = -3\rangle$ state. Decay to anything other than the target qubit state is dipole forbidden, so this state should be a perfect cycling transition.

However, this is an experimental thesis, and experiments are sometimes frustratingly non-ideal. The primary limitation on our state preparation fidelity is imperfect σ^- polarization. Residual π polarization in the optical pumping beam will excite the qubit state to $^2P_{3/2}|m_F = -2\rangle$. This will decay to one of three dipole-allowed hyperfine states in the $^2S_{1/2}$ manifold. One of these is the target state $|F = 2, m_F = -2\rangle$, in which case no infidelity is introduced. However, it can also fall to either $|F = 2, m_F = -1\rangle$ or $|F = 1, m_F = -1\rangle$. In order to clear out these states, we apply two repumping beams, resonant with the $|F = 2, m_F = -1\rangle \rightarrow ^2P_{1/2}$ and $|F = 1, m_F = -1\rangle \rightarrow ^2P_{1/2}$ transitions. The $^2P_{1/2}$ then decays to either one of the same two states or to our target qubit state, again due to selection rules. I will refer to this laser as the “repumping” laser or, for historical reasons, the “RD.” These transitions are shown on Figure 3.2 in red. This transition is also used as the dissipative step in sideband cooling, discussed in 3.4.2.

I will briefly and qualitatively note another issue limiting state preparation fidelity. Our magnetic field must be well-aligned with the optical pumping beam in order to generate pure σ^- light. This can be straightforwardly achieved with our magnetic field coil configuration, as shown in Section 4.2.7. However, there is an oscillating component of the magnetic field caused by the trap itself which causes mixing between the Zeeman sublevels. This effect is negligible in large three-dimensional Paul traps but is a larger effect in surface electrode traps due to their reduced symmetry. Still, this is a minor effect at the level of fidelity needed for the experiments discussed in this thesis.

3.3.3 Microwave transitions

We require a method of reliably transferring population between our two qubit states $|\uparrow\rangle$ and $|\downarrow\rangle$. This can be achieved with global microwave fields. Consider a ${}^9\text{Be}^+$ ion under an applied microwave field of amplitude \vec{B}_0 , angular frequency ω , detuned from our qubit frequency ω_{qu} by $\delta = \omega - \omega_{qu}$. The interaction Hamiltonian of the system is then (neglecting all levels besides $|\uparrow\rangle$ and $|\downarrow\rangle$)

$$\hat{H}_I(t) = \mu\vec{r} \cdot \vec{B}_0 \cos(\omega t) \quad (3.8)$$

By applying an appropriate rotating wave approximation, an ion that begins in the ground state $|\downarrow\rangle$ has a probability of being found in the upper qubit state after time t of

$$P_{\uparrow} = \frac{\Omega^2}{\Omega^2 + 4\delta^2} \sin^2\left(t\sqrt{\Omega^2 + 4\delta^2}\right) \quad (3.9)$$

[85] where the Rabi frequency Ω

$$\Omega = \frac{1}{2\hbar} \langle \downarrow | \mu\vec{r} \cdot \vec{B}_0 | \uparrow \rangle \quad (3.10)$$

is set by both the strength of the driving field and the strength of the transition between $|\downarrow\rangle$, $|\uparrow\rangle$. I will note a historical divergence between the trapped ion community and the broader atomic physics community in the definition of the Rabi frequency. Following the treatment in [1], the trapped ion community has defined Ω as I described above in Equation 3.10; much of the neutral atomic physics community instead defines the Rabi frequency as $\Omega' = \Omega/2$. The dynamics are unchanged by this definition, but it is worth keeping in mind when directly comparing Rabi frequencies.

In the particular case of $\delta = 0$, i.e. for a microwave field resonant with the qubit frequency $\omega = \omega_{qu}$, Equation 3.9 reduces to

$$P_{\uparrow} = \sin^2(\Omega t) \quad (3.11)$$

and, by normalization,

$$P_{\downarrow} = \cos^2(\Omega t) \quad (3.12)$$

Neglecting global phase, we can write the ion state with no loss of generality as

$$|\psi\rangle = \cos(\Omega t) |\downarrow\rangle + e^{i\phi} \sin(\Omega t) |\uparrow\rangle \quad (3.13)$$

for some relative phase ϕ . Notice that this is equivalent to an arbitrary rotation on the Bloch sphere parameterized by ϕ and Ωt . This rotation is a critical ingredient to quantum computing, which require single qubit rotations for computational completeness [26].

Due to their frequency of use, I will note two specific cases:

- The **π -pulse**: After applying a resonant microwave field for $t_\pi = \pi/(2\Omega)$, all of the population has swapped from the initial state $|\downarrow\rangle$ to $|\uparrow\rangle$. This is useful for transferring population coherently between target states.
- The **$\pi/2$ -pulse**: After applying a resonant microwave field for $t_{\pi/2} = \pi/(4\Omega)$, half of the population has swapped from the initial state $|\downarrow\rangle$ to $|\uparrow\rangle$, so the ion is in an equal superposition of $|\downarrow\rangle$ and $|\uparrow\rangle$, with the phase relationship defined by the phase of the magnetic field.

3.3.4 Stimulated Raman transitions

The previous treatment assumed that the transition from $|\downarrow\rangle$ to $|\uparrow\rangle$ is driven by a single frequency field. Because this field must oscillate at a frequency at or near the transition frequency ω_{qu} , the transition is mediated by a microwave photon. However, we would like to be able to use lasers to make these transitions. Lasers can be focused to tight spots, enabling single-qubit addressing. Additionally, to excite the ion's motion (which is necessary for two-qubit entangling gates), we require a large field gradient. This can be challenging to produce in microwaves, but comes for free in short wavelength light. Although a single optical or ultraviolet photon cannot be used to make the transition between our two qubit states, separated by ~ 1.25 GHz, *two* photons can make this transition.

For this, we use stimulated Raman transitions. We wish to make a two-photon transition between our two qubit states $|\downarrow\rangle$ and $|\uparrow\rangle$, separated by angular frequency ω_{qu} . In this treatment, we must also consider a third, auxiliary state $|e\rangle$, far separated from the two qubit transitions so that $\omega_{e,\uparrow}, \omega_{e,\downarrow} \gg \omega_{qu}$. The interaction Hamiltonian (Equation 3.8) can be modified to consider an




Figure 3.3: Raman transitions used in this experiment, adapted from [80]. We make two-photon transitions between our $| \downarrow \rangle$ and $| \uparrow \rangle$ states using light detuned by 80 GHz from the ${}^2P_{1/2}$ excited state. Note that this diagram of the ${}^9\text{Be}^+$ energy levels is not to-scale; although the Raman detuning and the ${}^2P_{1/2}$ hyperfine splitting look comparable as drawn, they are two orders of magnitude different.

ion interacting with an electromagnetic field containing two frequency components at $\omega_{1,2} \gg \omega_{qu}$:

$$\hat{H}_I(t) = q\mathbf{r} \cdot \left(\vec{E}_1 \cos(\omega_1 t) + \vec{E}_2 \cos(\omega_2 t) \right) \quad (3.14)$$

I will specifically discuss the case $\omega_1 = \omega_{e,\downarrow} + \Delta$, $\omega_2 = \omega_{e,\uparrow} + \Delta$: i.e. the first photon ω_1 is detuned from the $|\downarrow\rangle \leftrightarrow |e\rangle$ transition by Δ , and the second photon ω_2 is detuned from the $|\uparrow\rangle \leftrightarrow |e\rangle$ transition by the same Δ . In this case, ω_2 is detuned from ω_1 by the qubit frequency ω_{qu} . This is drawn for the specific case of ${}^9\text{Be}^+$ with the detunings used in this experiment in Figure 3.3.

Following [85], it can be shown that this acts as a single electric field oscillating at $|\omega_1 - \omega_2|$, with a modified Rabi frequency

$$\Omega_{eff} = \frac{\Omega_{\uparrow,e}\Omega_{e,\downarrow}}{2\Delta} = \frac{\langle \uparrow | q\vec{r} \cdot \vec{E}_2 | e \rangle \langle e | q\vec{r} \cdot \vec{E}_1 | \downarrow \rangle}{4\hbar^2 \Delta} \quad (3.15)$$

From here out the treatment in Section 3.3.3 is identical, with Ω_{eff} replacing Ω . Because this is a two photon transition, the probability of it occurring is reduced relative to a single photon dipole-allowed transition. This is reflected in the quadratic Rabi rate dependence on $\langle \psi | q\vec{r} \cdot \vec{E} | \psi \rangle$ terms. However, we can create large electric fields using tightly focused laser beams, so we can produce comparable π -times with both single-photon microwave and two-photon Raman transitions.

The ${}^9\text{Be}^+$ level structure is shown again in Figure 3.3 with the two Raman laser frequencies indicated, shown configured for a resonant $|\downarrow\rangle \leftrightarrow |\uparrow\rangle$ qubit transition. We use the ${}^2P_{1/2}$ excited state as our auxiliary state $|e\rangle$; this state is separated from the ${}^2S_{1/2}$ manifold by 950 THz, so the requirement that the Raman laser frequencies be much greater than the transition frequency $\omega_{qu} = 1.25$ GHz is well-satisfied. The transition $|F = 1, m_F = -1\rangle \rightarrow |F = 2, m_F = -2\rangle$ has a net angular momentum change of -1, so one beam is σ^- polarized and the other is π polarized.

3.4 Ion motional control

3.4.1 Sideband interactions

The other ingredient needed for quantum computational completeness is two-qubit entangling gates [26]. Two qubit gates are generally produced by coupling the ion's qubit state to the ion

motion. This requires a field *gradient*, so that the ion feels a state dependent force. The force felt by the ion should depend on the ion's qubit state, meaning that the field should vary over the ground state extent of the ion's secular motion. The wavelength of a 1261.8 MHz photon in vacuum is 0.24 m, 4×10^6 times the ground state wavefunction extent of ${}^9\text{Be}^+$ in a 2.0 MHz trap. It is thus challenging to create a substantive magnetic field gradient over this zero-point extent using microwaves. Challenging is not impossible, and I point to the results in [31][17][34] as evidence of what is possible using microwave gradients in ion traps. However, I will limit myself to discussing field gradients produced by UV lasers, which have wavelengths much more comparable in scale to the ground extent of trapped ions. We specifically use the stimulated Raman interaction, discussed above in Section 3.3.4 and illustrated for ${}^9\text{Be}^+$ in Figure 3.3, to drive ion motion.

As discussed in Chapter 2, an ion in an RF Paul trap can be treated as if it is trapped in a harmonic well

$$\Phi(x, r_1 r_2) = \frac{1}{2}m\omega_x^2 x^2 + \frac{1}{2}m\omega_{r_1}^2 r_1^2 + \frac{1}{2}m\omega_{r_2}^2 r_2^2 . \quad (3.16)$$

For simplicity, I will discuss a one-dimensional harmonic oscillator with trap frequency ω . This discussion is directly extendable to the three-dimensional case, since the potential is a linear sum of 1D harmonic oscillators.

Using Raman transitions, it is possible to transfer population both between our two qubit states and between motional states, as is shown in Figure 3.4. Instead of the specific case of ${}^9\text{Be}^+$, I have here shown the general trapped ion case, in which we have a two-level qubit system with a third, auxiliary level $|e\rangle$ enabling Raman transitions. Since the ion is trapped in a harmonic well, we also must consider the harmonic oscillator energy level ladder. If an ion in the lower qubit state $|\downarrow\rangle$ and in motional state $|n\rangle$ is illuminated with two lasers, detuned from the auxiliary level $|e\rangle$ by Δ and from each other by exactly $\omega = \omega_{qu}$, the difference frequency between our two qubit states, a coherent transition will be stimulated between $|\downarrow\rangle$ and $|\uparrow\rangle$, with the motional state $|n\rangle$ unchanged. This class of transition is denoted “carrier” transitions. If the two lasers are instead detuned from each other by $\omega = \omega_{qu} - \omega_x$, this transition is no longer resonant; however, a “red




Figure 3.4: Sample stimulated Raman transitions from $|\downarrow\rangle$ to $|\uparrow\rangle$, showing a carrier transition (black/solid), a red sideband (red/dashed), and a blue sideband (blue/dotted)

sideband” transition will be resonant, and a coherent transition is stimulated between $|\downarrow\rangle|n\rangle$ and $|\uparrow\rangle|n-1\rangle$. Similarly, if the lasers are detuned by $\omega = \omega_{qu} + \omega_x$, a “blue sideband” transition is resonant, coherently shifting population from $|\downarrow\rangle|n\rangle$ to $|\uparrow\rangle|n+1\rangle$.

This is illustrated in Fig. 3.4, but I will formalize it a bit mathematically. The interaction Hamiltonian for a trapped ion qubit’s interaction with laser light can be written

$$\hat{H} = \hbar\Omega\hat{S}_+e^{-i(\delta t-\phi)}\exp(i\eta[\hat{a}e^{-i\omega_x t} + \hat{a}^\dagger e^{i\omega_x t}]) + h.c. \quad (3.17)$$

where Ω is the Rabi rate from Eq. 3.15; \hat{S}_+ and \hat{S}_- are the raising and lowering operator for the ion’s qubit state; \hat{a} and \hat{a}^\dagger are the annihilation and creation operators, respectively, for the motional state of the ion; $\delta = \omega - \omega_{qu}$ is the detuning from resonance; ω_x is the secular frequency of the ion’s motion; and $\eta \equiv \Delta k x_0$ is the Lamb-Dicke factor, where x_0 is the spatial extent of the ion’s motional state [1].

The Lamb-Dicke factor can be considered a measure comparing the spatial extent of the ion’s wavefunction with the wavelength of the light. After Doppler cooling, the ion is in the Lamb-Dicke regime ($\eta\sqrt{\langle(\hat{a} + \hat{a}^\dagger)^2\rangle} \ll 1$), where the ion is sufficiently cooled that we can consider only first

order terms in our Hamiltonian. In the Lamb-Dicke regime, Eq. 3.17 can be approximated

$$H = \hbar\Omega \left\{ \hat{S}_+ e^{-i(\delta t - \phi)} + \hat{S}_- e^{i(\delta t - \phi)} + i\eta \left(\hat{S}_+ e^{-i(\delta t - \phi)} - \hat{S}_- e^{i(\delta t - \phi)} \right) \left(\hat{a} e^{-i\omega_x t} + \hat{a}^\dagger e^{i\omega_x t} \right) \right\} \quad (3.18)$$

[1]. Apply a rotating wave approximation that drops quickly varying terms, there are three particular detunings of interest:

- $\delta = 0$, i.e. **carrier transitions**: in this case, the Hamiltonian will reduce to

$$\hat{H} = \hbar\Omega \left(S_+ e^{i\phi} + S_- e^{-i\phi} \right) \quad (3.19)$$

Since the only operators in this Hamiltonian are raising/lowering operators on the ion's qubit state, this term enables transitions between the two qubit states. This transition changes the qubit state but leaves the ion's motion unchanged.

- $\delta = -\omega_x$, i.e. **red sidebands**: in this case, the Hamiltonian reduces to

$$\hat{H} = i\hbar\Omega\eta \left(S_+ \hat{a} e^{i\phi} + S_- \hat{a}^\dagger e^{-i\phi} \right) \quad (3.20)$$

This expression contains no pure raising/lowering or creation/annihilation terms. The only terms remaining are proportional to either $S_+ \hat{a}$ or $S_- \hat{a}^\dagger$, i.e. simultaneous flip of the qubit state from $|\downarrow\rangle$ to $|\uparrow\rangle$ and annihilation of a phonon *or* flip from $|\uparrow\rangle$ to $|\downarrow\rangle$ with creation of a phonon.

- $\delta = \omega_x$, i.e. **blue sidebands**: here, the Hamiltonian reduces to

$$\hat{H} = i\hbar\Omega\eta \left(S_+ \hat{a}^\dagger e^{i\phi} + S_- \hat{a} e^{-i\phi} \right) \quad (3.21)$$

and we have terms proportional to $S_+ \hat{a}^\dagger$, which flips the ion from $|\downarrow\rangle$ to $|\uparrow\rangle$ and creates a phonon, and $S_- \hat{a}$, which flips the ion from $|\uparrow\rangle$ to $|\downarrow\rangle$ and annihilates a phonon.

I have presented something of an oversimplification for readability. The Rabi rate Ω actually depends on the initial and final motional states $|n\rangle$ and $|m\rangle$ as

$$\Omega_{n,m} = \Omega |\langle m | e^{i\eta(\hat{a} + \hat{a}^\dagger)} | n \rangle| \quad (3.22)$$

[1]. Accounting for these varying Rabi rates, the probability of finding an ion in $|\uparrow\rangle$ (in any motional state) after applying a blue sideband for time t is given by

$$P_{\uparrow,bsb}(t) = \sum_{n=0}^{\infty} P_n \sin^2(\Omega_{n,n+1}t) \quad (3.23)$$

But since the ground state $|0\rangle$ is insensitive to the red sideband, the probability of finding an ion in $|\uparrow\rangle$ after applying a red sideband for time t is given by

$$P_{\uparrow,rsb}(t) = \sum_{n=1}^{\infty} P_n \sin^2(\Omega_{n,n-1}t) \quad (3.24)$$

[86].

The ability to drive sidebands and control the ion's motion is a powerful tool that can be used for a large number of applications. Ion motional control is key to the most commonly implemented two-qubit gates [87][88]. Sideband pulses have also been used to produce highly non-classical states of motion for quantum-enhanced sensing [89][70]. However, for this work, I will focus on the use of sideband transitions for ground state cooling.

3.4.2 Raman sideband cooling

The idea of sideband cooling arises from observing a fundamental asymmetry in the harmonic oscillator: while motion can be added to any state, it is not possible to remove motion from the ground state. Returning to the graphical picture of Raman sideband interactions, modified in Figure 3.5, consider an ion in the ground state of motion and the lower qubit state $|\psi\rangle = |0\rangle|\downarrow\rangle$. No energy level exists $\omega_{qu} - \omega_x$ detuned from this level. Therefore, when the ion is illuminated with lasers set to this detuning, no transition is driven. Returning to the red sideband interaction Hamiltonian Eq 3.20, noting $\hat{a}|0\rangle = 0$ and $S_-|\downarrow\rangle = 0$,

$$\hat{H}|0\rangle|\downarrow\rangle = i\hbar\Omega\eta \left(S_+ \hat{a} e^{i\phi} + S_- \hat{a}^\dagger e^{-i\phi} \right) |0\rangle|\downarrow\rangle = 0 \quad (3.25)$$

Fundamentally, the goal in ground state cooling is to move from some thermal distribution of states (as is achieved by Doppler cooling) to a pure ground state:

$$\sum_{n=0}^{\infty} \frac{e^{-n\hbar\omega_x/k_B T}}{1 + e^{-\hbar\omega_x/k_B T}} |n\rangle \rightarrow |0\rangle \quad (3.26)$$




Figure 3.5: An illustration of the idea at the heart of Raman sideband cooling: a red sideband on an ion in the ground state of motion has no resonant level to transfer population to.

The red sideband interaction plays a part in this, since it takes $|n\rangle \rightarrow |n-1\rangle$. However, the red sideband is a coherent, entropy-preserving interaction; since the ground state $|0\rangle$ has zero entropy, it cannot be reached with purely coherent interactions. Cooling requires a dissipative step.

This dissipative step is achieved with the repumping interaction. We repump on two transitions, shown on Figure 3.2: ${}^2S_{1/2}|F=1, m_F=-1\rangle \leftrightarrow {}^2P_{1/2}|F=2, m_F=-2\rangle$ and ${}^2S_{1/2}|F=2, m_F=-1\rangle \leftrightarrow {}^2P_{1/2}|F=2, m_F=-2\rangle$. An ion excited to ${}^2P_{1/2}|F=2, m_F=-2\rangle$ will spontaneously decay to either ${}^2S_{1/2}|F=1, m_F=-1\rangle$, ${}^2S_{1/2}|F=2, m_F=-1\rangle$, or to ${}^2S_{1/2}|F=2, m_F=-2\rangle$ (the $|\downarrow\rangle$ qubit state). All of these except $|\downarrow\rangle$ will continue to be cleared out by the repumping interaction, so the population will eventually be transferred to $|\downarrow\rangle$. Two features make this a good choice for use in sideband cooling. One is that in the Lamb-Dicke regime, the ion motional state $|n\rangle$ is unaffected by this interaction. The other is that the interaction is dissipative, so entropy is lost to the environment.

Red Raman sidebands and repumping give all the ingredients needed to cool the ion to its

motional ground state. The ion is cooled to the Doppler limit in the qubit state $|\downarrow\rangle$. (1) We apply a red sideband π -pulse, coherently transferring population from

$$|n\rangle|\uparrow\rangle \rightarrow |n-1\rangle|\uparrow\rangle \quad (3.27)$$

for $n \neq 0$, neglecting non-idealities. Any population in $|0\rangle|\downarrow\rangle$ will be insensitive to this interaction.

(2) Then we apply the dissipative step, which in the ideal case drives

$$|n\rangle|\uparrow\rangle \rightarrow |n\rangle|\downarrow\rangle \quad (3.28)$$

After many repetitions of (1) and (2), this process will cool the ion asymptotically close to the ground state $|0\rangle|\downarrow\rangle$ [1].

Chapter 4

Apparatus

In this chapter, I will present the work that has taken the bulk of my time during my thesis: the development of a robust, flexible, low-vibration cryogenic apparatus for trapped ion quantum information experiments. Operation at cryogenic temperatures enables excellent vacuum and reduced electric field noise at the ion. This apparatus also has a sufficiently low base temperature for the operation of integrated superconducting devices. Although in this work only a few trap designs were used (presented in Section 4.2.1), we anticipate that it will be used for many trap prototyping-type experiments.

4.1 The cryostat

There are multiple reasons to operate trapped ion experiments cryogenically, chiefly, improved vacuum, reduced anomalous heating, and potentially reduced overhead for testing new traps. There are two classes of cryostats that provide the relatively high cooling powers (> 1 W) needed for trapped ion experiments at around 4 K.

- *Wet cryostats*: In a wet cryostat, the system is cooled by contact with a cryogenic liquid. For 4 K operation, this must be liquid helium, since no other cryogen reaches such a low temperature. This category includes bath cryostats (which are cooled by thermal contact with a large bath of liquid helium) and flow cryostats (in which liquid helium continuously flows through the system). The primary challenge of a wet cryostat is cost. During continuous operation, the bath cryostat used in [47] requires helium fills approximately

twice weekly, with approximately 50 L of liquid helium per fill. Helium is not a renewable resource, and as the global helium supply dwindles, access to liquid helium may become tenuous for research groups without access to a helium liquification facility. However, wet cryostats can have reduced vibration relative to closed-cycle systems, which may make them ideal for sensitive or high-fidelity experiments. Apparatuses for trapped ion experiments cooled by bath cryostats are described in [90], [91], [47]. A flow cryostat system is described in [92].

- *Closed-cycle cryostats*: An alternative to the wet cryostat is a closed-cycle cryocooler. In such a system, the cooler does work on a closed volume of helium in order to cool it to cryogenic temperatures. Cryocoolers are typically based on either a Gifford-McMahon or a pulse tube cycle. The Gifford-McMahon cycle involves mechanically moving parts and thus introduces large-amplitude vibration into the system. Pulse tubes have lower amplitude vibration at higher acoustic frequencies but are typically more expensive. Pulse tubes have reduced vibration relative to Gifford-McMahon systems, but vibrate with amplitudes of multiple μm [93], still an inappropriate level for atomic physics experiments. Both Gifford McMahon and pulse tube cryocoolers can be operated in a low-vibration configuration like the one described below. Ion trapping apparatuses incorporating both pulse tubes [94][95] and Gifford-McMahon [96][25] cryocoolers have been demonstrated.

The experiments discussed in this thesis were performed in a two-stage Gifford McMahon cryocooler in a low vibration configuration (Janis SHI-4XG-15-UHV) with a specified cooling power of 1.5 W at 4 K.

4.1.1 Vibration

The Gifford-McMahon cycle relies on multiple moving mechanical parts [93][97]. This causes vibration with 10's of μm amplitude [93]. This is large on the scale of the laser beam diameters, which are focused to approximately 30 μm beam diameters.




Figure 4.1: Low-vibration cryocooler configuration. The cold finger, which vibrates with 10's of μm amplitude, is mounted on a steel frame which sits on the laboratory floor. The low-vibration experiment zone is mounted on a (non-floating) optical table. A thermal link between the two subsystems is made by filling the exchange space (pink/blue) with helium gas. The two subsystems are only directly mechanically linked by flexible rubber bellows.

Because of this large amplitude vibration, the cryocooler used in these experiments is operated in a low-vibration configuration, illustrated schematically in Figure 4.1. In this configuration, the cryogenic system is broken up into two zones, mechanically decoupled from each other: the cryocooler cold finger, which has vibration on the order of 10's of μm s, and the experiment zone. The experiment zone is mounted on a 4' \times 4' optical table, while the cold finger is mounted on a heavy steel frame that stands on the laboratory floor instead of being directly mounted on the optical table. The only direct mechanical connection between the two subsystems is made by flexible rubber bellows. The space between the experiment zone and the cryocooler is filled by helium gas at a pressure of 0.25 psi over atmosphere. This helium gas makes a strong thermal link but a weak mechanical link between the cold head and the experiment zone, enabling efficient cooling of the experiment zone without transfer of vibrations.

4.1.1.1 Vibration measurements

The system's vibration can be measured by a knife edge test, illustrated in Figure 4.2. A knife edge was mounted in the cryostat where a trap would be installed. A laser was focused on this knife edge so that the laser transmission was partially obscured. The laser used for this test was an Toptica DL Pro IR diode laser with wavelength 850 nm, chosen for its turnkey operation, reasonably low intensity noise, and availability when this experiment was performed. Assuming a Gaussian beamshape (a good assumption for the output of a single mode fiber), vibrations of the knife edge traverse to the laser beam translate into variations of the transmitted laser intensity by

$$I_{trans} = I_0 + \frac{I_{max}}{2} \operatorname{erf} \left(\frac{\sqrt{2}(x - x_0)}{w_0} \right) \quad (4.1)$$

for background intensity I_0 , full beam intensity I_{max} , and beam waist w_0 [98]. Near the center of the laser beam $x = x_0$, this is approximately linear:

$$\frac{dI_{trans}}{dx} \approx \frac{I_{max}}{\sqrt{2}w_0} . \quad (4.2)$$




Figure 4.2: Top view of the experimental configuration for vibration measurements. We are looking at a slice through the cryostat, with all vacuum and shielding layers drawn. A knife edge is placed at the center of the cryostat where an ion trap would otherwise be located. An 850 nm IR laser is focused onto the knife edge. Transverse position fluctuations of the knife edge can be extracted by measuring intensity variations of the transmitted beam. I measured both transverse directions using this method. I was also able to measure vertical vibration by using the bottom edge of the knife edge.

Time-domain vibration data taken using this scheme are shown in Figure 4.3. “Horizontal” vibration measurements were taken in the plane of the optical table and are likely due to the pendulum motion of the cryostat. “Vertical” vibration measurements were taken normal to the plane of the optical table. Horizontal vibrations dominate in this system.

There are a few characteristics of the cryostat vibrations to note:

- The cryostat’s cooling cycle (with a period of approximately 1 second) generates an approximately 1 Hz envelope in the horizontal data. However, within this envelope, there is

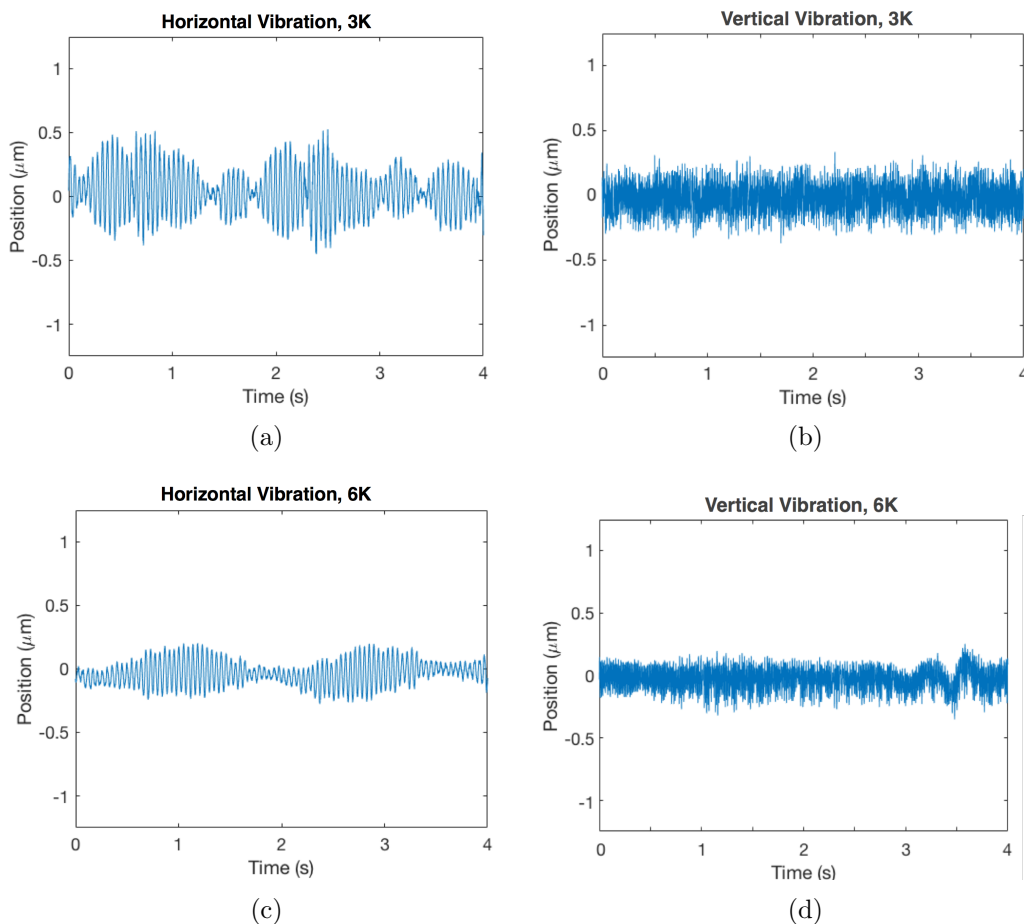


Figure 4.3: Cryostat vibrations in the time domain. “Horizontal” vibration (a, c) is in the plane of the optical table. “Vertical” vibration (b, d) is normal to this plane. (a, c) The horizontal vibration appears to be driven at the approximately 1 Hz from the cooling cycle of the cryostat. Within a 1 Hz envelope, an approximately 45 Hz resonance dominates. (b, d) Vertical vibrations are lower in amplitude than the horizontal vibration. The vibration at 6 K (c, d) is reduced relative to the vibration below the helium condensation temperature (a, b) in both directions.

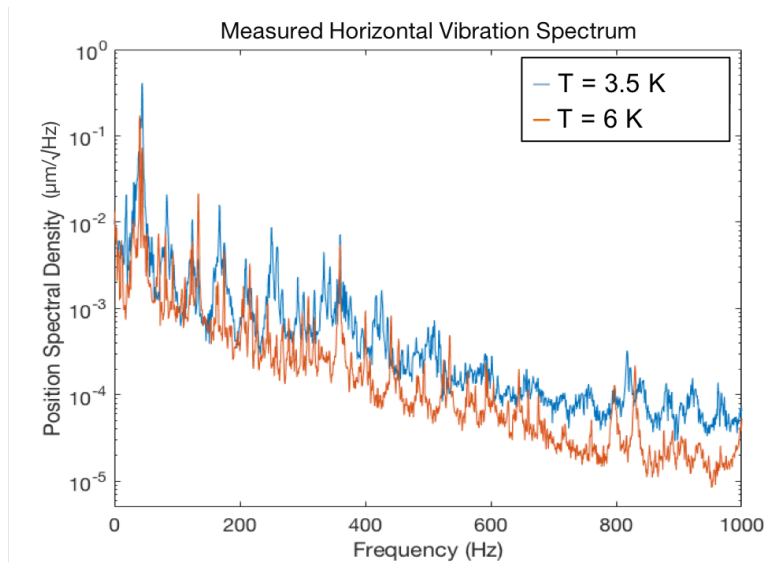


Figure 4.4: Cryostat vibrations in the frequency domain, measured with 1 Hz resolution bandwidth. The 45 Hz vibration dominates. The peak of this spectrum is $0.67 \mu\text{m}/\sqrt{\text{Hz}}$ at 3 K, below the helium condensation point, and $0.18 \mu\text{m}/\sqrt{\text{Hz}}$ at 6 K. All other modes are more than an order of magnitude reduced in amplitude.

a fast oscillation at around 45 Hz. We attribute this to a fundamental pendulum mode of the system that is driven by the impulse of the cold head motor.

- The vertical vibration does not possess a dominant frequency component when viewed the eye and is lower in amplitude than the horizontal oscillation.
- The vibration amplitude is reduced by heating the experiment zone from 3 K to 6 K. This can be attributed to a phase change in the thermal exchange helium. Above 4 K, the helium providing the thermal link between the cryocooler and the experiment zone is gaseous. Below this temperature, it condenses to liquid helium, which transmits more vibration to the experiment zone.

The horizontal vibration was also investigated in frequency space, shown in Figure 4.4. The data shown was measured on a HP89410A spectrum analyzer, with 1 Hz resolution bandwidth. A “position spectral density” was obtained by scaling a voltage spectral density ($\text{V}/\sqrt{\text{Hz}}$) by the same linear scaling factor (Eq. 4.2) used to convert time domain voltage fluctuations into position

fluctuations. This spectrum gives similar qualitative information to the time-domain plots, but enables a better quantitative understanding of frequency content. The 45 Hz frequency mode dominates, with all higher frequency modes (many of which are higher-order harmonics of the dominant 45 Hz mode) more than an order of magnitude smaller. The increase in vibration below 4 K is also visible in this measurement.

4.1.1.2 Estimate of vibration-induced errors

There are at least two major potential sources of error that could be introduced by an experiment vibrating with respect to the lasers: laser intensity errors and laser frequency errors induced by the Doppler shift of an ion moving relative to the laser. I have below roughly estimated the gate errors that might be introduced to a two-qubit gate by these effects, based on the discussions in [60] and [99].

Laser intensity errors can be estimated by assuming the ions are at the center of a perfectly Gaussian beam

$$I(r) = I_0 e^{-2r^2/w_0^2} \quad (4.3)$$

for a beam with peak intensity I_0 and beam radius w_0 [98]. In this apparatus, the lasers are focused to a waist of approximately $w_0 = 15 \mu\text{m}$. With a peak-to-peak vibration amplitude of $1 \mu\text{m}$, as can be seen in the 3.5 K data, this introduces a 0.2% variation in laser intensity. For a geometric phase gate, the gate error caused by Rabi rate error is

$$\epsilon_\Omega \approx 2.5 \times \left(\frac{\delta\Omega}{\Omega} \right)^2 \quad (4.4)$$

where Ω is the Rabi rate and $\delta\Omega$ is the Rabi rate error [60][99]. The Rabi frequency is directly related to the laser amplitude [83], so a 0.2% variation in laser intensity would introduce a 0.1% variation in Rabi rate, giving an estimated Rabi rate error of

$$\epsilon_\Omega \approx 2.5 \times (0.001)^2 = 2.5 \times 10^{-6} , \quad (4.5)$$

small on the scale of the lowest demonstrated laser gate errors to date.

Laser frequency errors would be caused by the Doppler shift of an ion moving within our laser beam. Consider an ion oscillating at 45 Hz with amplitude $0.67 \mu\text{m}$ (a reasonable lowest-order approximation for the cryostat's vibration at 3.5 K, as can be seen on in Figure 4.4). This ion has a peak speed of $3 \times 10^{-5} \text{ m s}^{-1}$. This introduces a Doppler shift of

$$\delta = 2 \frac{v_{ion}}{\lambda_{Ram}} = 192 \text{ Hz} \quad (4.6)$$

relative to the Raman lasers. From [99], the infidelity introduced to a simple two-qubit gate for small detunings is

$$\epsilon_{\delta} = \frac{1}{16} (\delta t_g)^2 \quad (4.7)$$

for frequency detuning δ from the ideal value and gate time t_g . Using such techniques as multi-loop gates [99] and dynamical decoupling [100], it is possible to further reduce sensitivity to detuning errors like this. If the vibrating ion introduces the worst-case value of $\delta = 192 \text{ Hz}$ detuning from the ideal gate detuning and assuming a $30 \mu\text{s}$ gate time (based on the gate time in [16]) this error is $\epsilon_{\delta} = 2 \times 10^{-6}$. It may still be useful to further reduce vibrations before using this apparatus for high-fidelity gate experiments. For all of the proof-of-principle experiments discuss in this thesis, intensity and detuning errors at this level are acceptable.

4.1.1.3 Future improvements

Since the vibration of the cryostat pendulum mode appears to dominate, mechanical bracing of this pendulum should reduce vibration. In previous cryogenic experiments at NIST, discussed in [47], vibration was reduced by tensioning the internal structure. Due to geometrical differences between the two apparatuses, this exact solution used in that apparatus would not be feasible in the closed-cycle system. However, a solution using springs to brace the cold stage to the outer vacuum wall could reduce vibration of the cold head.

Alternatively, the effects of the vibration could be mitigated by addressing the qubit in ways that are not sensitive to vibration. Gates can be performed by an on-chip radiation source, either using microwaves [31] or integrated optics [32]. Since the vibration would then be common mode

	Estimated heat load (mW)
Thermal links in cabling	100
Black-body radiation	47
Dissipated RF drive	3.2-7.9
Total	150 – 155 mW

Table 4.1: Estimated heat loads on the 4 K stage due to various experimental considerations

between the radiation source and the ion, it should not affect the gate fidelity. In one integrated optics experiment in a similar vibration-isolated cryostat, mechanically linking the cold finger to the experiment zone, effectively inducing vibrations on the scale of 10’s of μm at the trap, did not measurably increase qubit errors [101]. At NIST, microwave-driven gates are one of the major directions of investigation, but beyond the scope of this thesis; see [34][69] for more. We are also interested in pursuing tests of integrated optics at 313 nm, using samples identical to those in [32].

4.1.2 Heat load estimation and mitigation

There are three major sources of heat load in cryogenic trapped-ion experiments:

- Thermal flow through wiring
- Black-body radiation
- Heat dissipated due to the trap RF

For the experiments discussed in Chapter 6, the trap must be as cold as possible, so efforts were made to limit these heat loads. These efforts are discussed in detail in Appendix A. The estimated heat loads due to each source are summarized in Table 4.1.

4.1.3 Temperature measurement

There are two temperature sensors in this apparatus, both calibrated Lakeshore DT-670 silicon diodes. One is buried within the cold head itself (denoted T_{cryo}), and the other is mounted on the trap mounting block (denoted T_{trap}). The mounting of the trap temperature sensor T_{trap}

is shown in Figure 4.5(a). The temperature in each sensor is measured by a 4-wire measurement, illustrated schematically in Figure 4.5(b). Trap and cold head temperatures during a standard cooldown from room temperature to cryogenic temperatures are shown in Figure 4.6. At base temperature, the temperature on the trap mounting block is ~ 0.5 K elevated relative to the cold head temperature, but is cold enough for the applications discussed in this thesis.

4.1.4 Vacuum hygiene

One of the nice features of a cryogenic system is the ability to achieve excellent vacuum without going through the typical cleaning and baking procedure that achieving ultra-high vacuum (UHV) at room temperature requires. Room-temperature UHV systems must be rigorously cleaned to remove any oils and baked for days to remove water that has adsorbed onto the surface but which will outgas at high vacuum. Additionally, volatile materials like most plastics and many other synthetic materials cannot be used in room-temperature UHV setups, since they have high vapor pressures [102]. However, nearly every chemical besides helium has an extremely low vapor pressure at 4 K, so vacuum baking and careful material selection are unnecessary to generate UHV at cryogenic temperatures.

For this specific system, a few techniques were adopted from the preparation of UHV systems. Materials in the cryostat were limited to metals, ceramics, and low-outgassing synthetics like PEEK, Kapton, and Stycast 2850FT, and all parts were ultrasonically cleaned in propanol, methanol, and acetone. This is because, although any outgassed material will not have substantial vapor pressure at 4 K, they can adsorb onto cold surfaces. The previous cryogenic apparatus at NIST has been hampered by slowly-building contamination on cold windows. We attribute this contamination to exposure to a room-temperature vacuum space containing oily contaminants. At least one source of this contamination was made worse by exposure to high-intensity UV light. The exact source of this contamination is unknown. The closed-cycle system has been run cold for over 15,000 hours with frequent UV exposure, with no evidence of contamination to date.

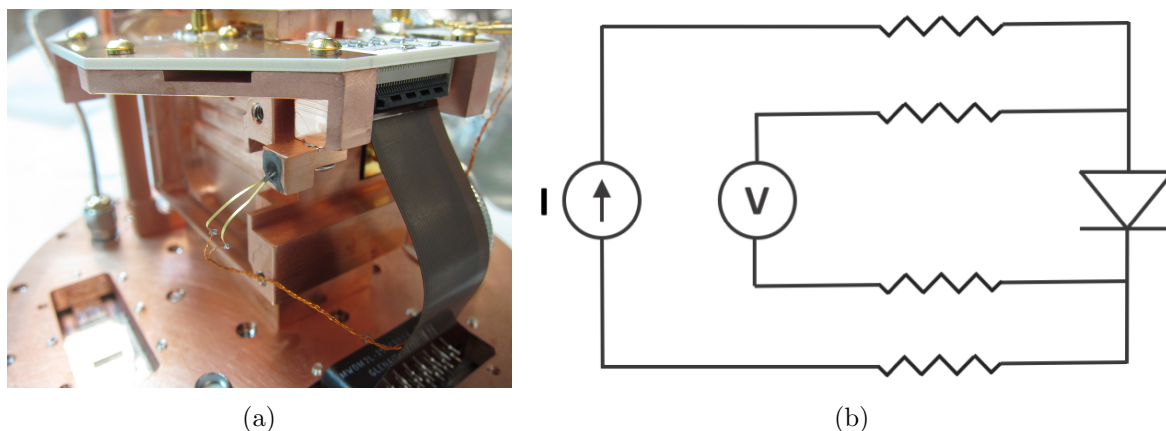


Figure 4.5: (a) Mounting of the “trap” temperature sensor. A silicon diode is potted in Stycast 2850FT in a copper mounting block, which is then mounted on the mounting block holding the trap. The silicon diode has a temperature-depended voltage drop. (b) Schematic for a four-wire measurement. By driving the current and measuring the voltage using different wires, it is possible to decouple voltage drop in the line due to non-negligible line-resistance from voltage drop in the sample.

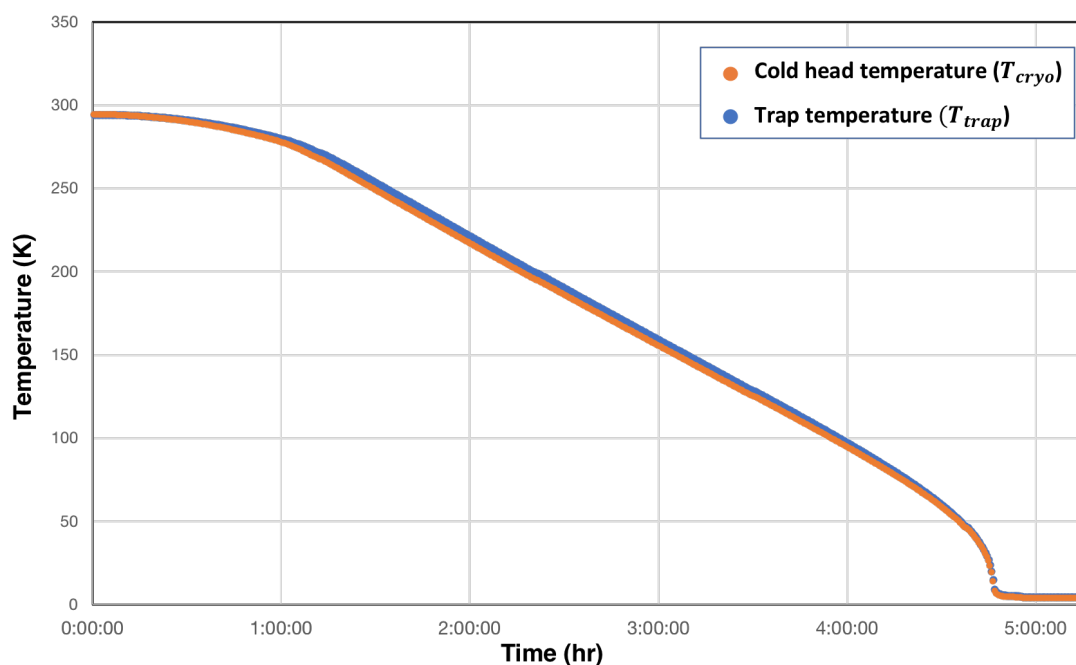


Figure 4.6: Trap and cold head temperatures while the cryostat cools from room temperature to 4 K. Due to finite thermal resistance between the trap and the cold head, T_{trap} lags slightly below T_{cryo} . At base temperature, the trap temperature T_{trap} is ~ 0.45 K elevated relative to the cold head temperature T_{cryo} .

4.2 Ion trapping system

The cryostat described above has low vibration, minimal heat load, and ultra-high vacuum. There is substantial additional experimental hardware needed to trap ions within this cryostat.

4.2.1 Traps

Three ion traps were used for the experiments discussed in this thesis. All are single-layer surface electrode traps fabricated by Daniel Slichter in the Boulder Microfabrication Facility at NIST unless otherwise noted.

4.2.1.1 Point trap

Initial loading tests in this apparatus were performed in a point trap with similar electrode design to that described in [3], but with no integrated detector. A photograph of the trap electrodes is shown in Figure 4.7. In this trap, the ion is trapped 40 μm above the trap surface by RF confinement in all three dimensions. Dust on this trap caused large, time-dependent stray electric fields in the presence of UV lasers, making micromotion minimization challenging, so this trap was not used for experiments requiring precise control of the ion. However, the point trap was useful for initial demonstration of the cryostat. It was also the trap used to demonstrate loading ${}^9\text{Be}^+$ ions with laser ablation, discussed in Section 4.2.8.1.

4.2.1.2 The detector trap

The trap used for the experiments discussed in Chapter 6 was fabricated by Daniel Slichter and Varun Verma and nicknamed the “detector trap.” The fabrication process was similar to what is described in [3]. The trap is fabricated on a 350 μm thick intrinsic silicon substrate, and trapping potentials are provided by voltages on 24 electroplated gold electrodes. The electroplated gold electrodes are 6 μm thick, with 3 – 4 μm gaps between electrodes. This high aspect ratio shields the ion from electric fields generated by any stray charge on the substrate surface.




Figure 4.7: Optical microscopy of the point trap, with false color added to identify electrodes. RF confinement is provided by the red electrode. The blue electrodes were grounded during all work that used this trap.




Figure 4.8: Layout of the detector trap electrodes. RF is applied to the red electrodes, producing a linear RF null $39 \mu\text{m}$ above the trap surface along the central axis of this trap. DC is applied to the blue electrodes to provide axial confinement. The two purple electrodes can be used to provide both further DC confinement (primarily used to rotate the radial modes in the yz plane for better overlap with the cooling beams) and a microwave field. Finally, a superconducting nanowire single photon detector (SNSPD) is integrated into the geometry of this trap. The SNSPD and its leads are colored green and yellow, respectively. Detail of the SNSPD is shown in Figure 4.9




Figure 4.9: (a) Detailed scanning electron microscope image of the trap-integrated SNSPD. False color is added to match Figure 4.8. The ion is trapped $35(1)\mu\text{m}$ above the detector surface for an effective numerical aperture (NA) of 0.32. (b) Detail on the meandered superconducting nanowire on a previous generation of this trap, from [3].

The electrode layout is shown in Figure 4.8. The trap is a “five-wire” trap, as described in Chapter 2. The trap was cut in a “bow-tie” shape, inspired by [72], so that lasers could be focused more tightly with reduced scatter on the trap surface. This bow-tie shape can be seen in Figure 4.14. RF applied to the two RF electrodes (colored red in Figure 4.8) will generate a RF quadrupolar null along a line located approximately $39\mu\text{m}$ above the trap surface along the x-axis. This provides confinement transverse to the x axis. The null is not uniformly $39\mu\text{m}$ from the trap. Due to end effects, it curves about $10\mu\text{m}$ closer to the trap towards the right hand end as oriented in Figure 4.8. Confinement in the axial direction is provided by DC voltages applied to the blue and purple electrodes. The calculation and generation of these DC voltages is described in detail in Chapter 2 and Section 4.2.6.2, respectively. These electrodes can also be used to compensate for stray electric field and to adjust the frequencies and mode orientations of the secular modes. In order to perform global spin flips, either of the central purple electrodes can act as a microwave antenna. The microwave input circuit illustrated is in Figure 4.10. This electrode is not a resonant structure, but acts as an antenna because the ion is much closer to the electrode than a single microwave wavelength ($\approx c/1250\text{ MHz} = 0.24\text{ m}$).

Finally, this trap incorporates a superconducting nanowire single photon detector (SNSPD),

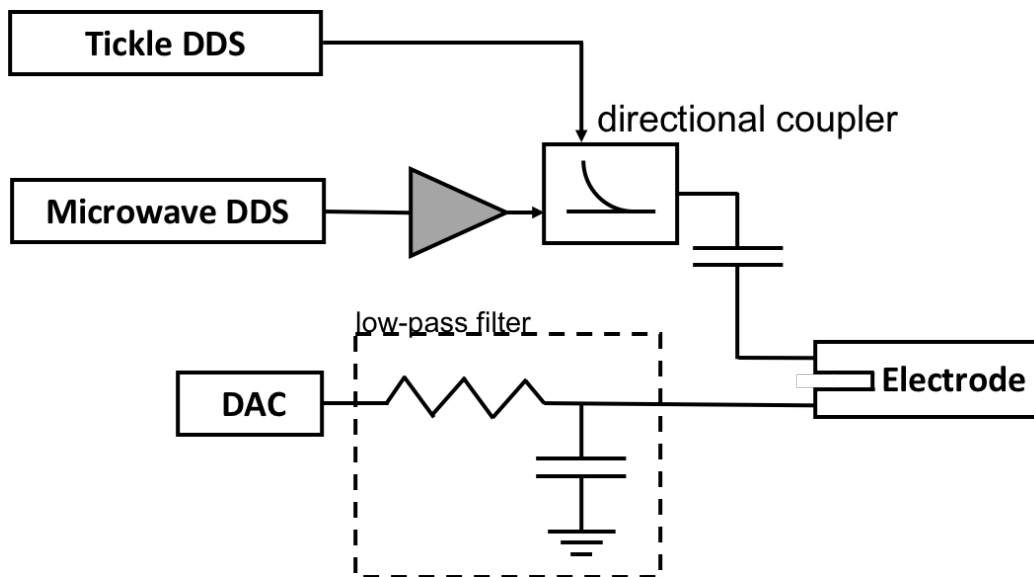


Figure 4.10: The circuit used to couple DC, microwave, and “tickle” signals onto a single electrode. The microwave electrodes are double ended, so the DC line is wirebonded to one end and the microwave to the other end. The DC signal passes through a low-pass filter before being applied to the electrode. The capacitor on the microwave side of the electrode acts as break for DC signals. At microwave frequencies, the capacitors provide minimal impedance, so the oscillating microwave signal is directly applied to the electrode. The other end of the electrode is microwave grounded through the capacitor in the low-pass filter. Applied microwaves will create a large microwave current and thus generate a large microwave magnetic field.

shown in green in Figure 4.8 and in more detail in Figure 4.9. The SNSPD consists of 8 nm thick amorphous $\text{Mo}_{0.75}\text{Si}_{0.25}$, which has $T_C = 5.2$ K. This superconductor is fabricated in a meandered nanowire pattern, shown in Figure 4.9(b), in an active area $20 \times 22 \mu\text{m}$ with nanowires of 110 nm width on a 170 nm pitch. The SNSPD requires a trap temperature $T_{trap} < 3.65$ K to operate.

4.2.1.3 The transport trap

After the detector trap delaminated (discussed in Section 4.2.3) it was replaced with a trap nicknamed the “transport trap.” This trap also is sometimes referred to as the “detectorless” trap because it is an identical geometry to the detector trap but has an evaporated gold electrode replacing the SNSPD. This allows for a simpler trap fabrication process with a higher yield of functional traps. This trap is appropriate for tests that do not require high fidelity on-chip detection.

The experiments discussed in Chapter 7 were performed in this trap.




Figure 4.11: Transport trap electrode detail. This trap is identical to the detector trap (shown and described in Figures 4.8 and 4.9) but the SNSPD is replaced by an evaporated gold electrode. This replacement simplifies fabrication and increases trap yield. There is an apparent color difference from the other electrodes on this picture, because the evaporated gold is 50 nm thick, while the other electrodes are 6 μm thick electroplated gold.

4.2.2 Experiment pillbox

All work involving an exposed trap (trap mounting, wirebonding, connecting cabling, installing ovens, etc.) is performed in a HEPA-filtered clean room to limit the trap's exposure to dust. After this is completed, the trap is enclosed in a copper "pillbox" that can be directly bolted to the cold head without exposing the trap to the lab. This pillbox, closed and mounted on the cold head, is shown in Figure 4.12. The pillbox is vented to the room temperature vacuum chamber and evacuated at the same time as that space. The venting holes do not have line of sight to the ion.

The pillbox also serves as a radiation shield, ensuring that the ion only has line of sight to 4 K surfaces.

4.2.3 Trap mounting

The trap mounting apparatus described below was only used with the detector and transport traps. The point trap was installed in an older generation of the cryostat 4 K mount. This mount had high thermal resistance to the cold head, vibration-sensitive RF coupling, and floppy,




Figure 4.12: The pillbox mounted to the cryostat cold head. This copper container protects the trap from exposure to the lab during installation acts as a 4 K radiation shield. It is vented to the room temperature vacuum chamber with no line of sight from the ion to any room temperature surfaces.

challenging to install DC wiring. The entire 4 K mounting apparatus was redesigned between installation of the point trap and the detector trap, and all of the experimental results discussed in this thesis were performed using the mounting apparatus described.

The trap is glued or bonded onto a 8×5 mm oxygen-free high conductivity (OFHC) copper mounting block. The first-generation detector trap was glued using a thin layer of Epo-tek H21D, which is electrically and thermally conductive and low-outgassing. Epo-tek H21D is also relatively brittle and the trap cracked upon cooling to 4 K, as shown in Figure 4.13. We attribute this to mismatched coefficients of thermal expansion (CTE). Silicon and copper have an order-of-magnitude mismatch in integrated coefficients of thermal expansion (CTE) between room temperature and 4 K

$$\frac{L_{4K} - L_{293K}}{L_{293K}}. \quad (4.8)$$




Figure 4.13: A trap mounted by gluing the trap directly onto a copper mounting block using Epo-tek H21D. The trap cracked upon cooling to cryogenic temperatures. We attribute this behavior to the thermal expansion mismatch between copper and silicon and to the fact that H21D is relatively brittle. In the experiments discussed in this thesis, traps were mounted using more flexible glues or solders and shim layers of materials that can better absorb thermal stresses.

Integrated CTE values for relevant materials are shown in Table 4.2. The effect of mismatched CTEs can be mitigated by installing a thin shim layer between the copper and the silicon trap chip, either something more able to absorb thermal expansion stresses or with an intermediate CTE.

Before any traps were installed in the cryostat, all mounting methods were tested by mounting 4-10 sample silicon chips on copper and repeatedly cycling these samples between liquid nitrogen (LN₂) and temperatures. It was important to investigate sample chips cut to the same “bow-tie” shape as the trap, as the internal strain that determines whether a trap will shatter depends on the chip geometry.

The detector trap (Section 4.2.1.2) was glued using Master Bond EP21TCHT-1, which is thermally conductive and low-outgassing but performed better in LN₂ submersion tests than the Epo-tek H21D. To further limit strain, an intermediate 500 μm thick sapphire shim layer was glued between the copper and the trap. This mounting scheme performed well upon initial cooling and for over a year of continuous cold operation. However, after the cryostat was warmed to room

Material	$L_{4K} - L_{293K} / L_{293K}$ (%)
Copper	-0.324
Indium	-0.706
Silicon	-0.022

Table 4.2: Integrated coefficients of thermal expansion for relevant materials for trap mounting. Values from [93]

temperature for routine maintenance, the trap delaminated when it was cooled back down to 4 K.

The above mounting recipe was deemed inappropriate for future traps. The transport trap (Section 4.2.1.3) was instead bonded to the copper mounting block using indium as a solder. If indium is soldered directly to copper, the indium diffuses into the copper, creating a brittle alloy [103]. In order to avoid this, the copper mounting block was gold electroplated. If the entire mounting block and trap were heated to the indium melting point ($> 150^\circ\text{C}$), oxidization on the copper surfaces would be accelerated. Additionally, the bond would be made at an elevated temperature, which would introduce further thermal expansion mismatch at 4 K. Because of this, traditional soldering techniques were considered inappropriate. Instead, the indium was locally heated to its melting point using sandwiched layers of NanoFoil[®], a material made out of layers of nickel and aluminum. When mixed, these materials react highly exothermically [104] and the indium is locally heated without heating or oxidizing the larger copper block.

The traps are mounted individually to a small mounting block by the methods described above. This mounting block is then bolted into a the octagonal copper block shown in Figure 4.14, nicknamed the “trap puck.” The trap puck also holds the “filterboard,” a circuitboard used for delivery of electronic signals. Wirebonds deliver signals from the filterboard to the trap. The filterboard also contains the low pass filters which reduce technical noise on the trap electrodes. Everything trap-specific inside the vacuum is mounted on this puck. This enables faster exchange of traps, since the trap puck can be separately assembled while another trap is installed in the pillbox and actively trapping ions. All connections to this filterboard are made by mechanically robust connectors so the trap puck can be quickly installed in the pillbox.




Figure 4.14: The trap puck and filterboard for the detector trap. The trap is glued or bonded to a 5×8 mm OFHC copper mounting block in the center. All hardware between the mounting block and the cold head is bulk OFHC copper, limiting thermal resistance between the trap and the cold head. The mounting block is bolted into the trap puck. Also bolted to the puck is a “filterboard,” so-called because it includes an RC low-pass filter for each DC electrode. Wirebonds deliver electrical signals from the filterboard to the trap. Electrical connection from the pillbox wiring to the filterboard are made by 3x FFC connectors (on the backside) for the DC voltages and 3x SMP connectors (shown on top) for two microwave antennae and the SNSPD. The RF is delivered from the step-up resonator by a spring contact on the backside before passing through a via on to the top layer and being wirebonded to the trap. The filterboard and trap puck are trap-specific and a new puck is used for every trap.

4.2.4 Beryllium ovens and ablation target

To generate ions for trapping, a flux of beryllium atoms is produced above the trap and photoionized by the two-photon process discussed in Section 3.1. A Doppler cooling laser is applied simultaneously to cool the ions to a low enough temperature for ion trapping. One challenge in this process is the generation of an adequate flux of neutral beryllium. Beryllium has a low vapor pressure and a high melting point compared to other atoms used for trapped-ion quantum computing purposes, shown in Table 4.3. This means that the oven (which must have line of sight to the trapping zone) must be heated to $\gtrsim 1500$ K for ion loading.

A sample beryllium oven is shown in Figure 4.15. The oven is constructed from 50 μm diameter beryllium wire tightly wrapped around a 100 μm diameter tungsten filament. This tungsten

Element	Vapor pressure (at 600 K)	Melting point
Be	$< 10^{-12}$ Torr	1560 K
Mg	1.1×10^{-4}	923
Ca	1.7×10^{-7}	1112
Sr	3.2×10^{-6}	1042
Yb	2.9×10^{-5}	1097

Table 4.3: Vapor pressures and melting points of common elements used for trapped-ion quantum computing, from [105]. Beryllium has a lower vapor pressure and a higher melting point than other elements, requiring more heating to generate the atom flux gas used to load an ion trap.

filament is coiled into a tight spring. The coil is then spot welded onto stainless steel wires which are connected to Accuglass socket contacts crimped onto copper wire. These stainless steel wires are mounted in alumina tubes to electrically and thermally shield them from the copper mounting block. The current required in a typical trap loading run is 1 – 1.5 A through this 1 – 2 Ω oven for 30 – 45 s, generating 10’s of joules of heat. To mitigate the effects of this heat load on the trap, the ovens are independently mounted with a independent all-copper path to the cold head from the trap mount. The ovens will still radiatively heat the trap mount. During a typical loading run, the cold head temperature T_{cryo} increases by ~ 0.55 K and the trap temperature T_{trap} increases by ~ 2.5 K.

An alternate option is to load using laser ablation. This enables very local heating of the target without heating other areas of the trap. Laser ablation has been demonstrated for ion trap loading in many previous systems [106][107][108][109][110][111], including surface electrode traps [112][113] and cryogenic ion trapping experiments [91][114]. In contrast to the ovens, which require 10’s of joules to load, ions were loaded into the point trap by hitting a 250 μm diameter beryllium wire with 1-5 laser pulses of 1064 nm light with 1.85 – 2.25 mJ per pulse. A quick note on nomenclature: although ablation generally refers to all processes where a surface is hit with sufficiently high laser fluence to expel material, many people shortcut ablation to refer specifically to processes where the laser fluence is high enough to generate a plasma at the surface. This second process generates ions in a single step. However, these ions have higher energy. It is also




Figure 4.15: Resistive beryllium ovens. The visible coil is a 100 μm diameter tungsten wire, coiled into a filament shape. The tungsten is then wrapped in 50 μm diameter beryllium wire and gap welded to stainless steel leads at each end. These stainless steel leads are mechanically mounted in alumina tubes. These tubes are held in a copper housing that is directly connected to the cold head to limit heating of the trap mount.

not species selective. A photoionization laser was still required to trap, indicating that the ablation laser fluence was low enough that it primarily generated neutral beryllium.

The resistive ovens generate beryllium traveling in all directions, almost all of which are not useful for loading. The ablation target generates a “plume” of beryllium atoms that is more directed than the ovens [115], but still broader than the tightly-focused photoionization beam. To limit both the radiative load on the trap and the risk of neutral beryllium coating and shorting electrical traces on the trap or filterboard, the ovens and ablation targets are both shielded by a copper nozzle. Line-of-sight from the trap zone to the ovens and ablation target, respectively, are shown in Figures 4.16(a) and 4.16(b).

4.2.5 Filters

Technical noise at the trap secular frequency heats the ion, with

$$\dot{n} = \frac{e^2}{2m\hbar\omega} S_E(\omega) , \quad (4.9)$$

where $S_E(\omega)$ is the power spectral density of the electric field noise at frequency ω , in units $\text{V}^2/\text{m}^2/\text{Hz}$ [1]. In surface electrode traps, the generally smaller ion-electrode distance means that a given voltage noise on an electrode will translate to a larger electric field noise at the ion lo-




Figure 4.16: View over the trap boards of (a) the ovens and (b) the ablation target. The resistive ovens generate neutral beryllium in all directions. The ablation target generates a more directed “plume,” but still generates some beryllium in undesired directions. Since beryllium is electrically conductive, this stray beryllium risks shorting electrodes if it is not shielded. Therefore the ovens are housed with a “nozzle” that blocks all but a direct line of sight from the ovens to the trap.

cation than it would for a trap with a longer ion-electrode distance. This is a separate effect from *anomalous* ion heating, discussed in Section 1.2.1. Anomalous heating arises from some unknown microscopic mechanism on the electrode surface that produces electric field noise. Technical noise arises from voltage noise on the voltage sources and Johnson noise from room temperature electronics and will also heat the ion.

Unfiltered, the digital-analog converter (DAC) voltage noise alone introduces an unacceptable level of heating into this experiment. In the fast DACs used in this experiment, discussed in Section 4.2.6.2 and in [116], the voltage noise is dominated by noise from the amplifiers, Analog Devices AD8250, which is quoted as $180 \text{ nV}/\sqrt{\text{Hz}}$ when the amplifiers are operated with a gain of 10.

The electric field noise at the ion location due to this voltage noise can be simulated using the same Biot-Savart analogue discussed in Chapter 2. Voltage noise at a level of $180 \text{ nV}/\sqrt{\text{Hz}}$ on every electrode would produce heating rates of greater than one quantum per microsecond on all modes. To reduce this, low-pass filters attenuate the noise at high frequencies on each electrode. These low-pass filters are first-order RC filters with a cutoff frequency

$$f_{3dB} = \frac{1}{2\pi\sqrt{RC}}, \quad (4.10)$$




Figure 4.17: A first order RC filter, used to limit electric field noise in this experiment, with a 3 dB cutoff frequency $f_{3dB} = 1/(2\pi\sqrt{RC})$.

schematically drawn in Figure 4.17. The voltage noise will be attenuated by

$$V_{att} = V_{in} \left| \frac{1}{1 + i\omega RC} \right|. \quad (4.11)$$

For the trap integrated-SNSPD experiments, discussed in Chapter 6, the filters had a relatively low cutoff frequency of 7.1 kHz, since this was an untested trap technology. However, for the fast ion transport experiments, the voltages on each electrode were time dependent with non-negligible frequency components at or near the secular frequency. Therefore two different sets of filters were used with the transport trap: a “tight” filter with a 3 dB cutoff frequency of 4.8 kHz on electrodes whose voltages do not vary much during fast transport, and a loose filter with a 3 dB cutoff frequency of 23.4 kHz on the electrodes whose voltages vary more during transport. Filter values for each experiment are shown in Table 4.4.

The values listed in are listed in Table 4.4 are room-temperature resistances and capacitances. These values will shift at cryogenic temperatures. The resistors used in this experiment are thin film NiCr resistors, with a measured resistance shift of 3.6% at cryogenic temperatures. The capacitors use C0G dielectrics, which do not exhibit substantial change in dielectric contrast at cryogenic temperatures.

Future work in this experiment might explore higher-order low-pass filters or notch filtering

Filter	R (k Ω)	C (nF)	f_{cutoff} (kHz)
Detector trap	15	1.5	7.1
Transport trap (loose filter)	6.8	1.0	23.4
Transport trap (tight filter)	15	2.2	4.8

Table 4.4: RC filter components used for the experiments discussed in this thesis. The detector trap was installed with filters with a relatively low cutoff frequency of 7 kHz, since it incorporated a previously untested trap technology. The transport experiments (discussed in Chapter 7) required quickly varying waveforms with non-negligible frequency components at or near the trap secular frequency. This trap was installed with filters with a higher cutoff frequency (labelled “loose filter”). The voltages on some electrode do not vary much or at all during transport. Electrodes for which this was the case (CUL, CLL, CUR, CLR, DET, and the mesh) were installed with filters (labelled “tight filter”) with a lower cutoff frequency to reduce ion heating due to technical noise.

at the trap frequency. Since the trap mounting was designed to enable efficient exchange of traps, these could be explored with relatively low experimental downtime.

4.2.6 Generation of trapping voltages

4.2.6.1 RF source

A Holzworth HS9002A digital RF synthesizer provides 5.0–9.0 dBm of RF at around 67 MHz. In a 50 Ω circuit, this would correspond to 0.56 – 0.89 V, inappropriately low for ion trapping. We use a resonant circuit to step-up the RF voltage amplitude. The resonator is on a printed circuit board, following inspiration from [117]. The circuit is an LC resonator with a SMT capacitor and an inductor printed on the circuitboard. The printed inductor has reduced vibration sensitivity over a freestanding coil. The traps used in these experiments have a measured capacitance of 6 pF so with a coupling capacitor of 1.5 pF, a parallel shunting capacitance of 12 pF, and an inductance of 315 nH, the circuit has resonant frequency

$$f_{res} = \frac{1}{2\pi\sqrt{LC}} \approx 67 \text{ MHz} \quad (4.12)$$

These resonators have a measured Q factor of 169 at 4K and impedance $Z = \sqrt{L/C} = 132 \Omega$. The resonator step-up factor is

$$\frac{V_{out}}{V_{in}} = \sqrt{(Z/Z_0)Q} \quad (4.13)$$




Figure 4.18: A sample LC resonator of the design used in this experiment. Following inspiration from [117], the inductor is printed on the circuitboard, using vias through the circuitboard to achieve a coiled structure. The capacitors are surface mount. Contact to the filterboard is made by a spring contact. This structure is more compact and mechanically stable than the quarter-wave helical resonators that have been historically used in ion trapping experiments [119]. The resonators used in the experiments discussed in this thesis had inductance $L = 315$ nH and $C = 18$ pF (combining the trap, resonator, and stray capacitance), giving a resonant frequency of approximately 67 MHz. When connected to our cryogenic traps, the resonator has a quality factor $Q = 169$ and impedance $Z = 132 \Omega$, giving a step-up factor of 18.5.

[118]. This gives a step-up factor of 18.5, enabling 7.4 – 11.6 V amplitude RF drive voltage at the trap with the described RF power input. This is appropriate for trapping of ${}^9\text{Be}^+$, a low-mass ion.

4.2.6.2 DC voltage generation

For all of the experiments discussed in this thesis, dynamical control over the trap “DC” voltages is necessary. (“DC” in quotes, because truly DC voltages have no dynamical components; however, they are so-called to distinguish them from the RF voltages) More specifically, these are the voltages providing axial confinement. In all the experiments discussed in this thesis, time-dependent control of these voltages is required to compensate stray electric fields and to transport the ion along the trap axis. The most demanding requirements on the so-called DC voltages arise in the fast transport experiments, discussed in Chapter 7. In these experiments, the voltages on the trap electrodes must be on the scale of ± 1 V, varying on a time scale of tens of trap periods

($\approx 10 \mu\text{s}$). Further, since the low-pass filters appreciably attenuate the fastest frequency components of the fast transport waveforms and will warp the fast transport waveforms, DACs must output predistorted waveforms (discussed in Chapter 7) with voltage amplitudes much larger than $\pm 1 \text{ V}$. In these experiments, transport waveforms were generated by the fast digital-analog converters (DACs) discussed in [116]. These DACs have a measured maximum slew rate of $35 - 40 \text{ V } \mu\text{s}^{-1}$ (variable device to device) and a maximum amplitude of $\pm 10 \text{ V}$.

4.2.7 Magnetic field coils

An externally applied magnetic field is required to define the quantization axis and lift the degeneracy of the Zeeman sublevels. Since the optical pumping in this experiment is polarization dependent, this magnetic field must be well-aligned with the optical pumping laser (the BD). Five magnetic field coils generate these fields: two aligned along the BD laser axis, which provide the quantization field; two aligned perpendicular to the BD laser axis in the plane of the optical table; and one aligned normal to the optical table. The quantization field is generated by running 3 A through the first two coils, which creates a 5.7 G magnetic field at the ion location. The other three coils cancel out Earth's magnetic field and any other stray global magnetic field in the laboratory.

The magnetic field orientation can be aligned by observing the depumping rate of an ion prepared in the bright state $|F = 2, m_F = -2\rangle$. If the magnetic field is well aligned with the BD and the BD laser polarization is purely σ^- , we will measure a Poissonian histogram of counts per detection. Magnetic field misalignment will depump ions to a dark state. The BD direction and quantization axis are well aligned when dark state counts are minimized. We generate bright states with $< 1\%$ infidelity, with residual infidelity attributed to imperfect polarization in the BD laser. Although the light is polarized before the vacuum using a Glan-laser polarizer and waveplates, the laser passes through three fused silica windows before reaching the ion. Strain in these windows will induce spatially-dependent birefringence, giving unavoidable polarization imperfections at the ion.




Figure 4.19: Laser beam configuration. The trap secular modes are oriented along $\hat{q}_{ax} = (1, 0, 0)$, $\hat{q}_{r1} = (0, \sin(15^\circ), \cos(15^\circ))$, and $\hat{q}_{r2} = (0, \cos(15^\circ), -\sin(15^\circ))$ so the BD coming in at 45° from the trap axis will have overlap with all three modes. Raman 1 and Raman 2 are counterpropagating, so $\Delta\vec{k}$ is oriented along this same axis; this also has overlap with all three secular modes, so sidebands of all three modes of motion can be driven using this laser configuration. The PI and the Raman lasers never need to be operated at the same time, so the final mirror for ‘Raman 1’ is on a magnetic kinematic base that can be removed from the beam path when loading ions, allowing PI light through. The filterboard surface mount components are placed out of the way of these beam paths.

4.2.8 Laser sources

There are three major wavelengths needed for these experiments:

- 1064 nm, for laser ablation
- 235 nm, for photo-ionization
- 313 nm, for cooling, optical pumping, state detection, repumping, and Raman transitions.

The laser beam configuration on the trap is shown in Figure 4.19.

4.2.8.1 Ablation laser

As discussed in Section 4.2.4, laser ablation can be used to generate a flux of neutral beryllium atoms over the trap zone for ion loading. Laser ablation depends strongly on peak pulse energy [115]. In this experiment, laser ablation was performed at a wavelength 1064 nm, where commercial Nd:YAG systems are widely available, using a Continuum Minilite, which can provide peak pulse energy 0.2 – 50 mJ. Initial rough alignment was performed with a co-aligned 532 nm laser pointer. The alignment wavelength was chosen for ease of use with most cameras and because 532 nm laser diodes are widely available as cheap, off-the-shelf laser diodes. Since this is the second harmonic of 1064 nm light, co-coated mirrors and lenses for 532 nm and 1064 nm light are available off-the-shelf with no need for custom coating runs. The beam was aligned by observing scatter on the ablation target with a CCD camera and a zoom lens. If the ablation laser is not incident on the target, it has a clear path through the vacuum system. Final alignment was performed by centering the pulsed 1064 nm source (run in low power mode) on the ablation target using a knife edge measurement.

Laser ablation was tested in this apparatus in the point trap, discussed in Section 4.2.1.1. With similar RF frequencies and amplitudes to loading conditions using the resistive ovens, ions were reliably loaded with 1-5 1.85 – 2.25 mJ pulses focused to 75 microns beam radius, or a peak fluence of 21 – 25 J/cm².

4.2.8.2 Photo-ionization laser

Previous work with ⁹Be⁺ at NIST has photoionized neutral Be using a pulsed, frequency tripled (one step of frequency doubling plus one step of sum frequency generation) Ti:Sapph laser to generate 13 – 18 mW of 235 nm light [47]. Pulsed lasers have a higher peak intensity than a CW laser of the same average power, which is convenient for the nonlinear frequency conversion necessary to generate 235 nm light. However, this high peak intensity also caused laser-induced damage on optical components, particularly in the frequency conversion setup.

For the experiments discussed in this thesis, photoionization light was instead generated by

a continuous wave laser. The CW 235 nm light was generated by two stages of resonant frequency doublers from a CW Ti:Sapph running at 940 nm. (This is close to a water absorption line, so the first stage of resonant frequency doubling must be enclosed within a box containing desiccant.) The frequency was measured using a wavemeter and manually set 319.01924 THz in the IR but not locked to a physical feature. The power broadened full width at half maximum of the photoionization transition at typical powers used for loading is 1.7 GHz (measured by Hannah Knaack, publication forthcoming), so the laser needed to be within ~ 0.85 GHz of the center frequency to load. Loading is ideally a quick and infrequent process, so it is sufficient to just occasionally check this frequency.

To ensure beam pointing stability, the PI light is fiber coupled into a single-mode LMA-10 photonic crystal fiber with 10 μm mode field diameter, hydrogen-loaded and cured as described in [120]. At 235 nm, these fibers have approximately 20-25% transmission for a 20 – 30 mm fiber length. After the fiber, we measured 0.5 – 2 mW of 235 nm light on the cryostat optical table, sufficient for loading. This is much lower power than the average power that was needed to load with the pulsed laser.

4.2.8.3 Resonant and Raman lasers

All of the 313 nm laser beams (the BD and RD resonant lasers and the Raman lasers) are generated via multiple stages of nonlinear optics. A single 1050 nm fiber seed, amplified with a Koheras Boostik 10 W fiber amplifier, is used to generate 6 W of infrared laser light at 1050.132 nm. This single beam is split into three 2 W beams to feed each of the BD, RD, and Raman beamlines. In each beamline, the 1050 nm light is combined with approximately 1550 nm light in a single-pass periodically poled lithium niobate (PPLN) waveguide to generate red light at 626 nm via sum frequency generation (SFG). The BD, RD, and Raman beamlines each need different frequencies, so three different seed lasers with different frequencies are used. These seed lasers are each amplified with a 5W fiber amplifier to generate ~ 3 W IR light at 1550 nm. The exact wavelengths of each beam are listed in Table 4.5. The SFG setup is discussed in more detail in [47] and in [121].

The BD and RD laser beams both drive resonant transitions, so they must be stabilized to

an absolute wavelength reference. About 10 mW of the red light from each beam are picked off for frequency stabilization. The picked-off light is shifted using acousto-optic modulators (AOMs) by -670 MHz and -279 MHz (for the BD and the RD, respectively) into resonance with an iodine transition. The frequency of these auxiliary beams is measured relative to a hyperfine transition of molecular iodine using Doppler-free saturated absorption spectroscopy [122] and locked by feeding back on the piezo of the respective seed lasers. The frequency stabilization setup is discussed in more detail in [47]

Ultraviolet light on each beamline is then generated with second harmonic generation in a resonant cavity containing a beta barium borate (BBO) crystal. The cavity resonance frequency is locked to the laser using a Hänsch-Couillaud lock circuit [123]. These resonant cavity doublers are described in more detail in [47] and in [121]. A pickoff and an AOM in a noise-eater configuration are used to stabilize the intensity of the UV output light.

Certain of the frequencies are close enough that they can be produced from the same UV source. The frequencies that are separated only by the ground state hyperfine splitting - i.e. the BD and the BDD (split by approximately half the hyperfine splitting, 625 MHz), the two repumping frequencies (split by approximately the hyperfine splitting, 1.25 MHz), and the two Raman frequencies (split by approximately the hyperfine splitting) can each be produced from a single UV doubler. The light can then be shifted to the desired frequencies by AOMs which can shift frequencies by hundreds of MHz. The use of AOMs also enables fast on/off switching of each beam through fast switching of the RF driving the corresponding AOM. (These respective classes of transitions are separated from each other by 80 – 200 GHz, so they cannot be bridged by AOMs.) This frequency shifting is performed by AOMs in double-pass configurations. The BD and RD AOM configuration that gives the resonant frequency desired for the BD, BDD, and two repumping transitions is shown in Figure 4.20. The BD, BDD, and RD beams are fiber coupled into the same UV single mode fiber (prepared using the method described in [120]) ensuring that they are well co-aligned on the ion.

The Raman transitions are first-order insensitive to the laser wavelength and thus are not

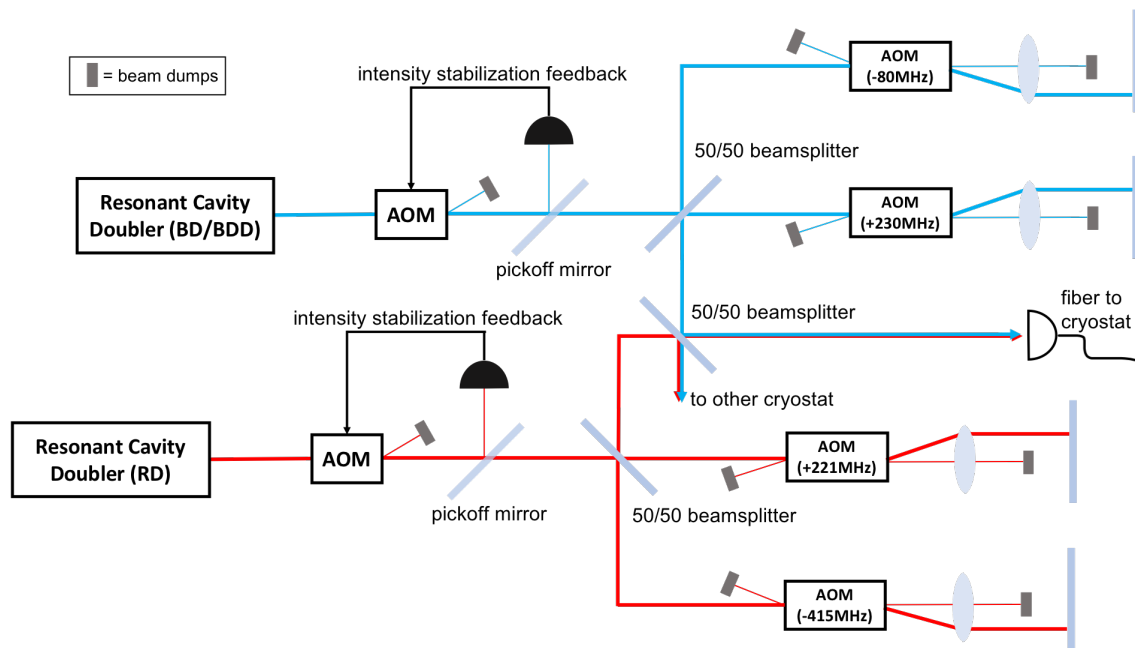


Figure 4.20: UV frequency shifting setup for the BD, BD detuned, and repumping beams. The BD and BDD beams are generated from a single cavity doubler. Similarly, both RD beams are generated from the same cavity doubler. The output of each doubler is intensity stabilized. Both beamlines are split with a 50/50 beamsplitter, resulting in four beams, each of which is frequency shifted by a double-pass AOM at the frequencies shown. Each beamline is then recombined on the same beamsplitter, then the two are combined together on a beamsplitter which is fiber coupled and sent to the cryostat. The other port of the beamsplitter goes to the other cryostat.

locked to an absolute reference; however, the two-photon process (discussed in Section 3.3.4) is highly sensitive to the *relative* frequency between the two beams, which must be quite stable. The Raman beam configuration is shown in Figure 4.21. The frequency offset between the two beams is generated by splitting the beam on a Glan-laser polarizer and sending each of the two resultant beams through double-pass AOMs driven by the same RF frequency. The +1 order of one AOM and the -1 order of the other AOM are taken, so in total, the beams are shifted relative to each other by four times the drive frequency. The two beams are recombined back on the same polarizer, but slightly offset vertically. This produces a beam with different frequencies in each polarization. The beam is picked off and fiber coupled to be sent to the trap. On the output, the different frequency components are split with a Glan-laser polarizer and sent to two different experimental ports as shown in Figure 4.19.

Application	Wavelength range	Wavelength
Photoionization (pre-quadrupling)	IR	939.75 nm
Photoionization [82]	UV	234.9329 nm
Sum frequency generation (all laser lines)	IR	1050.132 nm
BD sum frequency generation	IR	1549.408 nm
BD red light	Vis	626.267 nm
BD resonant light (cooling, optical pumping, detection)	UV	313.133 nm
RD sum frequency generation	IR	1550.192 nm
RD red light	Vis	626.835 nm
RD resonant light (repumping)	UV	313.198 nm
Raman sum frequency generation	IR	1550.032 nm
Raman red light	Vis	626.46 nm*
Raman laser (sideband transitions)	UV	313.23 nm*

Table 4.5: Exact laser wavelengths for ${}^9\text{Be}^+$ manipulation (in vacuum). Raman transitions are only second-order dependent on wavelength, so these lasers are free-running and not locked to an absolute reference.

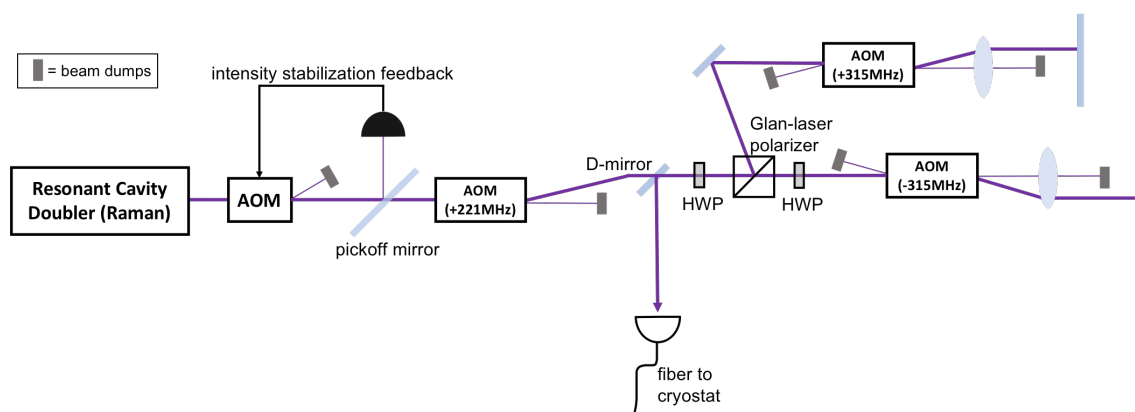


Figure 4.21: UV frequency shifting setup for the Raman lasers. Like the BD/RD (Figure 4.20), the Raman UV is intensity stabilized after the resonant cavity doubler output. It then passes through a switch AOM driven at 221MHz. The beam is split into two beams with opposite polarization by a Glan-laser polarizer. The two beams are shifted by equal and opposite frequencies by taking the ± 1 orders of two double-pass AOMs driven by a single RF drive. The whole setup together is nicknamed the “quad pass,” since the total frequency shift between the two beams is four times the AOM drive frequency. The beams are recombined using the same polarizer but there is a slight vertical offset. The output beams are picked off and sent to the experiment through a single mode fiber. This process produces a beam where the frequencies are different in each polarization. On the output of the fiber, the two frequency components of the beam can be split with a Glan-laser polarizer and sent to two different experimental ports as shown in Figure 4.19. This reduces path length fluctuations between the two Raman beams.

4.2.9 Imaging system

Ion trapping imaging typically employs custom objectives designed to maximize numerical apertures at the ion detection wavelengths, in the UV for most species. Imaging in this apparatus was instead performed by an objective built using off-the-shelf optics. All of the imaging optics are outside vacuum, enabling modification of the objective without warming the cryostat to room temperature and breaking vacuum.

In order to shield the cryogenic system from room-temperature blackbody radiation, there are multiple levels of radiation shielding between the ion and the outside world. This means the imaging system must have a working distance of more than 35mm, relatively long for a high NA system. A pair of commercial aspheric lenses (Asphericon AFL50-60-S-U) with 60 mm effective focal length and 0.39 NA form a relay image with no magnification. This relay image is magnified by a 40x Thorlabs LMU-40X-UVB with NA 0.5. A ray diagram of this imaging system is shown




Figure 4.22: Zemax ray tracing of the imaging system used for qubit state readout, shown with 313 nm light.

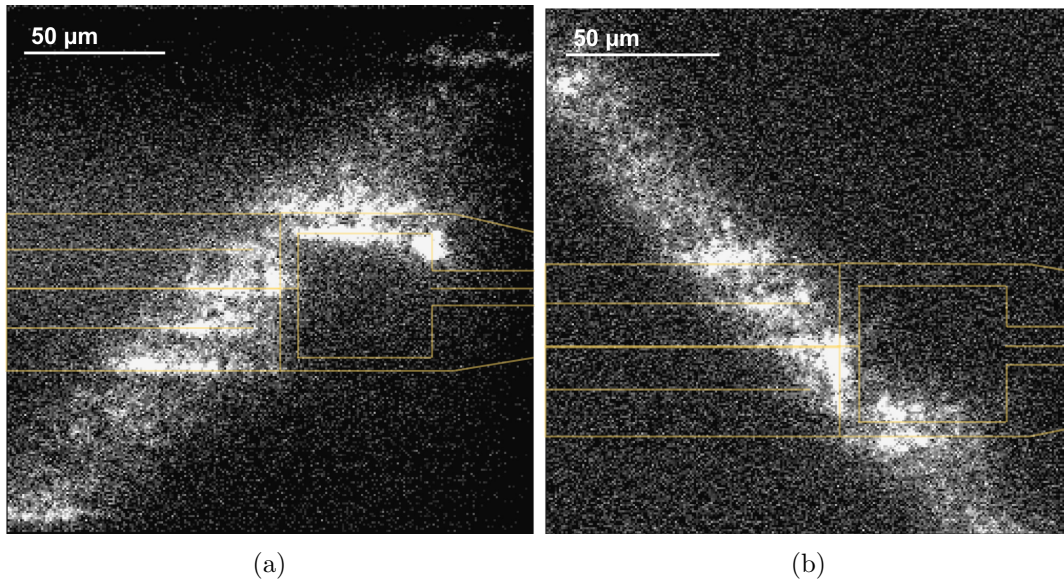


Figure 4.23: Typical images of laser scatter off the detector trap surface that would be used for laser alignment. In both cases, the laser is moved 40 μm closer to the surface and attenuated heavily so it can be aligned in reference to the trap electrodes. Scattered light off the edge of electrodes is imaged onto the EMCCD camera. The trap electrode geometry is overlaid in yellow on these images as a guide to the eye. (a) The BD beam, with the objective positioned to image 313 nm light. (b) The PI beam, with the objective positioned to image 235 nm light.

in Figure 4.22. Based on Zemax simulations, the system has an overall NA of 0.38.

Although the system is primarily designed to have high NA and low distortion for the 313 nm light used for state-dependent fluorescence imaging, this imaging system also functions at 235 nm, which is useful for laser beam alignment. Alignment is performed by scattering the laser off the trap and imaging scatter off the electrode edges on the camera. The beam can then be positioned in the plane of the trap based on these trap features before it is moved 40 μm away from the trap

to the ion height. While scattered light from the BD, RD, and Raman beams are close enough in wavelength to be imaged at the same working distance, the chromatic shift in the objective requires a different working distance to image 235 nm light when aligning the PI beam. This imaging can be performed by moving the lens 7 mm closer to the trap. Figures 4.23(a) and 4.23(b) show typical images of 313 nm and 235 nm light, respectively, during this alignment process.

Depending on the experiment, we may wish to image ion fluorescence onto either a camera or a fast, high-quantum efficiency photo-multiplier tube (PMT). To achieve this, the imaging system includes a motorized mirror that can be switched between two configurations. In one configuration, the mirror is moved out of the way and fluorescence is imaged onto an electron-multiplying CCD (EMCCD) camera DU-885K-CS0-#VP with a quoted quantum efficiency of 25% at 313 nm. The camera is used for beam alignment but not for readout. The other flip mirror configuration sends fluorescence to a Hamamatsu R7600P-203 PMT with a quoted quantum efficiency of 38%. The PMT is used for qubit state readout. Future experiments (particularly the array trap experiments discussed in [47]) may wish to make use of the spatial resolution available on the EMCCD. For the experiments discussed in this thesis, the PMT is sufficient.

The imaging system is mounted in a bore underneath the optical table, making it challenging to access. In order to control the position of the lens (for both switching between the 235 nm and 313 nm working distances and for imaging an ion in different trap zones), compact stepper motors control a 3-axis translation stage on which the objective is mounted. The stage is mechanically mounted on a small platform within the optical table bore. The relay optics are made from identical aspheres with minimized spherical aberration. However, they are not corrected for off-axis aberrations like coma. Because of this, the objective performs best when the ion, objective, and camera sensor are all aligned along a single axis. If the objective is shifted off-axis, the resolution will be degraded. Because of this, the camera is also mounted on a two-axis translation stage. This is only rarely adjusted, so it is not motorized.




Figure 4.24: A 5 ion chain loaded in the detector trap, fluorescing under BD illumination and imaged on the EMCCD camera.

4.2.10 Experimental control system

The experiments described in this thesis require microsecond to nanosecond resolution experimental control. High fidelity qubit state detection requires sub-microsecond photon arrival time information and the transport experiments require sub-microsecond synchronization between all DACs. These systems must also be well synchronized to each other. This is a nontrivial timing task beyond the capabilities of many conventional laboratory control systems. Experimental control is performed using a KC705 field programmable gate array (FPGA) evaluation board to control a simultaneously clocked set of digital logic (TTL) and direct digital synthesizer (DDS) signals with 1 ns timing resolution. This FPGA is programmed using the Advanced Real-Time Infrastructure for Quantum physics (ARTIQ) language [124], a compiled Python language developed specifically for quantum physics applications. “Slow” applications, such as moving motors, are not controlled from the FPGA but can be integrated with the ARTIQ software or triggered by a TTL pulse from the ARTIQ hardware.

4.3 Trapping ions in the new apparatus

Ions were first trapped in the detector trap in June 2018. A sample 5 ion ${}^9\text{Be}^+$ chain is shown in Figure 4.24.




Figure 4.25: Heating rates as a function of frequency for ions in the (a) detector and (b) transport traps. These are measured in trap zone 3, as labelled on Figure 4.8. The heating rate on the axial mode is measured for multiple axial frequencies by scaling the axial trapping voltages. The axial mode heating rates are higher in the transport trap than in the detector trap. This can be attributed this to the higher corner-frequency filters on the transport trap, which will transmit a higher level of electric field noise at the trap secular frequencies.

4.3.1 Ion lifetime

One of the goals of this cryogenic apparatus was a sufficiently good vacuum level that ions are not lost frequently due to background gas collisions, even with the reduced trap depths of surface traps. In fact, the ions are so infrequently lost due to background gas collisions that it is impractical to measure a trap lifetime. Doppler-cooled ions in both the detector and transport traps were regularly held for days at a time in all trapping zones, with the lifetime limited by user error and not background gas collisions.

4.3.2 Heating rates

The heating rate was measured in both the detector trap and the transport trap and the results are shown in Figure 4.25. In both cases, the heating rate was measured for an ion trapped zone 3 over electroplated gold (as indicated on Figure 4.8). The heating rate for an ion trapped over the SNSPD is discussed in Chapter 6.

The ion heating rate and electric field noise spectral density $S_E(\omega)$ are related by

$$\dot{n} = \frac{q^2}{2m\hbar\omega} S_E(\omega) \quad (4.14)$$

[1] so heating rates can be used as a probe of electric field noise. $S_E(\omega)$ is the power spectral density of the electric field in units $(\text{V}^2/\text{m}^2)\text{Hz}^{-1}$. Figure 4.26 shows the product of the frequency and the electric field noise for these measurements. These can be compared to the canonical “heating rate scatter plot” in Figure 4.27. The scatter plot compiles $\omega S_E(\omega)$ from a large set of trapped ion experiments and shows a roughly $1/d^4$ dependence on ion-electrode spacing. It shows a ~ 100 x order of magnitude reduction in cryogenic heating rates over room temperature heating rates.

The choice of $\omega S_E(\omega)$ as a figure of merit assumes that the electric field noise arises from a single noise source with a power-law scaling $S_E(\omega) \propto \omega^{-1}$. This is not known to be generally true, and a compilation of frequency scalings measured in different experiments has shown frequency dependencies varying between $S_E(\omega) \propto \omega$ and $S_E(\omega) \propto \omega^{-6}$, which $S_E(\omega) \propto \omega^{-1}$ the most frequently measured [125]. Nevertheless, $\omega S_E(\omega)$ is a frequently used figure of merit within the ion trapping community and is useful for comparing heating rates between ions of different masses or with different motional frequencies.

The measured heating rates in these experiments give $\omega S_E(\omega)$ between 7×10^{-6} and $3 \times 10^{-5} \text{ V}^2/\text{m}^2$. The electric field noise for the axial modes is elevated by around a factor of two in the transport trap relative to the detector trap. This can be attributed to the higher corner frequency f_{3dB} on the low-pass filters used in the transport trap, which will transmit more voltage noise at the trap secular frequencies. This range of values for $\omega S_E(\omega)$ is plotted compared to previously measured values of $\omega S_E(\omega)$ in Figure 4.27. The measured electric field noise in these experiments is comparable to other cryogenic surface electrode traps with $\approx 40 \mu\text{m}$ ion-electrode distance.




Figure 4.26: Electric field noise $\omega S_E(\omega)$ as a function of frequency in the (a) detector and (b) transport traps. The radial modes values are statistically indistinguishable. The axial heating rates are elevated in the transport trap relative to the detector trap, but are comparable to other experiments with short ion-electrode distance, shown in Figure 4.27.




Figure 4.27: Electric field noise vs ion electrode distance as measured in a large number of experiments. $\omega S_E(\omega)$ is plotted, which assumes that the electric field noise spectral density varies with ω^{-1} . Ion heating is not entirely understood, but this plot is commonly pointed to as evidence that the heating rate depends on ion-electrode distance with a roughly $1/d^4$ scaling, and cooling traps to cryogenic temperatures introduces a $\sim 100x$ reduction in heating rate. The range of $\omega S_E(\omega)$ measured in this experiment is shown as a purple rectangle. Modified from [50].

Chapter 5

Methods

5.1 Qubit state detection

Any pure two-level quantum state can be completely described by

$$|\psi\rangle = \cos\theta|\uparrow\rangle + e^{i\phi}\sin\theta|\downarrow\rangle \quad (5.1)$$

I will describe a method for identifying the *probabilities* of being found in $|\uparrow\rangle$ or $|\downarrow\rangle$, i.e. $P_{\uparrow} = |\langle\uparrow|\psi\rangle|^2$ and $P_{\downarrow} = |\langle\downarrow|\psi\rangle|^2$. Detection of the phase ϕ to fully identify $|\psi\rangle$ requires quantum state tomography (QST), beyond the scope of this thesis. For more, see [126] and [60].

The qubit state is detected by state-selective fluorescence. We drive the $|\downarrow\rangle \leftrightarrow {}^2P_{3/2}$ transition with the BD laser tuned to resonance. An ion in the $|\downarrow\rangle$ qubit state will be excited into the ${}^2P_{3/2}$ and will quickly decay back to $|\downarrow\rangle$, emitting a photon into 4π , where it can be detected by a single photon detector. This light is > 1 GHz off-resonance for an ion in $|\uparrow\rangle$. Because of this, I will refer to $|\uparrow\rangle = {}^2S_{1/2}|F=1, m_F=-1\rangle$ as the “dark” state and $|\downarrow\rangle = {}^2S_{1/2}|F=2, m_F=-2\rangle$ as the “bright” state.

We perform detection by illuminating the ion with our detection laser for some chosen period of time t_{detect} and counting ion fluorescence during this time. The challenge is accurately identifying, from the number of photon events if the ion is in the bright state or dark state. I will discuss two techniques for making this identification: thresholding, a simple way of identifying the ion state, and adaptive Bayesian methods, which we can use to achieve higher detection fidelities but which require photon arrival time information.

For both of these, it is useful to have a precise and consistent definition of detection fidelity.

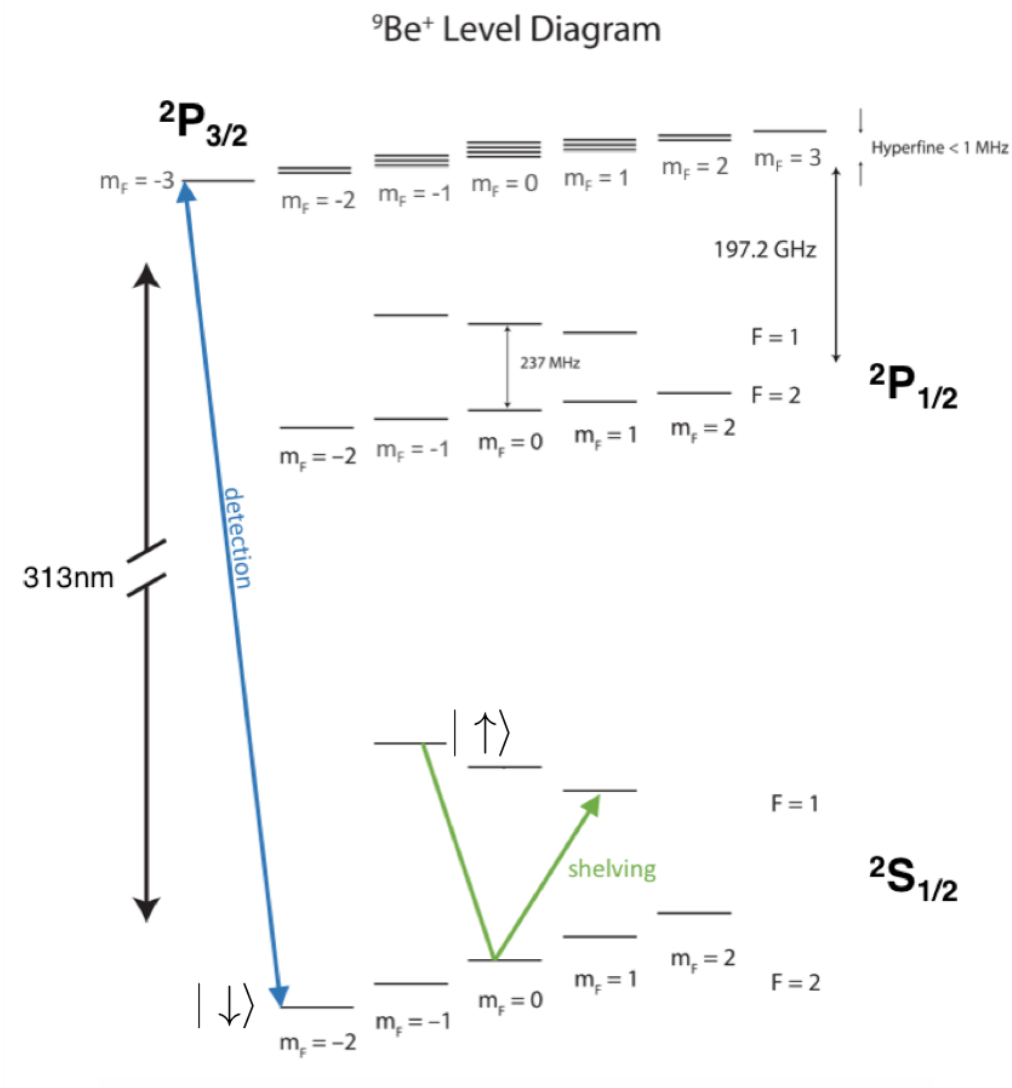


Figure 5.1: Relevant transitions for qubit state detection in ${}^9\text{Be}^+$, adapted from [80]. The detection laser is resonant with the $|\downarrow\rangle \leftrightarrow {}^2P_{3/2}$ transition, which drives state-dependent fluorescence, causing an ion in $|\downarrow\rangle$ to fluoresce, but not an ion in $|\uparrow\rangle$. To maximize detection fidelity, dark ions are shelved in the $|F = 1, m_F = 1\rangle$ level.

Following definitions in [127], I will define the detection fidelity

$$\mathcal{F} = 1 - \epsilon = 1 - \frac{\epsilon_b + \epsilon_d}{2} \quad (5.2)$$

where ϵ_b is the bright state error, i.e. the probability of mis-labeling a bright ion as dark, and ϵ_d is the dark state error, i.e. the probability of mislabeling a dark ion as bright. This is dependent both on the apparatus and on the analysis.

5.1.1 Thresholding

In the absence of errors caused by off-resonant transitions or imperfect closing of the cycling transition, minimization of state detection errors would be a relatively straightforward process. Consider an ion in the state

$$|\psi\rangle = a|\downarrow\rangle + b|\uparrow\rangle \quad (5.3)$$

If the ion is measured in the bright state $|\downarrow\rangle$, it will scatter photons that will be detected at some rate R_b . In the ion is measured in the dark state, it should not scatter photons; however, there will be some rate of dark counts R_d due to detector dark counts and background scatter. Upon many repetitions, the number of counts measured during detection time t_{detect} will take a double-Poissonian distribution:

$$P(n) = |a|^2 \left(\frac{\lambda_b^n e^{-\lambda_b}}{n!} \right) + |b|^2 \left(\frac{\lambda_d^n e^{-\lambda_d}}{n!} \right) \quad (5.4)$$

where $\lambda_b = R_b t_{detect}$ and $\lambda_d = R_d t_{detect}$. This distribution is shown in Figure 5.2. Two cases are shown, with identical R_b and R_d , but with two different choices of t_{detect} .

Under these ideal conditions, the ion state can be identified with high fidelity using a simple thresholding method. We choose a threshold cutoff value n_c and count runs. Any run in which $n \geq n_c$ counts were detected is labelled bright; any run in which $n < n_c$ counts were detected is labelled dark. The optimal threshold is the number of photons at which the bright and dark state Poissonian distributions intersect

$$n_c = \frac{\lambda_b}{\ln\left(1 + \frac{\lambda_b}{\lambda_d}\right)} \quad (5.5)$$




Figure 5.2: Ideal Poissonian histograms for an ion in an equal superposition of the bright and dark states, shown on a logarithmic scale. In (a), Poissonian parameters are $(\lambda_B, \lambda_D) = (10, 1)$. In (b), Poissonian parameters are $(\lambda_B, \lambda_D) = (30, 3)$. This is equivalent to the case in which (a) and (b) are detected from systems with equal scattering rates, but the detection time for (b) is three times as long as for (a). Depumping and repumping are entirely neglected.

[127] where λ_b and λ_d are the means of the bright and dark state Poissonians, respectively. However, there is some statistical probability that a purely bright state Poissonian will produce fewer counts than the threshold and conversely that a purely dark state Poissonian will produce more counts than the threshold. This introduces bright state and dark state errors ϵ_B and ϵ_D :

$$\epsilon_b = P(\text{measure dark}|\text{bright}) = \sum_{i=0}^{n_c} \frac{\lambda_b^i e^{-\lambda_b}}{i!} \quad (5.6)$$

$$\epsilon_d = P(\text{measure bright}|\text{dark}) = 1 - \sum_{i=0}^{n_c} \frac{\lambda_d^i e^{-\lambda_d}}{i!} \quad (5.7)$$

In the examples shown in Figure 5.2, Eq. 5.5 gives cutoff thresholds of $n_c = 4.2$ and 12.5 for the short and long detections, respectively. In practice, this means that when thresholding the short detection, runs in which $n \leq 4$ counts were detected would be labelled as dark, and all other would be labelled as bright. For the long detection, runs in which $n \leq 12$ counts are detected would be labelled dark, with all others labelled bright. In the first case, there is substantial overlap between the bright and dark Poissonians, so many runs will be mislabelled by this method. From Eqs. 5.6 and 5.7, $\epsilon_b = 0.029$ and dark state error $\epsilon_d = 0.004$, for an overall detection fidelity $\mathcal{F} = 98.3\%$. If we instead consider the second case (the long detection case), we label all runs in which 13 or more

counts are measured are bright ions; all runs with 12 or fewer counts are labelled dark. This gives a probability $\epsilon_b = 1.7 \times 10^{-4}$ of mislabeling the bright state as dark and a probability $\epsilon_d = 1.6 \times 10^{-5}$ of mislabeling the dark state as bright, for an overall detection fidelity $\mathcal{F} = 99.99\%$.

From this, it appears that high fidelity detection is simply a matter of choosing t_{detect} such that the bright and dark state histograms are well separated. However, I have critically neglected a major source of error. The histograms are in fact non-Poissonian, primarily due to off-resonant transitions and imperfections in our cycling transition.

5.1.2 Depumping and repumping

The leading source of error in our detection is caused by errors that flip ions between our qubit states during the detection process. There are two types of these errors, depending on where the ion starts.

- **Repumping:** an ion that is dark at the start of the detection flips to bright during t_{detect} . These events are primarily caused by off-resonant transitions.
- **Depumping:** an ion that is bright at the start of the detection flips to dark during t_{detect} . These events are primarily caused by imperfect σ^- polarization and by misalignment between the detection beam and the quantization axis.

The rate of depumping and repumping events is constant per unit time; therefore, with increasing detection times, the probability that a depumping or repumping event has occurred grows.

These errors will negatively impact our ability to accurately label an ion's state using thresholding. For an illustration, consider the data shown in Figure 5.3. This plot shows the times at which photons were detected for two different detections. In both cases, 15 photons were counted in a 500 μs detection period. However, the photon arrival time illustrates that these ions are likely in different states. In Figure 5.3(a), the ion begins bright and scatters photons at a high rate; however, at around 100 μs , the ion appears to be depumped and is dark for the rest of the detection interval. In 5.3(b), the ion scatters no photons until before 200 μs , after which point it scatters photons at

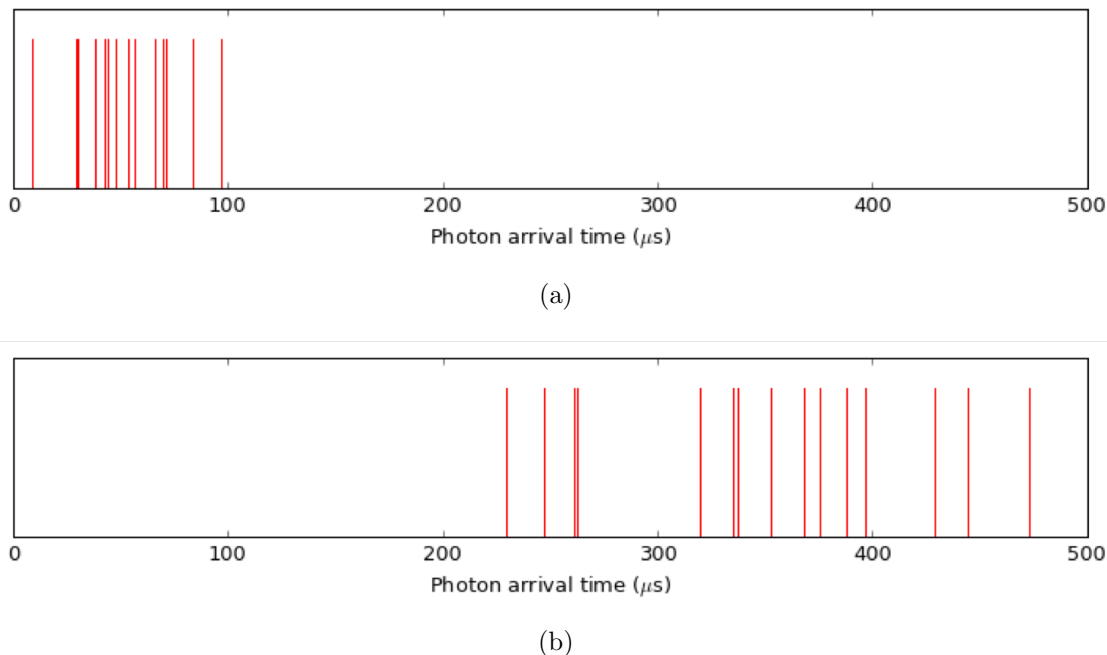


Figure 5.3: Photon arrival times for (a) an initially bright ion which is depumped dark during detection and (b) an initially dark ion that is repumped bright during detection. Both plots show runs in which 15 photon counts were collected, so state identification using thresholding would label them as identical.

a high rate, likely indicating that the ion was repumped from the dark state to the bright state at this time. Thresholding would identify both of these cases as identical.

In practice, repumping and depumping introduce a “tail” in between the two peaks of the double Poissonian, as can be seen in Figure 5.4. Since the probability that an ion has been depumped or repumped grows with increasing detection time, we now have two competing factors. We need to detect for long enough that the two Poissonians are well-separated. After they are separated, we should detect for as short a detection time as possible, in order to limit depumping and repumping errors.

There several techniques we use to mitigate errors due to repumping and depumping:

- Microwave shelving - as I will discuss below, we can reduce the repumping rate by transferring dark ions to a state that requires multiple off-resonant scattering events to transition back to the bright state.




Figure 5.4: Simulated histograms for a perfect cycling transition with no depumping or repumping (blue) compared with real, experimental histograms (red). Notice the tail between the two histogram peaks in the data, introducing error into qubit state detection using thresholding.

- Improvements to the readout beam polarization and its alignment with the quantization axis, ensuring that the detection transition is as close to a cycling transition as possible
- Improved rate of detection: if, by increasing R_b , the bright and dark histograms separate from each other faster, we can use a shorter t_{detect} , reducing the probability of depumping or repumping.
- Statistical methods that take advantage of time stamping information to better identify bright and dark ions.

With the combined use of all of these techniques, including a detector with higher quantum efficiency than the best photomultiplier tubes at 313 nm, as I will discuss in Chapter 6, we are able to detect ${}^9\text{Be}^+$ states with detection fidelity $\mathcal{F} > 99.9\%$.

5.1.3 Shelving

Microwave shelving allows us to reduce repumping errors. Although the probability of an ion in the dark state $|\uparrow\rangle$ being excited by the detection laser is small, if such an event *does* occur, the ion will be excited to the ${}^2P_{3/2}|m_F = -2\rangle$ state, which has a non-negligible probability of decaying to the $|\downarrow\rangle$ qubit state, which is bright. Instead consider a “shelved” ion, which stored in the $|F = 1, m_F = 1\rangle$ state by applying two microwave π -pulses to transfer the population. If such an ion is excited off-resonantly by the detection laser, it will require at least one additional off-resonant scattering events to decay to the bright state. This shelving is shown on the ${}^9\text{Be}^+$ level diagram in Figure 5.1.

5.1.4 High fidelity state detection using adaptive Bayesian statistics

The thresholding method discussed above uses a single piece of information - counts collected during some detection window - to identify an ion as bright or dark. However, we can also use information about photon arrival times. Intuitively, one can look at the arrival time data in Figure 5.3 and see that this can give us information about whether the ion was the dark or bright state at the start of the detection. We can formalize this instinct for a method that gives both improved detection fidelity and shorter average detection times. This is not a new method but instead follows the work described in [127].

In practice, to perform this measurement, we break the detection window up into a large number of time bins. Say in the k th time bin, n_k photons are measured. There are two sets of probabilities we must consider:

- $P(n_k|b)$: the probability of measuring n_k photons given the ion is in a bright state
- $P(n_k|d)$: the probability of measuring n_k photons given the ion is in a dark state

In the absence of depumping and repumping, we can then write the overall probability that an ion is bright after i measurements $\{n_1, \dots, n_i\}$ as

$$P(b|\{n_1, \dots, n_i\}) = \frac{1}{N} P(b|\{n_1, \dots, n_{i-1}\}) P(n_i|b) \quad (5.8)$$

and correspondingly

$$P(d|\{n_1, \dots, n_i\}) = \frac{1}{N} P(d|\{n_1, \dots, n_{i-1}\}) P(n_i|d) / N \quad (5.9)$$

where N is a normalization factor. However, this assumes that if we had a bright ion in the i th time bin, this ion was also bright in the $(i-1)$ th time bin. This is usually true, but does not account for repumping or depumping. A depumping event during the k th time bin, with probability P_{DP} will take the ion from the bright state to the dark state. Correspondingly, a repumping event, with probability P_{RP} , will take the ion from dark to bright. We can use these to modify the probability that an ion is bright given measurement history $\{n_0, \dots, n_i\}$ as

$$P(b|\{n_0, \dots, n_i\}) = \frac{1}{N} ((1 - P_{DP})P(b|\{n_0, \dots, n_{i-1}\})P(n_i|b) + P_{RP}P(d|\{n_0, \dots, n_{i-1}\})P(n_i|b)) \quad (5.10)$$

We have here assumed that depumping and repumping events only happen at the beginning of every time bin, so this holds in the limit of small time bins. Similarly, accounting for the probability of depumping during the i th time bin, the dark state probability is modified as

$$P(d|\{n_0, \dots, n_i\}) = \frac{1}{N} ((1 - P_{RP})P(d|\{n_0, \dots, n_{i-1}\})P(n_i|d) + P_{DP}P(b|\{n_0, \dots, n_{i-1}\})P(n_i|d)) \quad (5.11)$$

To normalize, we require

$$P(b|\{n_0, \dots, n_i\}) + P(d|\{n_0, \dots, n_i\}) = 1 \quad (5.12)$$

These probabilities $\{P(n_i|b), P(n_i|d), P_{RP}, P_{DP}\}$ are dependent on the count rate from a bright ion, the dark count rate, and the depumping and repumping rates. I will discuss in Chapter 6 how we measure each of these rates, but for now, assume that these are measurable. Note that these measurements must also be *independent* - i.e. we cannot use the same dataset for calibration and for state measurement.

For high detection fidelity measurements, we keep track of continuously updating probabilities as described by Eqs. 5.10 and 5.11. After one of them hits a desired threshold P_{th} , the state is assigned.

To summarize: we break our measurement into a large number of time bins, where each bin contains some number of detector counts. This gives us updated information about the *true* state of the ion. Depending on the number of counts in a bin, we update our probabilities to indicate that it is more or less likely that the ion is in a bright state. If either $P(b)$ or $P(d)$ exceeds some threshold, we stop the measurement and assign the state; otherwise, we repeat with the next time bin. This method enables higher fidelity and shorter detection times than thresholding.

It is possible to use this method in real time in an experiment, provided the bright count, dark count, depumping, and repumping rates are well-calibrated [127]. However, this has not to date been performed on this experiment. We instead used this method in post-processing to determine the maximum detection fidelity achievable using the SNSPD; for most other experiments, we used thresholding, which is good enough for the proof-of-principle experiments I describe in this work.

5.2 Ion motional state measurement

For a thermal state of ion motion, a measurement of the mean phonon number can be efficiently made with only two measurements. In such a state, the probability P_n of the ion being found in the n -th harmonic oscillator state is given by

$$P_n = \frac{e^{-n\hbar\omega_{sec}/k_B T}}{1 + e^{-\hbar\omega_{sec}/k_B T}} \quad (5.13)$$

Following the derivation from [86] and plugging this into Equation 3.24,

$$I_{rsb}(\tau) = \sum_{n=1}^{\infty} \frac{e^{-n\hbar\omega_{sec}/k_B T}}{1 + e^{-\hbar\omega_{sec}/k_B T}} \sin^2(\Omega_{n,n-1}\tau) \quad (5.14)$$

If we reindex the sum, we see

$$\begin{aligned} I_{rsb}(\tau) &= \sum_{m=0}^{\infty} \frac{e^{-(m+1)\hbar\omega_{sec}/k_B T}}{1 + e^{-\hbar\omega_{sec}/k_B T}} \sin^2(\Omega_{m+1,m}\tau) \\ &= e^{-\hbar\omega_{sec}/k_B T} \sum_{m=0}^{\infty} \frac{e^{-m\hbar\omega_{sec}/k_B T}}{1 + e^{\hbar\omega_{sec}/k_B T}} \sin^2(\Omega_{m+1,m}\tau) \\ &= e^{-\hbar\omega_{sec}/k_B T} I_{bsb}(\tau) \end{aligned} \quad (5.15)$$

where I have substituted $I_{bsb}(t)$ from Equation 3.23. Note that, for phonons in a thermal distribution,

$$\bar{n} = \frac{1}{e^{\hbar\omega/k_B T} - 1} \quad (5.16)$$

[128]. So

$$I_{rsb}(\tau) = \frac{\bar{n}}{1 + \bar{n}} I_{bsb}(\tau) . \quad (5.17)$$

Thus the average phonon occupation of the harmonic oscillator can be straightforwardly extracted for a thermal state by measuring the ratio of the red and blue sidebands intensities, by

$$\bar{n} = \frac{1}{I_{bsb}(\tau)/I_{rsb}(\tau) - 1} \quad (5.18)$$

Notice that this ratio is independent of interrogation time τ .

In practice, we don't exactly measure what I have described. The lasers, discussed in Chapter 4 and in more detail in [47], can have some intensity dependence on frequency, primarily due to finite AOM bandwidth. In order to eliminate avoid this dependence, we compare bright and dark ion populations after the two following pulse sequences:

- (1) Microwave π -pulse, transferring population from $|\downarrow\rangle$ to $|\uparrow\rangle$; (2) Raman π -pulse at $\omega = \omega_{qu} - \omega_{sec}$, transferring population back to $|\downarrow\rangle$. Although the Raman laser beatnote is at the red sideband frequency, since we start in the upper qubit state, it acts as a “motion adding sideband” (MAS).
- (1) Raman π -pulse at $\omega = \omega_{qu} - \omega_{sec}$, transferring population to $|\uparrow\rangle$; (2) Microwave π -pulse, transferring population from $|\uparrow\rangle$ to $|\downarrow\rangle$. This is the traditional red sideband, which I will refer to as the “motion subtracting sideband” (MSS).

We then compare the populations in $|\downarrow\rangle$ for these two. The physics here is equivalent to comparing red and blue sideband pulses. However, we are able to keep the Raman laser detuned at the same frequency for both pulses, eliminating frequency-dependent intensity variations. An example of this measurement is shown in Figure 5.5 for a sideband cooled ion with average occupation $\bar{n} = 0.17(4)$.




Figure 5.5: Sample sideband cooling data, demonstrating a sideband imbalance. The motion adding sideband is in blue, and the motion subtracting sideband is in red. Error bars are smaller than the data points shown. The data is fit to a sinc^2 lineshape $P(f) = A \text{sinc}^2(w(f - f_0)) + P_0$. This gives amplitude $A_{MAS} = 0.71(2)$ for the motion adding sideband and $A_{MSS} = 0.10(2)$ for the motion-subtracting sideband, giving a measured temperature $\bar{n} = 0.17(4)$.

Also note that this measurement scheme assumes a thermal state of motion. In general, the ion may not be in such a state; particularly, sideband-cooled ions are *not* generally in a thermal state described by Eq. 5.13. In such a case, the ion temperature can be measured by taking a Fourier decomposition of carrier flops, since different Fock states have different Rabi rates. This is discussed in more detail in [89]. This is also useful for investigating the temperature of hotter ions, where the sideband imbalance I_{MAS}/I_{MSS} is close to 1. However, for the work in this thesis, I have assumed that the ion is in a thermal state with low average \bar{n} . This is a reasonable assumption based on the predicted Doppler-cooled motional mode occupation $n_{ax} \approx 9.2$ based on a typical secular frequency $\omega_{ax} = 2.0$ MHz for a mode oriented 45° degrees from the Doppler cooling laser.

5.3 Calibrations

This section describes the day-to-day methods of calibrating and tweaking up the experiment. I include it primarily as a reference for future users of the experiment, who may wish to be aware of some of its quirks.

5.3.1 Loading Ions

Ions are loaded in Zone 2 or 3, as labelled on Figure 4.8. One set of DC voltages that produces harmonic wells in the detector trap in each zone are listed in Appendix B. It is possible to load in the transport trap with the same set of voltages as those used in the detector trap with the “det” electrode grounded. The BD/BDD and PI beams are aligned over this zone with the frequencies set as listed in Table 4.5. Typical loading RF powers are 6.0 – 9.0 dBm. It is possible to keep an ion trapped with as low as 5.0 dBm RF power for long times, but it is challenging to load with such a low trapping potential. The ovens generate visible and IR light, which will generally swamp ion fluorescence. With a filter that blocks visible and IR light but transmits UV (Thorlabs FGUV11-UV) to reduce this background light, the ion’s fluorescence is visible on the EMCCD when it is trapped. At this point, the ovens and PI are turned off.

On a timescale of 1-3 months, the level of cryo-pumped gases in the pillbox slowly grows, either due to a slow leak or due to other parts of the apparatus outgassing. This doesn’t impact normal operation of the trap, as all gasses besides helium will adsorb to surfaces at 4 K. However, during loading, some surfaces are heated by the ovens and may release cryo-pumped gases. If enough gas is released in this process, we can see ions “pop in” to the trap and briefly fluoresce but not remain trapped. We attribute this to collisions with the background gas ejecting the ion. Since the 4 K pillbox is vented to the room temperature vacuum chamber, this background gas can be pumped up without warming to room temperature and breaking vacuum. The pillbox can be degassed by warm up the cold head to 25 K and pumping on the outer vacuum space with an ion pump. After this process, the trap reliably loads again upon cooling back to 4 K.

5.3.2 Stray fields

Ion loading and high-intensity UV scatter will generate stray electric fields. These stray fields can be measured and corrected for by adding “shims” potentials to the trapping potential. The shims are arrays of electrode voltages that are designed to adjust only one derivative of the trapping

potential without impacting any others. The generation of these shim potentials is discussed in Section 2.3.2.

Stray fields along either of the two radial directions will push the ion off of the RF potential null and introduce micromotion. Micromotion that overlaps with the cooling light will frequency modulate the cooling light in the frame of the ion, introducing sidebands at $f_{BD} \pm \Omega_{RF}/2\pi$. This effectively reduces the amount of light that is resonant with the cooling transition [129]. The micromotion can be minimized by applying a variable shim along \hat{y} (the radial direction in the plane of the trap) and maximizing ion fluorescence counts. The fluorescence counts vs shim magnitude are described by a Bessel function [129], so it is important to scan the shim magnitude broadly enough to confirm that the counts correspond to the maximum of the Bessel function and not simply a local maximum.

This method minimizes micromotion that overlaps with \hat{y} direction but does not address micromotion perpendicular to the trap in the \hat{z} direction, since that is not addressed by the BD laser. This requires a different technique. The microwave antenna introduces a small magnetic field gradient at the trap location. Micromotion in the plane perpendicular to the trap surface will modulate the amplitude of the magnetic field seen by the ion at the RF frequency $\Omega_{RF}/2\pi$. This will generate sidebands on the carrier microwave frequency at $(\omega_{qu} \pm \Omega_{RF})/2\pi$ [130]. When the micromotion is minimized, the amplitude of these sidebands will also be minimized. Therefore micromotion can be minimized by minimizing the Rabi frequency (or equivalently, by maximizing the duration of a Rabi π -pulse) for this sideband.

An axial stray field shifts the ion along the trap axis but does not generate any micromotion. Stray fields in the axial direction can be compensated for by changing the axial trapping potential. If an ion is at the center of a harmonic well, changes to the xx curvature of that well will not shift its equilibrium position. If it is pushed off-center by a stray electric field, its equilibrium position will shift if the curvature of the trapping potential is changed. An axial stray field can be seen by changing the curvature of the trapping well and watching the ion on the camera. If the ion is at the center of the well, it will not shift position. For more precise measurement, the readout

beam can be moved so that the ion is on the slope of the Gaussian beam profile. If the ion moves slightly here, the fluorescence will shift. To compensate for stray electric field along \hat{x} , the x shim is manually varied until changes in the trap xx curvature do not change the ion position.

The above methods of stray field compensation also serve as measurements of the stray electric field \vec{E}_s since the compensation shims will generate $-\vec{E}_s$. The magnitude of the electric field generated by a given shim amplitude is known from simulation. In the experiments discussed in this thesis, typical stray electric field magnitudes were $|E_{s,x}|, |E_{s,y}| \leq 300 \text{ V m}^{-1}$, $|E_z| \approx 750 \text{ V m}^{-1}$.

These stray fields are spatially varying, so they can cause a stray electric field curvature at the ion location, which will shift the trap secular frequencies. These stray curvatures can be compensated for by measuring the trap secular frequencies as described in Section 5.3.5. If the secular frequency ω_i is not at the desired values, it can be adjusted by applying an ii shim. In the experiments described in this thesis, only the axial frequency was controlled using this method.

The stray fields typically drift over the first 12-24 hours after loading an ion. We speculate that this is due to charging from the 235 nm light and from ions and electrons which end up on the trap surface during the loading process.

5.3.3 Polarization and magnetic fields

For best optical pumping, the polarization must be pure σ^- and the magnetic field must be well-aligned with the BD laser beam. Both of these can be optimized using the same process. This is performed by increasing the intensity of the BD detection beam (to increase the probability of depumping and repumping) and then adjusting both the BD beam polarization and current in the magnetic field offset coils to maximize the population in the bright state and minimize depumping. This is a multiparameter process and requires “walking” (a process in which one parameter is adjusted so that the state quality is degraded and a second parameter is scanned to find the best state quality. If the overall state quality was improved by this process, it was a “step” in the correct direction.).

5.3.4 Microwave pulse times and frequencies

In order to use microwave shelving (discussed in Section 5.1.3) to improve dark state fidelity, it is necessary to calibrate microwave π -pulse durations and frequencies on three transitions, all within the $^2S_{1/2}$ manifold: $|F = 2, m_F = -2\rangle \leftrightarrow |F = 1, m_F = -1\rangle$, $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = 0\rangle$, and $|F = 2, m_F = 0\rangle \leftrightarrow |F = 1, m_F = +1\rangle$. These can be tuned up by a combination of measurements: a time scan on a given transition with a constant microwave frequency, to calibrate the π -pulse duration; and a frequency scan with a constant π -pulse duration, to calibrate the transition frequency. This typically does not require frequent recalibration, but if conditions in the lab have changed frequencies may shift on the order of 10's of kHz.

5.3.5 Secular frequency measurement

Trap frequencies can be measured in two ways. The first is a “tickle” measurement. An oscillating electric field at the trap secular frequency will coherently displace the ion's secular motion. If this coherent displacement is sufficiently large, it will cause a dip in fluorescence counts. The secular frequencies can be measured by scanning the frequency of this “tickle” signal and fitting to the minimum of the fluorescence counts.

This oscillating field could be generated by applying an oscillating voltage to any trap electrode. However, the trap electrodes have low-pass filters that strongly attenuate signals oscillating the trap secular frequencies. The microwave antenna generates electric fields which overlap with all three secular modes of the ion and it is configured for the application of high-frequency voltage signals, as shown in Figure 4.10. Because of this, the tickle signal is combined with the microwave signal using a directional coupler and the microwave antenna is also used to apply the electric field for the tickle measurement.

The trap frequencies can also be measured by measuring the frequencies of Raman sidebands, discussed below. As discussed in Section 3.4.1, the Raman lasers can drive transitions both when the Raman beatnote is $\Delta = \omega_{qu}$ (carrier transitions) and when the beatnote is tuned to $\Delta = \omega_{qu} \pm \omega_i$

for the i th secular mode (sideband transitions). The trap secular frequencies can therefore be measured by preparing the ion in a bright state and applying a Raman pulse of an appropriate duration with a variable frequency.

5.3.6 Sideband cooling calibration

The first step in the calibration of sideband cooling is to find appropriate frequencies and π -pulse durations. As in the microwave pulse time and frequency calibration, these must be calibrated by the combination of two measurements: a time scan on a given transition with a constant Raman beatnote frequency, to calibrate the π -pulse duration; and a frequency scan with a constant π -pulse duration, to calibrate the Raman beatnote frequency. Once these values are known, the ion can be sideband cooled by the repeated application of motion-subtracting sidebands and repumping pulses, as discussed in Section 3.4.2

One detail that should be considered when calibrating the pulse time duration is that the Rabi frequencies $\Omega_{n,m}$ depend on the ion motional state [1]. When the ion is in a superposition of motional states, an effective π -pulse duration can be used that is a weighted average of the π -pulse durations for the individual states. However, this means that as the ion is cooled to a lower average motional state, this effective π -pulse duration will change. This should be considered and it may be desirable to have cooling pulse sequences with where subsequent pulses have longer durations. In these experiments, we observed a change the π -pulse duration after the ion was cooled, but did not reach a lower final \bar{n} by applying pulse sequences with variable pulse durations.

Chapter 6

Trap-integrated qubit state detection

Qubit state detection has been performed in ions in times as short as 11 μs [131] and with fidelity exceeding 99.99% [132], but these experiments have relied on components which may be challenging to scale to large numbers of ions. All previous high-fidelity results have been performed with bulky objectives with limited fields of view. The ratio of field of view to objective diameter is on the order of 0.004. In a large-scale quantum computer, it may be desirable to have simultaneous state detection of qubits spread over a trap active area that may be larger than the field of view achievable in a high-NA, high-magnification objective, as shown in Figure 6.1(b). In this chapter, I will discuss demonstration of a high-efficiency trap-integrated single-photon sensor. Larger, more complex traps could incorporate multiple of these detectors for parallel detection in spatially separated zones.

6.1 Maximizing detection efficiency

In order to perform fast, high fidelity qubit state detection, we wish to detect as many of the photons emitted by the ion as possible. I will define the overall detection efficiency η as the probability that a photon emitted by the ion is counted by the detector. For the purposes of this discussion, I will group the detection efficiency η into two components:

- Geometric efficiency (α): the percentage of emitted photons which are incident on our photon-counting sensor. This accounts for both solid angle over which photons are collected,




Figure 6.1: (a) A cartoon image of traditional fluorescence detection, in which ion fluorescence is collected by a large out-of-vacuum, high-NA objective and focused onto a high-quantum efficiency detector. The need for high-NA and high magnification are in conflict with a wide field of view, and this ultimately limits the number of ion qubits that can be imaged using such a scheme. (b) Scalable detection using on-chip single photon counting. In this scheme, nearby, trap-integrated single photon sensors convert photon detections into electrical signals on-chip, and the signal can leave the vacuum as an electrical signal. Here, the same numerical aperture can be achieved with a much smaller active area, since the detector is so close to the ion; this would enable readout from multiple neighboring ions in parallel. (b) adapted from [3].

the angular dependence of the fluorescence intensity, and other sources of loss such as fiber coupling inefficiency and surface reflections.

- Detector efficiency (β): the probability that a photon incident on a sensor will be counted by the detector.

Overall,

$$\eta = \alpha \times \beta . \quad (6.1)$$

Previous high-fidelity detection experiments [132][131] used custom imaging objectives with high numerical aperture (NA). These objectives are typically bulky and expensive. In this apparatus, we image using a simple 0.38 NA objective built from off-the-shelf parts, described in Section

4.2.9. An NA of 0.38 is equivalent to a solid angle of 0.45 steradians, meaning that this system would detect 3.6% of photons emitted from an unpolarized emitter. (The cycling transition which we detect on in ${}^9\text{Be}^+$ is, in fact *not* an unpolarized emitter. Since it is a σ^- transition, it emits photons in the “peanut” pattern shown in Figure 6.7, meaning that we detect an smaller percentage of photons from the ion than this estimate suggests.) Additionally, the light is transmitted through 12 optical surfaces before hitting the PMT, introducing further loss into our geometric efficiency.

Instead consider an ion trapped over the detector shown in Figure 4.9. The detector has an area of $20 \times 22 \mu\text{m}$. The ion is trapped a measured $29 \mu\text{m}$ above the electroplated gold surface; the detector is recessed an additional $6 \mu\text{m}$ from the electroplated gold surface, so the ion is $35 \mu\text{m}$ from the detector surface. This subtends a solid angle of 0.32 steradians, viewed from the ion position. Since there are no intermediate surfaces between the ion and the detector, the geometric efficiency is entirely accounted for by this solid angle.

6.1.1 Single photon detectors

Instead of improving detection efficiency by increasing geometric efficiency, one might focus on improvements to the sensor detector efficiency. Typical photon counting in our trapped ion experiments is performed by a Hamamatsu R7600P-203 photo-multiplier tube (PMT), with a quoted 38% quantum efficiency at 313nm. (This is integrated with the imaging system discussed in Section 4.2.9). To date, the best reported quantum efficiency of a single photon counter is 98% for 1550nm, measured using a superconducting nanowire single photon detector (SNSPD) [133].

High fidelity detection of trapped ion states have been made using high-NA imaging systems and sensitive single-photon counters in ${}^{40}\text{Ca}^+$ with 99.991(1)% fidelity [132] and in ${}^{171}\text{Yb}^+$ with 99.931(6)% fidelity [131].

6.1.2 Superconducting nanowire single photon detectors

SNSPDs are an appealing technology for qubit state detection, since they can count up to $\sim 10^7$ events per second, sufficient to handle trapped ion fluorescent rates. Additionally, SNSPDs

have been demonstrated with $> 90\%$ quantum efficiency at 1550 nm [134][135] and can have dark counts below 10^{-4} per second [136] which would be helpful for high-fidelity qubit state detection.

The principles of operation of an SNSPD are illustrated in Figure 6.2. Initially (i), a DC bias current I_{DC} is passed through a superconducting nanowire. Below the superconducting critical temperature T_C , the current sees no resistance and there will be no voltage drop. If, at a time $t = 0$, an incident photon hits the superconducting nanowire, it will create a local hotspot. The current will continue to follow the path of least resistance and will flow around the (non-superconducting, $R \neq 0$) hotspot. However, this will cause a local increase in current density. For a sufficiently high bias current I_{DC} , the current density in the nanowire will exceed the critical current density of the superconductor, causing a full “bridge” in which the superconductor goes normal across the entire width of the nanowire. The typical resistance of such a bridge is $k\Omega$ or more. We can view this as flipping open a switch, illustrated schematically in the lumped-element diagram in Figure 6.2(b). At this point, the current, instead of passing through the superconductor, passes through some external shunt resistance (as configured in this apparatus, simply the 50Ω impedance of the measurement circuit), creating a voltage drop. The timescale of this initial process is < 1 ns, making these sensors appropriate for experiments where photon arrival time information is critical. The timing jitter of SNSPDs can be made as low as ~ 3 ps [137].

With the current passing through the shunt impedance, the bridge rethermalizes and the nanowire goes superconducting again. The shunt resistance R and the superconductor kinetic inductance L form a parallel LR circuit. Once the nanowire is superconducting again, the bias current will re-enter the superconductor on a characteristic filter timescale $\tau = L/R$, and the output voltage on the shunt resistance will decay with this same timescale. A typical timescale is $\tau = 20 - 50$ ns, a reasonable timescale given typical trapped ion photon emission rates.

One question when considering any detector for ${}^9\text{Be}^+$ is what the detector efficiency is at 313 nm. Most of the demonstrations of SNSPDs have been performed at 1550 nm wavelengths, where the ubiquity of applications drives development. However, SNSPDs have been fabricated which operate with $> 70\%$ detection efficiency at 315 nm [3] and at 370 nm [63]. They have been

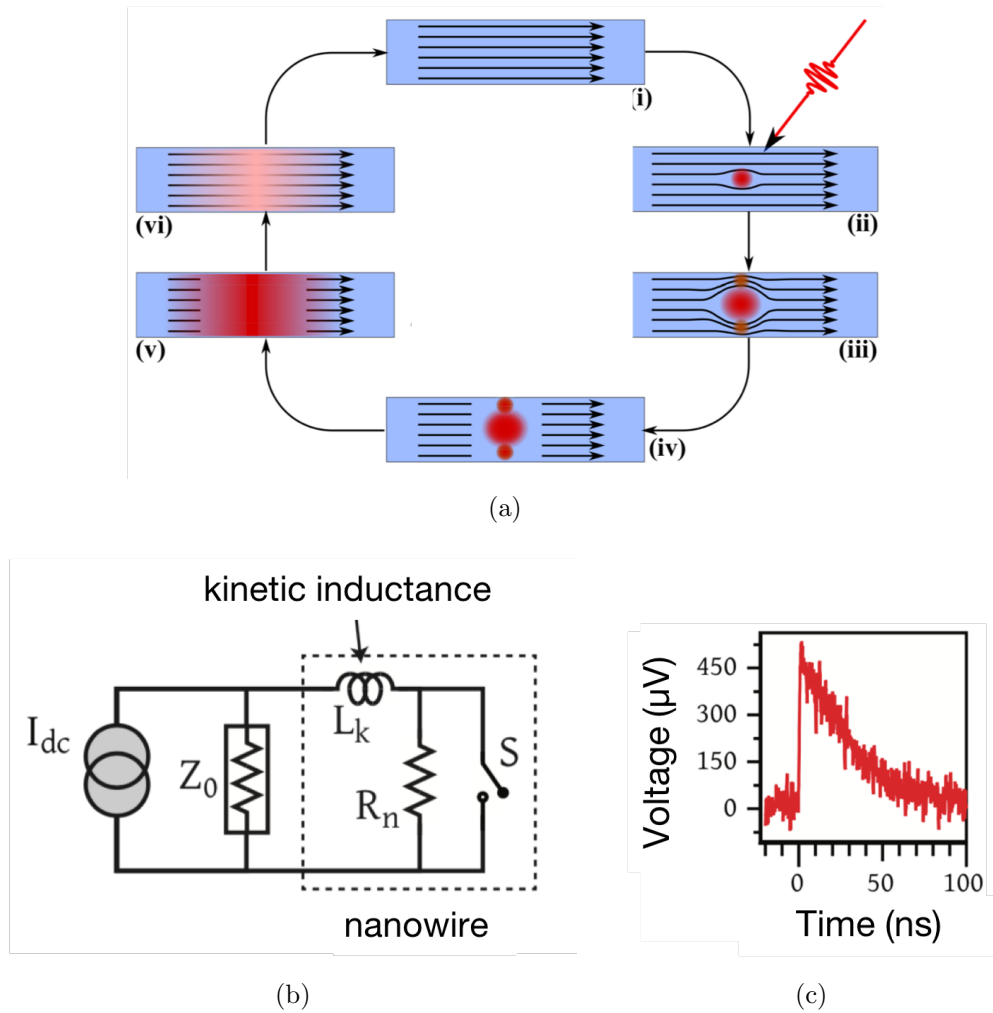


Figure 6.2: Principles of operation for a superconducting nanowire single photon detector (SNSPD). The SNSPD can be modeled as a lumped element circuit as shown in (b), where the nanowire detector consists of some kinetic inductance L_k and has some normal resistance R_n . (i) Below the superconducting critical temperature T_C , the nanowire can be modeled as a closed switch with zero resistance. A bias current I_{DC} will pass through it and experience no voltage drop. (ii) If a photon hits the nanowire, it will create a local hotspot where the superconductor goes normal. (iii) The current will bunch around the local hotspot, causing an increase in the local current density. (iv) For sufficiently high I_{DC} , the increase in the local current density will exceed the superconductor critical current, causing a full normal “bridge” in the nanowire. This can be modeled in the lumped-element picture in (b) as the switch suddenly opening. R_n is much larger than the $50\ \Omega$ impedance Z_0 , so the bias current will pass through the Z_0 , causing a measurable voltage spike. As the heat dissipates, the current will return back into the superconducting nanowire on a timescale given by the characteristic timescale of the circuit $\tau = L_k/Z_0$. A typical voltage pulse corresponding to a SNSPD photon click event is shown in (c). SNSPD operation schematic taken from [59].

used for high fidelity qubit state detection in $^{171}\text{Yb}^+$ at 369.5 nm [131].

SNSPDs introduce one major overhead challenge, which is the need for a cryostat. In order to operate, an SNSPD must be below the superconducting transition temperature of the nanowire. However, as discussed in Chapter 4, there are multiple independent reasons to operate trapped ion experiments cryogenically. The apparatus I built during my Ph.D. is a good platform for exploring the integration of SNSPDs with trapped-ion experiments, since it has a low base temperature compared to other trapped-ion cryostats.

6.2 Trap-integrated detection

The goal of this experiment is not only to demonstrate high fidelity state detection but to do so in a potentially scalable way that enables parallel readout from multiple ions. Other groups have proposed and fabricated systems which use microfabricated optical components such as microfabricated Fresnel lenses [138], diffractive lenses [38], or trap-integrated mirrors [33][39] for photon collection. However, these still require imaging using a remote, high-quantum efficiency sensor, and require the fluorescence photons pass through multiple optical surfaces. Our approach is to instead use a single photon sensor directly below the ion, and to collect a large enough fraction of photons simply by being very close. The photon count information will then exit the system as an electronic signal.

There are many reasons to push towards microfabricated surface-electrode traps, as discussed in Section 2.3. By using the same technologies integrate this detector into the ion trap itself, we can include more of the ion trapping system in a robust, compact package. However, this introduces new challenges; chiefly, that the detector must be able to operate in the high-RF environment of an ion trap, where there will be non-negligible coupling between the RF electrode and the detector. We must also consider the effect of the detector on the ion. Electric field noise will cause ion heating, so it is critical that this detector generate sufficiently low electric field noise that the ion can still be well-enough-controlled for quantum computing applications.

There are several major questions we wish to answer about a trap-integrated single photon

detector:

- What effect does the RF trap drive have on the SNSPD performance?
- What is the detector efficiency of a trap-integrated SNSPD? Do the further fabrication steps required to fabricate both an SNSPD and an ion trap on the same chip degrade the SNSPD performance?
- What is the spatial sensitivity of the SNSPD? We envision a scalable setup in which neighboring ions each have individual SNSPDs for parallel readout. By how far would these ions have to be separated so that the crosstalk between neighboring ion-detector pairs is acceptable?
- What is the effect of the SNSPD on the ion? Does it introduce unacceptable electric field noise for high-fidelity quantum computing applications?

I will answer all of these questions in this chapter, and demonstrate that trap-integrated SNSPDs are promising technologies for trapped-ion quantum computing applications and high-fidelity state detection.

6.2.1 Previous work

Investigations into trap-integrated detectors to date have been limited to just a few research groups. As a proof-of-principle experiment, a cloud of $^{88}\text{Sr}^+$ ions was been trapped over a transparent indium tin-oxide (ITO) substrate with a silicon diode underneath operating for detection [41]. Fluorescence from $^{24}\text{Mg}^+$ has also been observed using a multimode fiber tip located directly below a hole in a trap electrode [36]. Other work has been done to integrate avalanche photodiodes (APDs) into trap chip geometries, with the benefit that every step of this process can be performed in a commercial CMOS chip foundry [32]. However, initial results in this system had a large, time-varying noise floor and a low overall signal-to-noise ratio [139]. Traps with an integrated SNSPD have been fabricated, and photons were counted in the presence of trap-like RF amplitudes and

frequencies [3]. However, prior to this work, ions had never been successfully trapped in a trap with an integrated SNSPD. Further, high fidelity qubit state detection had never been performed in a system with a trap-integrated detector.

6.3 Hardware for SNSPD readout

I here report results collected in the so-called “detector trap,” discussed in Chapter 4, fabricated at NIST using a similar process to that discussed in [3]. This trap is a five-wire style surface electrode trap in which the ion can be trapped anywhere along the trap axis. An SNSPD made from a MoSi superconducting alloy is fabricated at one end. With the exception of the detector itself, all electrodes are made of electroplated gold 6 μm thick. This allows ions to be loaded at one end (the “load zone”) with no line-of-sight to the detector. We load ions on the side of the trap that is furthest from the SNSPD. By applying time-varying potentials to the trap DC electrodes, the ion can be shuttled from the load zone to anywhere else in the trap, including over the detector. In either the load zone or directly over the detector, the Doppler-cooled lifetime of ions in the trap is days.

The ion is trapped 39 μm above the electroplated gold trap surface. However, close to the detector, the pseudopotential null moves closer to the surface, dropping to 29 μm above the electroplated gold when centered over the detector. This represents a distance of 35 μm from the SNSPD, which is recessed relative to the thick electroplated gold. This drop is illustrated in Figure 6.3. We attribute this to a combination of end-effects from the end of the RF electrodes and failures of the Biot-Savart analog I discussed in Chapter 2. The Biot-Savart analog assumes all electrodes are in a two-dimensional plane [73], which is generally a good approximation, but fails over the detector, where the difference in electrode height between the SNSPD and the electroplated gold electrodes is large on the scale of the ion height.

The vacuum and ion trapping apparatus were discussed in Chapter 4. However, there is additional hardware that is specific to the operation of the SNSPD.




Figure 6.3: Schematic showing ion position over the detector. Due to end-effects of the end of the RF electrode and due to breakdowns in the assumption that the trap can be approximated as a two-dimensional surface, the ion moves $\sim 10 \mu\text{m}$ closer to the trap surface over the detector.

6.3.1 Readout circuit

The current supply and readout circuit is shown in Figure 6.4. For SNSPD operation, there must be some bias current running through the detector nanowire. This bias is typically on the order of a few μA . We wish to avoid the detector being biased at all times, since detector pulses may introduce electric field noise at the ion. We drive the detector bias using a pulse generated by an arbitrary waveform generator; this waveform is triggered by a TTL pulse from the ARTIQ experiment control system. This waveform is ramped on and off over $11 \mu\text{s}$, which limits ringing in the bias tee and capacitive coupling to other electrodes. This waveform is attenuated by 20 dB and then passes through a $10 \text{ k}\Omega$ resistor. This gives a supply voltage to bias current ratio of $100 \text{ mV} : 1.05 \mu\text{A}$.

The detector current supply and the readout circuit are connected to the detector via the same cryostat coaxial line. The bias current is a DC signal and the detector pulse is a quickly varying waveform with many high frequency components. This means the DC bias supply can be coupled into the detection circuit using a bias tee, as is shown on Figure 6.4. The readout circuit is connected to the detector by this bias tee; the other end of the detector is grounded. If there is a detected photon, the short to ground of the superconducting nanowire will “open” and the input side of the detector will quickly go to $V = I_{bias} \times 50 \Omega$. We can see this pulse this on the output

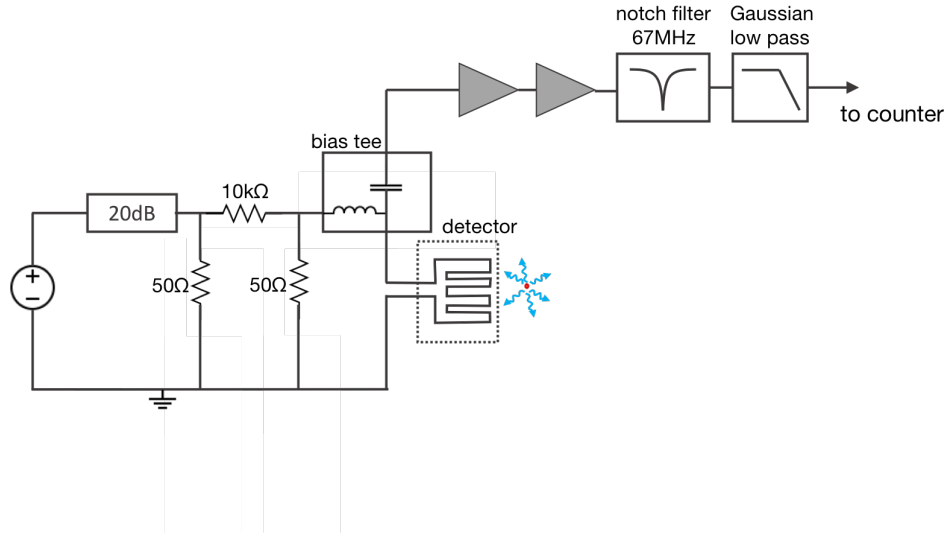


Figure 6.4: SNSPD readout circuit. The DC bias current and the readout signal are connected to the same cryostat port by a bias tee. The bias current is supplied by an AWG providing a pulse with 100 – 500 mV amplitude. This signal is attenuated by 20 dB and passes through a 10 k Ω resistor that converts the voltage bias into a current bias. It then passes through the DC port of the bias tee to reach the detector. We read out the detector signal at the high frequency port of the bias tee. It is amplified by two stages of amplification and then filtered with a notch filter at the trap RF frequency plus a second low-pass filter with a corner frequency of 25 MHz before going to a discriminator to be counted.

of the bias tee.

The SNSPD output pulse is small on the scale of the TTL logic we want it to eventually be converted to (< 1 mV vs 1–5 V), so we pass the detector output through two stages of amplification for a total 46 dB gain. Before amplification, the SNSPD output pulse is on the order of 250 μ V peak. This signal is small compared to the RF pickup on the SNSPD output. However, this RF pickup is at a specific, known frequency and can be removed using a narrow band notch filter with 40 dB attenuation at the RF frequency. For improved signal processing, we also apply a Gaussian low-pass filter with a corner frequency $f = 25$ MHz. This amplified, filtered signal is then sent to a discriminator where it is converted to a TTL signal that can be counted by the experimental control software. Filtered SNSPD pulses are shown in Figure 6.5. The amplitude of the pulses varies from pulse to pulse. This is because, due to RF pickup on the nanowire, the bias current in the nanowire will vary with the phase of the RF drive at the time of photon arrival.




Figure 6.5: Sample SNSPD output pulses. The height of the pulses is given by the bias current by $V_{pulse} = i_{bias} \times 50 \Omega$. Because of RF pickup on the SNSPD, the bias current I_{bias} is time dependent as $I_{bias}(t) = I_{DC} + I_{pickup} \cos(\Omega_{RF}t + \phi_{RF})$, so the height of the voltage pulse depends on the phase of the RF drive when the photon hits the detector.

6.4 Trap-integrated SNSPD performance

6.4.1 RF sensitivity

One of the main questions about the trap-integrated SNSPD is the effect of the RF drive. Paul traps require an RF drive to trap the ions, typically on the scale of 10 V – 100 V amplitude at frequencies of 10 – 100 MHz. Since the RF electrodes and the SNSPD are near each other, there is capacitive coupling between the two. We filter the RF pickup out of the output signal. However, the corresponding induced RF currents affect the detector efficiency.

The impact of the RF is shown in Figure 6.6. We plot four different curves showing the count rate of the SNSPD as a function of applied bias current, each for a different RF amplitude. The SNSPD is not illuminated by an ion (since in all but one case, the RF amplitude is too low to trap an ion) but instead by “ion-like” light levels. This was achieved by moving the detection laser closer to the trap surface until the stray scatter caused as many counts in a detection interval as a typical ion at around saturation parameter $s = 1/3$.

The blue curve shows performance with the RF off. Under these conditions, we see the classic

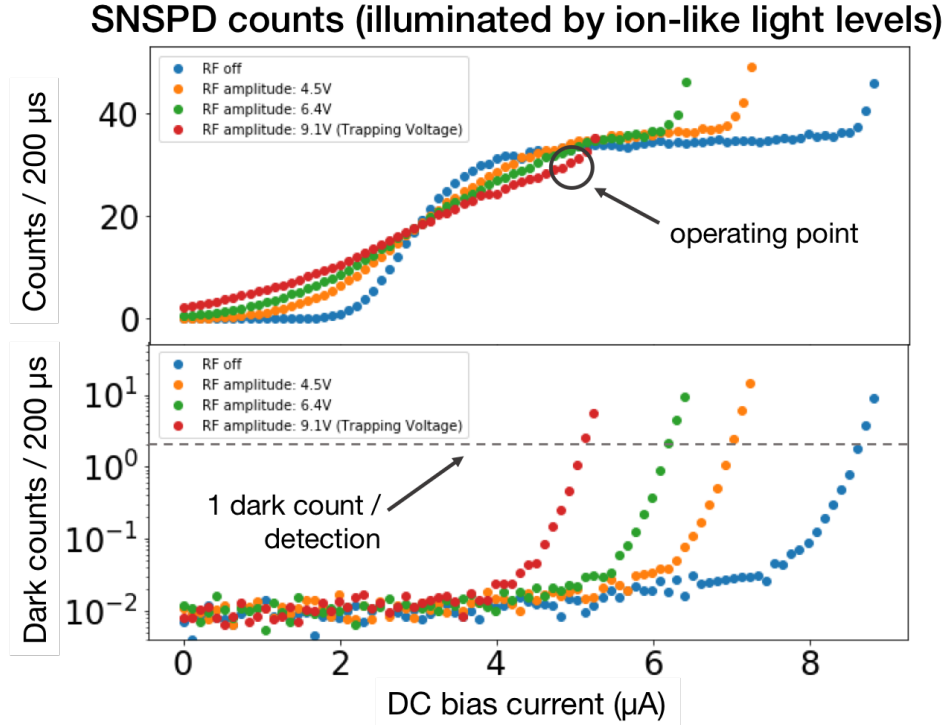


Figure 6.6: SNSPD click rate vs DC bias current for a variety of RF amplitude levels. It is impossible to trap a ${}^9\text{Be}^+$ ion at all but the highest RF level, so for a fair comparison, bright counts were created by adjusting the detection beam to a height at which the scattered counts in a single 200 μs detection period were similar to those produced by an ion in the bright state at around saturation parameter $s = 1/3$. Counts vs. bias current for the RF amplitude used for ion trapping are shown in red. We operate at a bias current point level that gives about 1 dark count per detection period. At higher RF amplitudes than our trapping RF, with peak voltage $V_{RF} = 9.1$ V, the detector did not operate at all.

SNSPD response. At lower bias current levels, there are no counts at all. Even if a photon hits the nanowire, the increase in current density caused by the bias current bunching around the hot spot is still low enough that the detector does not go normal across the full width of the nanowire. Then there is some “turn-on” range, here between 2 – 4 μA , in which whether the photon is counted or not depends on where the photon hits the nanowire. At high bias current, the detector exhibits a “plateau,” where the count rate is largely independent of bias current. Finally, as the bias current approaches the nanowire critical current, the dark counts start to rise. At higher bias current than the critical current, the detector is no longer superconducting.

As the RF amplitude is turned up, we can see two effects. First, the bias current that can

be applied before the SNSPD goes normal decreases with increasing RF amplitude. Second, the slope of detector count rate vs. bias current in the “turn-on” range decreases with increasing bias current. We attribute this to the effect of RF pickup on the nanowires. Figure 6.6 shows bright counts vs. applied DC bias current; however, the local bias current “seen” by the SNSPD is

$$I_{bias}(t) = I_{DC} + I_{pickup} \cos(\Omega_{RF}t + \phi_{RF}) \quad (6.2)$$

Depending on the phase of the RF pickup when a photon hits the nanowire, the actual bias current thus may be larger or smaller than I_{DC} by as much as I_{pickup} .

In previous samples[3], it was possible to cancel this effect by applying a cancellation tone directed down the SNSPD output line at the RF frequency $\Omega_{RF}/2\pi$ with a suitable phase. However, this did not reduce the effect of the RF pickup in this sample. We attribute this to the fact that there is increased coupling between the RF electrode and the SNSPD meander in this sample. Based on simulations, RF pickup coupled onto the SNSPD leads will have the same effect on the current in the nanowire as a cancellation tone applied to the SNSPD output line and can therefore be cancelled coherently. However, RF pickup due to capacitive coupling to the nanowire itself will induce currents in the nanowire with spatially varying amplitude and phase, which cannot be cancelled by a cancellation tone applied to the SNSPD output line.

The count rate vs bias current for the operating RF amplitude at which we are able to trap an ion is shown in red in Figure 6.6. I emphasize two features of this curve:

- (1) The maximum useable DC bias is sufficiently reduced that the current limit is reached before the count rate reaches the “plateau” of the maximum count rate.
- (2) There is a nonzero count rate even with no applied DC bias current, so this detector cannot be truly “turned off” when the RF is at a trappable level.

Although these will both affect the performance of the SNSPD, it measures at a high enough count rate for ion state detection. We operate at a bias current level that produces ~ 1 dark count per detection. At this bias current, the bright counts are 83% of the maximum no-RF count rate.




Figure 6.7: Photon emission pattern for a σ^- transition. The quantization axis is shown in the black arrow. Notice that perpendicular to the quantization axis, where we detect fluorescence photons, the probability of photon emission is lower than along the quantization axis.

6.4.2 Detector efficiency measurement

Typically, detector efficiency measurements are made using well-calibrated light sources. A saturated single ion serves as an high quality self-calibrating light source. The scattering rate of an ion in a laser field is given by

$$\gamma_p = \frac{s_0\gamma/2}{1 + s_0 + (2\delta/\gamma)^2} \quad (6.3)$$

where s_0 is the saturation parameter, γ is the transition linewidth, and δ is the detuning of the laser frequency from the transition frequency [84]. For high saturation parameters s_0 , this approaches

$$\gamma_p \rightarrow \frac{\gamma}{2} \quad (6.4)$$

The linewidth of the detection transition $^2S_{1/2}|F = 2, m_F = -2\rangle \leftrightarrow ^2P_{3/2}|F = 3, m_F = -3\rangle$ is $\gamma = 2\pi \times 19.4$ MHz [82].

The measured height of the ion above the electroplated surface is 29 μm , and the thin SNSPD wires are further recessed from the surface of the electroplated electrodes by 6 μm . The detector has dimensions 20×22 μm , so this gives an effective solid angle of 0.32 steradians. To know the absolute number of photons scattered on the detector, we must also take into account the non-uniform scattering pattern for a σ^- transition (illustrated in Figure 6.7). Taking both of these factors into account, 2% of photons scattered by the ion will be incident on the detector.

These calculations give the absolute number of photons that are incident on the detector in a 200 μs detection period for a saturated ion, with the leading source of error being the uncertainty

in the measurement of the ion height over the detector (approximately $\pm 1 \mu\text{m}$). We can then determine the detector efficiency by measuring the saturated detector count rate. One might naïvely measure this by scanning the laser intensity and measuring the detector counts; however, this will also cause an increase in background counts due to laser scatter which will grow linearly with the laser intensity. We instead prepare the ion in

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

by applying a $\pi/2$ -pulse before detection and measuring full histograms at each laser intensity, which will give a double Poisson distribution of observed counts. The difference between the Poissonian means $n_{\text{bright}} - n_{\text{dark}}$ is purely due to photons scattered from a bright ion. These counts $n_{\text{bright}} - n_{\text{dark}}$ are plotted as a function of detection laser power in Figure 6.8 and level off with increasing laser intensity, as would be expected for a saturated ion. From this, we can fit to determine the detector count rate due to photons scattering from a saturated ion. Since we know from the geometry and the transition linewidth the rate at which photons are incident on the SNSPD, we can extract a quantum efficiency of 0.44(2) at our operating RF level of 9.1 V. The error is dominated by the $1 \mu\text{m}$ uncertainty in the ion height.

6.4.3 Spatial sensitivity

The long-term vision for this technology involves traps with multiple integrated SNSPDs. We first need to be confident that detector counts come from an ion directly over the detector instead of from a neighboring ion. To quantify the sensitivity of the detector to fluorescence from ions in different trap zones, we measured the saturated count rate for an ion in various locations along the trap axis relative to the detector. By measuring the saturated count rate, we can ensure that imperfect beam alignment on the ion does not impact the measurement. The saturated count rate as a function of well position is shown in Figure 6.9(b), normalized to the saturated count rate of an ion centered on the detector. An ion $80 \mu\text{m}$ from the center of the detector (less than the width of two trap electrodes) counts photons at less than 2% of the rate of an ion on the center of the

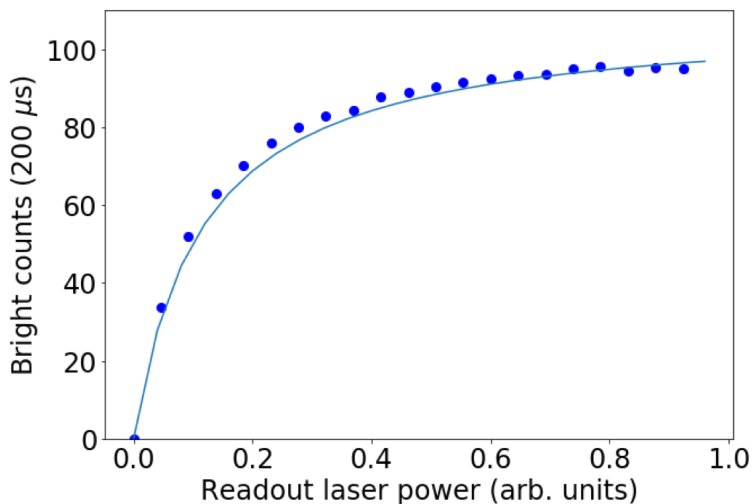


Figure 6.8: Detector bright count rate vs laser power. In order to cancel out the effect of increasing dark counts due to laser scatter, this was measured by preparing an ion in a superposition of the bright and dark states and measuring the difference between the bright and dark Poissonian means. By fitting the data to Eq. 6.3, we calculate a detector efficiency of 44(2)%.




Figure 6.9: Normalized saturated ion counts as a function of ion distance from the detector center. The colors of the points on (b) match the position of the ion shown on (a). With the ion translated less than 80 μm from the detector center, a distance of less than two electrode widths, the detector count rate is less than 2% of the count rate from an ion over the center.

detector. This would have only minor impacts on state identification using either a thresholded or adaptive Bayesian readout method.




Figure 6.10: Simultaneous $200\mu\text{s}$ detection of an ion prepared in $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$ using the SNSPD and the PMT with the traditional imaging system. Threshold values for each histogram are indicated as dashed lines. The SNSPD bright count average n_{bright} is 2.5 times larger than that of the PMT, allowing for improved state discrimination in a shorter time.

6.5 High fidelity state detection using a trap-integrated SNSPD

Using the ARTIQ experimental control system, we can count photons during the same detection period using both the SNSPD and the PMT integrated with the external bulk imaging system. This enables apples-to-apples comparisons of the two detection methods, since all calibrations - beam alignment, micromotion, polarization, etc - will be identical. I show photon count histograms for an ion prepared in $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$ in Figure 6.10. In the same detection period, the SNSPD counts about 2.5 times as many bright counts than the PMT does. A quick visible inspection of this histograms indicates that the fact SNSPD should have improved detection fidelity, since the bright and dark state histograms are further spread apart in the same detection time. I will formalize this instinct and give a detection fidelity number for the SNSPD.

6.5.1 Post-selection

Recall from Chapter 5 that detection fidelity is defined

$$\mathcal{F} = 1 - \frac{\epsilon_B + \epsilon_D}{2} \quad (6.6)$$

where ϵ_B , is the bright state error and ϵ_D is the dark state error. Typically, we would measure ϵ_B by preparing an ion in a pure bright state and counting the number of runs that are mislabelled as dark. Similarly, ϵ_D is measured by preparing an ion in a pure dark state and counting the number of runs that are mislabelled as bright. In systems in which the state preparation error is small compared to the detection error, this works well. Our experiment shows state preparation errors at the $\sim 0.5\%$ level, dominated by errors from imperfect optical pumping. These are larger than the measurement errors we seek to characterize. In order to get a clean estimate of our detection fidelity, we use post-selection to improve the state preparation fidelity.

Since we have timestamps for all photon arrivals, we can choose the time interval used for qubit state detection in post processing. For post-selected state preparation, we consider only the first 50 μs of detection. If in that first 50 μs , we measure a number of counts at or above a certain threshold n_b , we are confident that the ion has been prepared in the bright state. If it at or is below a different threshold n_d , we are confident that the ion has been prepared in the dark state. If the number of counts is between the two thresholds, we do not use this trial in further analysis. In this data, we used $(n_b, n_d) = (8, 0)$. We then consider the detection time to start at the end of this 50 μs heralding period, and do not use any data from the heralding period in the subsequent analysis of readout fidelity. This technique gives improved state preparation fidelity, allowing for more accurate characterization of readout fidelity. The bright and dark state histograms collected using this method are shown in Figure 6.11(a) for a detection time of 125 μs .

6.5.2 Calibrations for high fidelity detection measurements

The detection fidelity was evaluated using the protocol from [132] described in Section 5.1.4. We require four independently measured rates for high fidelity state detection:

- $P(n_i|b)$: the probability that n_i counts will be measured from a bright ion during a given time bin
- $P(n_i|d)$: the probability that n_i counts will be measured from a dark ion during a given

time bin

- P_{RP} : the probability of a repumping event during this bin
- P_{DP} : the probability of a depumping event during this bin

All of these rates are per time bin, where the time bin is the small step used to generate the detection fidelity as discussed in Section 5.1.4. In the following discussion, we used 1 μ s time bins. We must use a different set of data to calibrate these rates than the data set that we will use to measure the detection fidelity.

6.5.2.1 Depumping and repumping rates

Depumping and repumping rates were calibrated using data from the conventional imaging system and PMT. This data was collected simultaneously with the SNSPD data. Since we have photon arrival time data, these “experiments” can be performed in post-processing.

The depumping rate can be measured by preparing an ion in the bright state and waiting a variable time. The rate at which the average counts decrease is related to the depumping rate. We pick a time $t_{cal} > 0$ and investigate the histograms at this time. We measure the populations using thresholding to measure $P_B(t = t_{cal})$ and $P_D(t = t_{cal})$. We then can scale the measurements of bright counts and dark counts versus time to find $P_B(t)$. The slope of $P_B(t)$ is exactly the depumping rate per unit time, assuming the rate is small enough that we can approximate an exponential decay as a line.

The method for finding the repumping rate is basically the same. We measure counts after preparing an ion in the dark state and waiting a variable time. We measure the bright and dark populations using thresholding at some $t_{cal} > 0$ to obtain a scaling factor to convert counts into $P_D(t)$. The slope of $P_D(t)$ is the repumping rate per unit time.

In previous work, the final high fidelity results were relatively insensitive to the exact repumping and depumping rates [132] and that holds in this case as well. Varying these rates by

10% in the detection fidelity calculation below gave only minor corrections to the final detection fidelity, well within the statistical uncertainty.

6.5.2.2 Bright and dark measurement rates

We also need to independently measure $P(n_i|b)$ and $P(n_i|d)$, the probabilities that a bright and dark ion, respectively, will produce n_i detection events in a given time bin. These will each be Poisson distributed, so we simply need to find the mean of the respective distributions as a function of time to extract a count rate. We cannot use the same dataset that we use for determining the readout fidelity. However, recall that in the post-selection process, 21% of the nominally “bright” runs and 48% of the nominally “dark” runs were dropped. These unused data sets can be used to calculate the bright and dark Poisson parameters needed. Our goal is to find $\lambda_B(t_d)$ and $\lambda_D(t_d)$, the Poissonian parameters for the bright and dark states as a function of integrated detection time. We can build a histogram for every t_d and fit to a Poissonian to obtain $\lambda_B(t_d)$ and $\lambda_D(t_d)$. Then

$$P(n_i|b) = \frac{\lambda_B^2(t_{bin})e^{-\lambda_B(t_{bin})}}{n_i!} \quad (6.7)$$

and correspondingly,

$$P(n_i|d) = \frac{\lambda_D^2(t_{bin})e^{-\lambda_D(t_{bin})}}{n_i!}, \quad (6.8)$$

where $\lambda_B(t_{bin})$ and $\lambda_D(t_{bin})$ are obtained by fitting to the extracted slopes of $\lambda_B(t_d)$ and $\lambda_D(t_d)$, which are linear in detection time.

6.5.3 SNSPD readout fidelity

Eqs. 5.10 and 5.11 give $P(b|\{n_0, \dots, n_i\})$ and $P(d|\{n_0, \dots, n_i\})$ where $\{n_0, \dots, n_i\}$ is a sequence of the numbers of clicks in each of i time bins. These probabilities depend on the probabilities of depumping and repumping P_{DP} and P_{RP} and on the probabilities $P(n_i|b)$ and $P(n_i|d)$, which can be measured as described above. In general, one would calculate $P(b|\{n_0, \dots, n_i\})$ and $P(d|\{n_0, \dots, n_i\})$ for an increasing number of bins until one of them reaches a desired error threshold, at which point the calculation would terminate and the state would be declared to be determined. Using this




Figure 6.11: (a) Histograms generated from the bright- and dark-ion preparations after post-selection, with a 125 μs detection time. (b) Detection fidelity as a function of readout time for both a threshold method and adaptive Bayesian scheme, following the scheme described in [132]. Thresholding gives a maximum detection fidelity of $\mathcal{F} = 99.88(1)\%$, while the adaptive Bayesian method gives $\mathcal{F} = 99.91(1)\%$ with approximately 1/3 the average detection time

method, we can calculate bright state and dark state errors as described in Section 6.5.1. The infidelity as a function of time is shown in Figure 6.11(b). This method reaches a higher fidelity than the threshold method in a shorter time. We ultimately achieve $\mathcal{F} = 99.88(1)\%$ fidelity in 125 μs using thresholding and $\mathcal{F} = 99.91(1)\%$ fidelity in 46 μs using the adaptive Bayesian method.

6.6 Ion heating due to the SNSPD

One of the major questions about this detector is what effect it will have on the ion. Traditional detection involves a sensor that is very distant from the ion. Here we discuss a detector situated just one ion-electrode distance away. Further, the detection process involves the production of a fast voltage pulse which may heat or excite the motion of the ion. We are not concerned with heating during the detection process, as the detection later will be the dominant source of heating or cooling during readout. However, if other laser-based operations, such as stimulated Raman transitions, induce an elevated heating rate, this will limit our ability to perform interesting quantum algorithms over the detector.

The average motional occupation of an ion at the Doppler limit depends on the frequency as

$$\omega \left(\bar{n} + \frac{1}{2} \right) |\hat{x}_i \cdot \hat{k}_{Dop}|^2 = \frac{\gamma}{2} \quad (6.9)$$




Figure 6.12: Heating rate as a function of the rate of detector events during the delay. The trend is roughly linear, indicating a rate of 0.010(4) quanta of heating per detector event.

[84]. By working at a high enough trap secular frequency, we can start at a lower average motional occupation that gives the measurable sideband imbalance we need for a heating rate measurement. For heating rate measurements over the detector, we Doppler cooled to approximately $\bar{n} = 2$ in a trap with axial frequency $\omega_{ax} = 5.34$ MHz and a 45° angle between the cooling beam and the axial mode.

The heating rate in the conventional way: cool the ion motion, wait for a variable delay time, and measure the the final temperature using a motion-adding/motion-subtracting sideband imbalance. By fitting \bar{n} versus time to a linear fit, a heating rate was measured to be 108(14) quanta/s. This corresponds to an electric field noise level $\omega S_E(\omega) = 1.56 \times 10^{-5} \text{ V}^2/\text{m}^2$, consistent with other cryogenic heating rates given the trap with an ion-electrode distance of $35 \mu\text{m}$, as shown in Figure 4.27.

However, as I have mentioned, we were not able to sideband cool to the motional ground state while the ion was over the detector. We attribute this to heating due to detector “clicks.” Since the RF pickup in the nanowire is such that the detector cannot truly be turned off, any stray light incident on the detector will cause detection events, albeit with reduced efficiency relative to when an external bias current is applied. These events are fast voltage spikes, so they should

provide fast delta-function impulses to the ion motion, heating it. To test this, we turned on an off-resonant Raman beam during the variable delay in a heating rate measurement. Although we tightly focus all laser beams, there is still some residual stray light scattered off the trap; this residual stray light caused an elevated rate of detector counts during the delay, arising from residual RF current in the SNSPD. The heating rate is plotted as a function of this rate of detector counts in Figure 6.12. Based on a linear fit, we estimate $0.010(4)$ quanta of heating per detector click. This is a manageable heating level when working with relatively dim beams as are needed for Doppler cooling and state detection. However, it becomes more challenging when working with high-intensity Raman beams.

The highest intensity off-resonant Raman beam I used in this measurement generated 21 clicks per millisecond. However, this beam intensity is 300-1000 times reduced in intensity from the powers that were typically used for Raman sideband cooling in this experiment. A heating rate of 3-10 quanta per millisecond during sideband cooling pulses likely explains why we were unable to sideband cool over the detector. The axial mode was sideband cooled using 25 $10\mu\text{s}$ pulses, which would heat the ion 0.75-2.5 quanta. We would likely also be unable to perform laser-mediated two-qubit gates over the detector, since these also require high-intensity beams.

6.7 Outlook

This is a promising first demonstration of a trap-integrated single photon detector, showing not only that an on-chip detector continues to operate but that it does so at a level capable of producing high-fidelity qubit state detection. To return to the four open questions that I identified previously:

- *What effect does the RF trap drive have on the SNSPD performance?* It degrades the performance, but not to a level where the SNSPD fails to operate.
- *What is the quantum efficiency of a trap-integrated SNSPD?* We measured this as $0.44(2)$, a reduction from reported efficiencies in the absence of RF, but an improvement over the

best commercially available PMTs.

- *What is the spatial sensitivity of the SNSPD?* The SNSPD count rate is low enough from ions in neighboring zones that high fidelity parallel readout is not precluded.
- *What is the effect of the SNSPD on the ion?* The SNSPD *does* cause an elevated heating rate, but only in the presense stray laser light. This means we are unable to Raman sideband cool or perform laser gates over the detector, but the heating rate is similar to that of the rest of the trap when the light is turned off.

Still, a few issues remain that must be addressed before this technology can be broadly adopted for quantum computing applications. In particular, the low operating temperature and sensitivity to on-chip RF would make this a challenging technology for heavier ions than ${}^9\text{Be}^+$, which require higher RF trapping amplitudes that also typically heat trap chip more. However, for systems using ${}^9\text{Be}^+$, the temperature and RF pickup can be managed to a level appropriate for trapped ion quantum computing applications. The elevated heating rate due to laser scatter is a more concerning issue, as cooling to near the ground motional state is a necessary component of all high-fidelity gates demonstrated to date [16][15]. There are a few approaches that future experiments could take to address this issue. One approach might be to explore working with a smaller detector, in which we may be successfully able apply an RF cancellation tone [3]. This would let us reduce the induced RF currents to prevent detector counts from occurring when the external bias current is turned off. Another option would be to add a layer of an electrically conductive but transparent thin film to a future device to shield the ion from the electric fields generated by the SNSPD.

Another option is to operate the SNSPD in an architecture where there is no need to apply Raman beams over the SNSPD. The quantum CCD offers a clear solution. In a QCCD-type device, the SNSPD could be operated in a detection zone, far from sideband cooling, two qubit gates, or any other operations requiring high-intensity lasers.

Chapter 7

Cold, fast ion transport

Two-qubit gates for ions in a shared harmonic well can be performed on few microsecond timescales. Two-qubit gates have been demonstrated with 99.8% fidelity in 1.6 μs [43] and with 99.9% fidelity in 30 μs [16]. These gate durations are still long on the timescale of superconducting qubits, in which gates can be performed on the order of 10's to 100's of nanoseconds [140], but are not the most time-intensive step in a typical trapped-ion quantum algorithm. In the case of the QCCD, major contributions to the duration of a quantum algorithm come from transport, separation, and merge steps, each of which are typically performed on hundred's of microsecond timescales [64].

To reduce the duration of transport steps in the QCCD, one might naïvely move the well more quickly. However, accelerating the ion from a well at rest to a moving well imparts momentum $\Delta p = m\Delta v$ to the ion, heating it. If the transport duration is long, Δv is small and the energy imparted by momentum kick will be small relative to the energy scale of a harmonic well $\hbar\omega_{sec}$. For a faster transport, this momentum kick is proportionally larger and the ion can end up in a highly excited state after the end of transport. Performing two-qubit gates requires ions to be cooled to near the ground state of motion [16][15]. We wish to avoid require potentially lengthy recooling after transport, so transport should end with the ion still close to the motional ground state.

The solution lies in thinking carefully about transport boundary conditions. Starting and stopping the transport will both give a momentum kick. If the ion were not oscillating in its

harmonic well, this would not be a problem. The ion would receive equal and opposite momentum kicks at the start and end of the transport and return to its initial motional state. However, the ion *is* oscillating in a harmonic well, and thus the timing of the kicks relative to the phase of this oscillation must be considered.

7.1 Previous work

Faster-than-adiabatic transport has been performed in three dimensional microstructured traps that are large compared to surface traps. Ions were transported $370\ \mu\text{m}$ in $8\ \mu\text{s}$ [4] with 0.1 quanta of excitation and $280\ \mu\text{m}$ in $3.6\ \mu\text{s}$ with 0.1 quanta of excitation [5]. Both of these distances are the width of a single trap electrode (i.e. a single zone in the QCCD) in their respective experiments. These experiments were performed by transporting the ion with nominally constant velocity in an integer number of trap periods, in which case the ion theoretically sees exactly opposite momentum kicks from starting and stopping transport and ends up back in the ground state. This method is highly sensitive to the exact duration of the transport and requires careful calibration, making it more complicated to use in a larger QCCD device.

7.2 Theory of fast transport

7.2.1 Classical treatment

For simplicity, start by considering transport classically. I also make the simplifying assumption that the the ion's motional modes are well-decoupled from each other and we can consider only the center-of-mass motion along the direction of transport. This is a good assumption for a single ion or a single-species crystal but breaks in the case of a multi-species ion crystal [141]. The case of a single ion is discussed in detail, but much of this analysis is identical for the center of mass mode of a single-species chain.

Consider ion is in a well that is rigidly (i.e. with unchanging frequency) moving. The ion

sees the potential

$$\Phi(x, t) = \frac{1}{2}m\omega^2(x - X(t))^2 \quad (7.1)$$

where $X(t)$ is the time-dependent center of the well, and ω is the well's (unchanging) secular frequency.

I will take a moment here for a quick note about nomenclature. There are two positions here that we must keep track of: the position of the *ion* and the position of the *harmonic well*. The position of the harmonic well is a control parameter that can be changed by applying appropriate voltages to trap electrodes. The position of the ion is a dependent parameter, which we seek to understand and control. I denote the position of the ion as $x(t)$ and the position of the well as $X(t)$.

The ion's motion in such a potential can be described by the classical equation of motion

$$F = m\ddot{x} = -\frac{\partial\Phi(x, t)}{\partial x} \quad (7.2)$$

Consider the simplest case of a well that moves with constant velocity $X(t) = vt + X_0$ from time $t = 0$ to $t = T$:

$$\ddot{x} = -\frac{1}{m} \frac{\partial\Phi(x, t)}{\partial x} = -\omega^2(x - vt - X_0) \quad (7.3)$$

(For the time being, I will neglect the fact that a well instantaneously changing from sitting at rest at time $t < 0$ and moving with constant velocity at $t > 0$ requires infinitely acceleration.) For an ion that starts at rest ($\dot{x}(0) = 0$) at the center of the harmonic well ($x(0) = X_0$) (i.e. a well-cooled and micromotion-compensated ion), this is solved by

$$x(t) = \left(\frac{v}{\omega}\right) \cos\left(\omega t + \frac{\pi}{2}\right) + vt + X_0 \quad (7.4)$$

We wish to transport a distance L in time T . In this simple linear transport scheme, then $v = L/T$, and at time T the ion has kinetic energy

$$E_K = \frac{1}{2}m(\dot{x}(T))^2 = \frac{mL^2}{2T^2} \left(1 - \sin\left(\omega T + \frac{\pi}{2}\right)\right)^2 \quad (7.5)$$

We wish to find solutions in which E_K is minimized at the end of the transport time. This classical kinetic energy can be expressed in units of ‘‘motional quanta’’ to get a classical estimate of the final

quantum occupation of the motional mode after transport by

$$\bar{n} = \frac{E_K}{\hbar\omega}. \quad (7.6)$$

The final excitation, in “quanta,” is plotted in Figure 7.1(a) for a single ${}^9\text{Be}^+$ ion transported 45 μm (approximately the width of a single electrode in the transport trap) in a 2 MHz trap as a function of the transport duration. Around 250 μs , the adiabatic transport duration used in [60], the final excitation is less than one quantum, independent of precise transport duration. As transport becomes shorter, the excitation is larger. For transport of approximately 100 μs duration, the final excitation is generally more than 1 quantum, a level which would require cooling before performing high fidelity gates. At shorter time periods ($< 10 \mu\text{s}$), timescales which are comparable to the timescale of gates, the excitation can be hundreds or even thousands of quanta.

As can be seen by zooming in in Figure 7.1(b), the final transport excitation is quickly varying with periodicity equal to the trap secular period $2\pi/\omega$. For transport times equal to an integer number of ion oscillations in the trap $T = 2\pi N/\omega$, the ion can have low final excitation ($\bar{n} < 1$) after transport.

I have discussed the simple case of a well moving with constant velocity for purposes of illustration. For more general transport profiles $X(t)$, we can evaluate the profile classically by numerically solving Eq. 7.2 and evaluating $\bar{n}(T)$.

7.2.2 Quantum Treatment

In a quantum mechanical picture, the ion motion is described by the Hamiltonian

$$\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega^2(\hat{x} - X(t))^2 \quad (7.7)$$

where, again, we are considering a rigid well with unchanging trap frequency ω with time-dependent well minimum position $X(t)$. As is derived in [71], the equation of motion of an ion in state $|\tilde{\psi}\rangle$ relative to the moving frame of the well is identical to the equation of motion of a displaced coherent state

$$\frac{\partial}{\partial t}|\tilde{\psi}\rangle = i\hbar\sqrt{\frac{m\omega}{2\hbar}}\dot{X}(t)\left(\hat{a}^\dagger e^{i\omega t} - \hat{a}e^{-i\omega t}\right) \quad (7.8)$$




Figure 7.1: Final excitation $\bar{n} = E_K/\hbar\omega$ after a constant velocity, 45 μm transport of a 2 MHz, as a function of transport duration, (a) plotted from $T = 0$ to 250 μs , the transport time used in [60], and (b) zoomed in to see the fast oscillations. The final excitation oscillates quickly at the secular frequency ω . For transport times shorter than 100 μs , the ion's final excitation will be more than 1 quantum unless the transport time is carefully calibrated. If the transport time is exactly an integer number of trap periods $T = 2\pi N/\omega$, the final excitation can still be well below one quantum.

with displacement parameter

$$\alpha(t) \equiv \dot{X}(t)e^{i\omega t} \sqrt{\frac{m\omega}{2\hbar}} \quad (7.9)$$

[142]. This can be integrate to find

$$\alpha(t) = \sqrt{\frac{m\omega}{2\hbar}} \int_0^t \dot{X}(t')e^{i\omega t'} dt' \quad (7.10)$$

To find transport profiles with low net excitation, the goal is to find $X(t)$ such that $\alpha(T) = 0$ at the end of the transport. We allow nonzero $\alpha(t)$ for values of t during transport, as long as net displacement at the end of the transport is zero.

The expected value of the ion's final position in the lab frame is

$$x(t) = X(t) - \int_0^t \dot{X}(t') \cos[\omega(t-t')] dt' \quad (7.11)$$

[71]. In the simplest case of a constant velocity transport, $\dot{X}(t) = v = L/T$, where L is the distance of transport and T is the transport duration, so

$$x(t) = \frac{Lt}{T} - \frac{L}{T} \int_0^t \cos[\omega(t-t')] dt' = \frac{L}{T} \left(t - \frac{1}{\omega} \sin(\omega t) \right) \quad (7.12)$$

The ion velocity is then

$$\dot{x}(t) = \frac{L}{T} (1 - \cos(\omega t)) \quad (7.13)$$

This is identical to the classical result. The quantum treatment is illustrative of how we can find more transport profiles $X(t)$ that give a low final motional excitation. In particular, this gives us a place to look for solutions which are relatively insensitive to calibration errors and which might scale more favorably to larger systems with more degrees of freedom.

7.2.3 Transport of ion chains

For transport of multi-ion crystals, we must consider not only the externally applied harmonic potential but also the Coulomb interaction between ions:

$$\hat{H} = \sum_i \left(\frac{\hat{p}_i^2}{2m_i} + \frac{1}{2} m_i \omega^2 (\hat{x}_i - X(t))^2 + \frac{1}{4\pi\epsilon_0} \sum_{j>i} \frac{e^2}{|x_i - x_j|} \right) \quad (7.14)$$

In general, solving the Hamiltonian fully accounting for the Coulomb interaction is a hard problem. In the limit of small oscillations about the ion equilibrium position, we can work in the normal mode-picture, discussed in Section 2.2.3. In this picture, the Hamiltonian of an N ion crystal can be approximated as a sum of $3N$ harmonic oscillators. The treatment of a transported ion crystal should be identical to the treatment of a single ion discussed above, with minor adaptations to replace harmonic oscillator operators with sums over $3N$ operators. In the case of a rigid harmonic well that is translated axially, only the center-of-mass (COM) mode should be excited by the transport, so transport profiles that translate a single ion with low net excitation should also transport a single-species ion chain with low net excitation.

Much of the ion trapping community has begun to move towards mixed-species operation [143][61]. Mixed-species crystals are challenging because the modes can no longer be treated as independent harmonic oscillators, so this treatment of ion transport would not apply to a mixed-species crystal. Although theoretical proposals have suggested methods for fast multi-species transport [141], these have not been implemented to date.

7.2.4 Finding transport profiles

The above discussion presents methods for understanding ion motion during transport and for predicting the final ion motional state occupation \bar{n} after applying a transport profile $X(t)$. However, we wish to have a method to derive optimal transport profiles $X(t)$ such that \bar{n} is minimized after transport is minimized.

We choose desired boundary conditions for transport profiles, where $X(t)$ is the position of the well as a function of time:

$$X(0) = 0 \quad (7.15)$$

$$X(T) = L \quad (7.16)$$

$$\dot{X}(0) = 0 \quad (7.17)$$

$$\dot{X}(T) = 0 \quad (7.18)$$

namely, the well should move from $X = 0$ to $X = L$ in time T . Further, the well should start and end at rest. This requirement is *not* fulfilled by constant velocity transport, which assumes that the well can instantaneously jump from sitting at rest to moving with constant velocity. The final boundary condition is that the ion should experience no net displacement relative to the frame of the well, from Eq. 7.10:

$$\alpha(t) = \sqrt{\frac{m\omega}{2\hbar}} \int_0^t \dot{X}(t') e^{i\omega t'} dt' = 0. \quad (7.19)$$

To find an $X(t)$ that fulfills these boundary conditions, make the ansatz

$$X(t) = \frac{L}{2} + \sum_{n=1}^N a_n \cos\left(\frac{(2n-1)\pi t}{T}\right) \quad (7.20)$$

for some maximum number of terms N . This automatically fulfills the derivative boundary conditions $\dot{X}(0) = \dot{X}(T) = 0$. Eq. 7.15 and Eq. 7.16 both set the constraint

$$\sum_{n=1}^N a_n = -\frac{L}{2}. \quad (7.21)$$

For readability, I will make the definition

$$C_n(\omega, T) \equiv -\frac{\pi(2n-1)}{T} \int_0^T \sin\left(\frac{(2n-1)\pi t}{T}\right) e^{i\omega t} dt = \frac{1 + e^{i\omega T}}{(\omega T / (2n-1)\pi)^2 - 1}. \quad (7.22)$$

In this case, the displacement boundary condition Eq. 7.19 can be rewritten

$$\sum_{n=1}^N a_n C_n(\omega, T) = 0 \quad (7.23)$$

Noting that $1 + e^{i\omega T}$ is a common factor, I will make the further definition

$$\tilde{C}_n(\omega, T) \equiv \frac{1}{(\omega T / (2n - 1)\pi)^2 - 1} \quad (7.24)$$

from which we can rewrite Eq. 7.19 as

$$\sum_{n=1}^N a_n \tilde{C}_n(\omega, T) = 0. \quad (7.25)$$

Eqs. 7.21 and 7.25 make a system of two linear equations. This can be analytically solved for $N = 2$, and the solution $\{a_1(\omega, T), a_2(\omega, T)\}$ defines the transport profile $X(t)$ that I have used in these experiments. If more terms are included in the sum, i.e. $N > 2$, the problem is underconstrained. Optimization of this underconstrained problem could yield more robust or more physically achievable transport profiles. As a first demonstration I have chosen to work with the $N = 2$ profiles.

7.3 Waveform Generation

I have discussed finding transport profiles $X(t)$ that are optimized to minimize ion excitation at the end of the transport. However, these profiles must be physically implementable in the lab. I will describe the generation of a set of voltage waveforms $\{V_{el}(t)\}$ for each of the 23 electrodes of the transport trap that produce the desired transport profile.

7.3.1 Fast DACs

These waveforms are generated on fast arbitrary waveform generators, nicknamed ‘‘PDQs,’’ as described in [116]. The limiting factor in our ability to produce quickly changing voltages is set by the AD8250 output amplifiers. These amplifiers have a maximum amplitude of ± 10 V. The slew rate was measured to be > 35 V/ μ s for every PDQ. These amplifiers, combined with the low-

pass filters on the trap electrodes, are the limiting factor in our ability to produce fast transport waveforms.

7.3.2 Finding potentials

The transport profile $X(t)$ is continuous in time. In order to find the transport potentials, we break this profile up into discrete time steps $X_i(t_i)$. In the work described, we have used 16 discrete time steps per “trap zone,” i.e. 16 steps per 44 μm of travel. At each time step, we generate a set of voltages $\{V_{el,i}(t_i)\}$ that produces a well with the desired position $X_i(t_i)$ using the method described in Section 2.3.1. As discussed in that section, solving to find the voltages that produce a well at the desired location with appropriate curvatures is an underconstrained problem, and multiple sets of voltages can produce first-order identical harmonic wells. We can then choose within this class of voltage sets to find an optimum solution. The waveforms were optimized to maximize transport speed using two cost functions:

- **Voltage smoothness:** The voltages on each electrode should vary as slowly as possible in time. This is achieved by cost function

$$C_{VS} = \sum_{i=0}^{i_{max}-1} |V_{el,i+1} - V_{el,i}| \quad (7.26)$$

that applies a penalty for large variations between the voltage on electrode el at timestep i and $i + 1$.

- **Electrode contribution strength:** The voltages on each electrode should have a non-negligible impact on the shape of the potential. This is achieved by cost function

$$C_{ECS} = \sum_{el,i} \frac{|V_{el,i}|}{|\vec{E}_{el,i}|^2} \quad (7.27)$$

where $\vec{E}_{el,i}$ is the electric field produced at the ion position $X_i(t_i)$ in time step i due to the voltage $V_{el,i}$ on electrode el . This applies a penalty for large voltages that produce small electric fields at the ion location and tends to keep voltages at zero if they have little impact on the properties of the confining well.




Figure 7.2: A first order RC filter, used to limit electric field noise in this experiment. This has a 3 dB cutoff frequency $f_{3dB} = 1/(2\pi\sqrt{RC})$.

The full cost function is a weighted sum of both of these

$$C = w_{VS}C_{VS} + w_{ECS}C_{ECS} . \quad (7.28)$$

The weighting is chosen by manually varying w_{VS} and w_{ECS} to find the shortest transport duration where the necessary voltages could be produced by the fast DACs.

There is also a constraint that requires $|V_{el,i}| \leq 10$ V for all electrodes and times.

7.3.3 Predistortion

As discussed in Section 4.2.5, there are low-pass filters on each electrode to reduce electric field noise at the ion. The 3dB cutoff frequency of the filters used in these experiments are 5 – 23 kHz, so they will attenuate the fastest-varying components of the waveforms and distort the waveforms produced by the fast DACs. We must instead generate waveforms that are predistorted, i.e. that have the effect of the filter taken into account. This ensures that the waveform after the filter will look as designed.

The filter circuit is shown in Figure 7.2. By consideration of Kirchoff's current law, we can determine the relation between $V_{in}(t)$ and $V_{out}(t)$ as

$$V_{in}(t) = V_{out}(t) + RC\dot{V}_{out}(t) . \quad (7.29)$$

We wish to use this to derive $V_{in}(t)$, the predistorted waveform that produces $V_{out}(t)$, the desired




Figure 7.3: Predistorted and target waveforms on electrode DC4 for a transport profile in which the ion is transported four trap zones (176 μm) and back in 5 μs . The target waveform is shown in green. The predistorted waveform assuming a first-order RC filter is shown in blue. The predistorted waveform for a second order filter, accounting for line resistance and parasitic capacitance, is shown in orange.

waveform on the electrode. This is simply a matter of calculating $V_{out}(t)$ and $\dot{V}_{out}(t)$ and plugging them into Eq. 7.29. A predistortion of a sample waveform (green) for $(R, C) = (6.8 \text{ k}\Omega, 1 \text{ nF})$ is shown in Figure 7.3 in blue.

This relationship can also be considered in the frequency domain by noting that the circuit can be modeled as a voltage divider. The voltage on the electrode $V_{out}(\omega)$ is then

$$V_{out}(\omega) = \frac{Z_C}{Z_C + Z_R} V_{in}(\omega) = \frac{1/i\omega C}{1/i\omega C + R} V_{in}(\omega) = \frac{1}{1 + i\omega RC} V_{in}(\omega) \quad (7.30)$$

This can be used to determine the 3 dB cutoff frequency for the filter, i.e. the frequency f_{3dB} at which the output power (which is proportional to $|V_{out}|^2$) is reduced from the input power by 3 dB. In the first order filter shown in Figure 7.2,

$$f_{3dB} = \frac{1}{2\pi RC} \quad (7.31)$$

Two sets of filters were used in this trap: “loose” filters, with $(R, C) = (6.8 \text{ k}\Omega, 1.0 \text{ nF})$ for

a corner frequency 23 kHz, and “tight” filters, with $(R, C) = (15 \text{ k}\Omega, 2.2 \text{ nF})$ for corner frequency 4.8 kHz. The tight filters are connected to electrodes whose waveforms only vary by small amounts during transport. The remaining electrodes are filtered by loose filters.

The predistortion is the limiting factor in our ability to generate fast transport profiles. The most quickly varying components of the waveform overshoot the target voltage, as can be seen in Figure 7.3. We are unable to generate waveforms with larger than $\pm 10 \text{ V}$, so we are speed-limited to the transport duration at which the overshoot hits this $\pm 10 \text{ V}$ limit. There is some variation depending on the specific waveforms, but we are limited by the fast DACs to waveforms that move an ion with an approximate average speed $15 \text{ }\mu\text{m}/\mu\text{s}$. For faster waveforms, the maximum voltages on the predistorted waveforms exceed the $\pm 10 \text{ V}$ limit of the DACs.

7.3.3.1 Higher order effects

The prior discussion holds for a filter that can be well-described as a first order RC filter. The cryostat wiring has non-negligible parasitic capacitance and line resistance $(R_s, C_s) \approx (150 \text{ }\Omega, 125 \text{ pF})$. This stray resistance and capacitance are small compared to the filter values and only have a minor effect on the filter behavior at low frequencies. At higher frequencies, the stray resistance and capacitance have more of an effect on the filter attenuation.

To lowest order, the filter circuit accounting for stray resistance and capacitance can be modeled as the second-order RC filter drawn in Figure 7.2. By consideration of Kirchoff’s voltage and current laws, we can derive the characteristic relationship between input and output waveforms in this circuit as

$$V_{in}(t) = V_{out}(t) + (RC + RC_S + R_S C_S) \dot{V}_{out}(t) + (RC)(R_S C_S) \ddot{V}_{out}(t) \quad (7.32)$$

This can also be expressed in the frequency domain as

$$V_{out}(\omega) = T(\omega)V_{in}(\omega) = \frac{V_{in}(\omega)}{1 + i(RC + RC_S + R_S C_S)\omega - (RC)(R_S C_S)\omega^2} \cdot \quad (7.33)$$

A predistorted waveform taking second order effects into account is shown in Figure 7.3. A comparison of $T(\omega) \equiv \text{Re}(V_{out}(\omega)/V_{in}(\omega))$ between the the first order and second order filter




Figure 7.4: Transfer functions for the (a) loose filters $R = 6.8 \text{ k}\Omega$, $C = 1.0 \text{ nF}$ and (b) tight filters $R = 15 \text{ k}\Omega$, $C = 2.2 \text{ nF}$ as a function of frequency. First order transfer functions are shown in blue, neglecting parasitic capacitance and line resistance effects. Second-order transfer functions, assuming $(R_S, C_S) = (150 \text{ }\Omega, 125 \text{ pF})$, are shown in red.

approximations is shown in Figure 7.4. This effect is small but should be taken into account for faster transport profiles, which have more high frequency components.

One question that remains is whether this second-order approximation is appropriate. The second-order RC filter modeled in Figure 7.2(b) assumes that the stray resistance and capacitance can be treated as lumped circuit elements; however, they are in fact distributed throughout the cryostat wiring. The exact transfer function of this system is unknown, and it may impact our ability to predistort the fastest waveforms accurately.

7.4 Results

We do not have the ability to move our laser beams on the timescales we are interested in investigating for fast transport, so we investigated waveforms in which we transport a desired distance, wait some variable wait time, and return. The figure of merit is the change in the ion’s motional state occupation, measured using the sideband imbalance method. All measurements were compared to a “calibration scan,” in which the ion was identically prepared and waited in the start well without being moved for a time equal to the time of transport, wait, and return.

Before all fast transport tests, we compensated for stray fields and axial curvature at every




Figure 7.5: (a) Raman red sideband scan after an ion in a 2.0 MHz well was transported 88 μm in 25 μs transport with 0.03(2) quanta of heating per transport. The sideband frequency is shifted after transport by 7.5 kHz for unknown reasons. This behavior requires further investigation. (b) Net excitation per transport after an optimized transport for transport durations between 25 μs and 25.5 μs , scanning the duration over a full trap period. For comparison, net excitation per transport for an ion transported the same distance with constant velocity is shown in teal. There is only a small range of durations from 25.1 μs to 25.15 μs where the ion is cold enough after constant velocity transport that the sideband imbalance method of temperature measurement is useful. At all transport durations, the ion transported using the optimized transport profiles had lower excitation due to transport than the lowest net excitation achieved with a constant velocity transport.

$X_i(t_i)$ in the transport profile. Particular attention was paid when compensating for axial stray fields, since axial stray fields push the ion out of the design transport profile. These shims were added to the transport waveforms $\{V_{el,i}\}$ to generate a set of transport voltage arrays accounting for stray fields. “Buffer points” were added to the end of these arrays to generate a desired wait time, and then the reverse waveform was attached. These voltage arrays were smoothed by a Gaussian filter with a width of one time step and then predistorted as described above. Finally, the arrays were converted into cubic splines to be generated by the fast DACs [116].

Raman sidebands after the ion was transported 88 μm with 25 μs duration are shown in Figure 7.5(a). This transport is equal to the width of two trap zones, requiring more complex waveforms than transporting a single zone does. The calibration scan Raman sidebands (in which the ion was cooled and then waited an identical delay) indicate an ion motional state occupation $\bar{n} = 0.17(4)$. The sidebands after transport indicate an ion motional state occupation $\bar{n} = 0.22(2)$.

This gives a “per transport” net excitation of $0.03(2)$ quanta. To confirm that the ion was actually transported, we also turned on the laser and measured counts during the delay between transporting and returning, when the ion is not in the initial well. These were indistinguishable from background laser counts, indicating that the ion was successfully transported out of the initial trap zone.

One unanswered question from these measurements is apparent in the Raman frequency scan in Figure 7.5(a). The sideband frequency after transport is shifted $+7.5$ kHz relative to the sideband frequency in the calibration measurement. This is consistent across measurements and unexplained. One possible explanation is that the electrode voltages were slightly different at the times of the two measurements. If this is the case, it would likely be an effect of the RC filters, indicating that the waveforms were not correctly predistorted. This is likely due to the higher order filtering effects of stray resistance and capacitance. This behavior merits further investigations.

The final transport temperature did not strongly depend on transport duration. Net excitation per transport for transport durations of $25 \mu\text{s}$ to $25.5 \mu\text{s}$ (scanning over a full period of trap motion) are shown in purple in Figure 7.5(b). The temperature added due to transport is less than one quantum for all transport durations. In contrast, the motion added per transport after transporting an ion with constant velocity is shown in teal. Using this method, there is only a small range of ~ 50 ns transport duration for which the final trap temperature is cold enough to be accurately measured using the sideband imbalance technique. Furthermore, the motional excitation using our new optimized method at all transport durations is less than that of the best constant velocity transport.

Since we only measured ion temperatures after transporting and return, there is a possibility that only the “out and back” transport profile is cold but the ion is heated after one-way transport. To ensure that this was not the case, we scanned the delay time between the two transport segments. If the ion were excited to a coherent state of motion at the end of the first transport, then the final temperature would depend on delay time. The change in ion temperature due to transport is plotted as a function of delay time for a $25 \mu\text{s}$ transport in Figure 7.6. The data are consistent along the delay times, indicating that the ion motion is not excited after one-way transport.





Figure 7.6: Change in \bar{n} for an ion transported $88 \mu\text{m}$ in $25 \mu\text{s}$, delayed variable wait time, and transported back to the readout zone in $25 \mu\text{s}$. The change in temperature due to transport does not depend on delay duration outside of the error bars, indicating that one-way transport is cold.


At shorter transport durations than $25 \mu\text{s}$, the transport was no longer insensitive to transport duration. The behavior at these shorter transport durations resembled the behavior of ions transported with constant velocity, in that it was possible to find specific transport durations for which the transport was cold. An example is shown in Figure 7.7(a), in which an ion was transported $88 \mu\text{m}$ in $9.1 \mu\text{s}$ with $0.28(4)$ quanta of heating due to transport. However, after a $9.2 \mu\text{s}$ transport, the ion was heated so much that the ion temperature could not be measured with a sideband imbalance, shown in Figure 7.7(b). We attribute this to failures in the predistortion arising from the fact that the exact filter transfer function is unknown. At short transport durations, we are no longer applying the design waveforms to the electrode so we are no longer able to controllably transport the ion with no net displacement.

7.5 Summary

We have demonstrated a method for fast, multi-zone transport of ions in surface electrode traps. This method requires less careful calibration of transport duration than the previously demonstrated faster-than-adiabatic transport, making it more useful for practical application as part of a longer quantum algorithm. We have used this method to demonstrate faster-than-adiabatic transport of ions with less net excitation than has been previously demonstrated in other



(a)



(b)

Figure 7.7: (a) Raman scan after an ion in a 2.0 MHz well was transported 88 μm in 9.1 μs . The calibration Raman sideband scan indicates $\bar{n} = 0.14(4)$ and the Raman sideband scan after transport indicates $\bar{n} = 0.70(7)$. This corresponds to 0.28(4) quanta of heating per transport. (b) Raman scan after an ion in a 2.0 MHz well was transported 88 μm in 9.2 μs . The ion is heated so much that the sideband imbalance method for measurement of motional state occupation is not useful.

faster-than-adiabatic transport experiments. At the shortest transport durations, we are limited by our ability to accurately apply the design waveforms to the trap electrodes.

Chapter 8

Conclusions and future directions

The quantum CCD is just one of many promising visions for scaling the trapped-ion quantum computer. I have presented work motivated by a QCCD approach. However, it is my hope that the technical achievements presented in this thesis are of use to the broader ion-trapping community, beyond the subset of people studying the QCCD.

8.1 Trap-integrated detection

Before this work, SNSPDs appeared to be a good technology for integration into surface electrode traps, but successful integration had not yet been demonstrated. We have seen that it is possible to reliably trap ions over a SNSPD with no degradation to ion lifetime. Further, even with some performance degradation caused by RF pickup, the SNSPD count rates exceed those of the best commercially available PMTs. This enables qubit state detection in short times and with high fidelities.

8.1.1 Next-generation trap-integrated SNSPDs

Major challenges still remain before trap-integrated SNSPDs are appropriate for all trapped-ion quantum computing applications. In particular, the behavior of the detector in the presence of RF introduces a major challenge for many applications. It is not clear that at present, the detector could be integrated with any ion other than ${}^9\text{Be}^+$ due to the higher RF amplitudes needed to trap heavier ions. One possibility to mitigate this effect is to fabricate the detector under a grounded,

electrically-conductive transparent film such as ITO or similar. This could act as a shielding layer, protecting the detector from RF pickup. Such a layer would also serve to shield the ion from the electric field noise caused by detection pulses.

It also may be possible to explore different superconducting materials, possibly with a higher critical temperature. A higher T_C would be helpful for trapping other ions, since the higher amplitude RF used to trap heavier ions would likely heat the cryostat past the 3.65 K limit at which we were able to operate this device. The higher T_C would also be helpful for limiting dependence on RF pickup and reducing stray detector clicks. If an SNSPD is operated further below its T_C , the “plateau” where detector efficiency is not dependent on bias current level becomes broader. With a sufficiently large plateau, the reduction in efficiency caused by RF pickup could be reduced or even eliminated. Additionally, it might be possible to operate in a way where the detector can be fully turned off by not applying a DC bias current, meaning that stray clicks could be eliminated outside of the detection window. This would eliminate the ion heating due to detector pulses caused by stray light and would enable operation of high-intensity lasers over the SNSPD.

8.1.2 Integration into a larger device

Many groups are pursuing integration of more and more components of a trapped-ion quantum computer onto a single chip. On-chip detection using trap-integrated SNSPDs is a major piece of that puzzle, since it eliminates the need for a bulky objective, camera, and PMT. Other pieces might include on-chip delivery of cooling and detection laser light or integration of microwave antennae for two-qubit gates. Integrated optics for delivery of cooling and detection light is a naturally complementary technology to the SNSPD, since SNSPDs must be operated at cryogenic temperatures and on-chip laser light delivery has been demonstrated to reduce an ion’s sensitivity to vibration-induced errors [101].

A larger device could also include multiple detection zones by incorporating multiple spatially-separated detectors. This would allow us to measure crosstalk from ions simultaneously trapped over different trap-integrated SNSPDs

8.2 Fast manipulation of ion chains

In this work, I have presented optimized waveforms for fast transport of ions with reduced dependence on transport duration that introduced very little motional excitation. This addresses one piece of the puzzle for fast manipulation of ions chains in the quantum CCD. However, recall that in order to fully be able to arbitrarily rearrange qubits in the quantum CCD, we need four operations: transport, separation, merge, and reorder.

8.2.1 Separation

To date, separation has not been performed at timescales comparable to those of laser gates. Separation has been demonstrated in a “quasi-adiabatic” regime, but has still been within the limit

$$\frac{d\omega}{dt} \frac{1}{\omega^2} \ll 1 . \quad (8.1)$$

The fastest cold separation of ions to date has been performed in 55 μs with ~ 2 quanta of excitation [4] and in 80 μs with ~ 4 quanta of excitation [144]. These are both slower and introduce more excitation than fast transport experiments in the same systems.

One proposal for fast separation that has to-date not been implemented is the idea of a “release and catch” method of separation. This would involve quickly dropping axial confinement and letting the ions separate under Coulomb repulsion. These ions would separate coherently. They could then be caught in harmonic wells after separating a desired distance. The harmonic wells would have to be a shallower frequency because the ion waveforms will spread during separation. If this method could be implemented, it could be used to separate two ${}^9\text{Be}^+$ ions in less than 1 μs . This experiment would require the ability to apply very quickly changing trap potentials and to carefully control the stray field in this trap. Both of these features were demonstrated in the transport trap, making it an ideal platform for testing this idea.

Bibliography

- [1] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped ions. J Res NIST, 103(3):259–328, 1998.
- [2] D Kielpinski, C Monroe, and D J Wineland. Architecture for a large-scale ion-trap quantum computer. Nature, 417(6890):709–711, 2002.
- [3] D. H. Slichter, V. B. Verma, D. Leibfried, R. P. Mirin, S. W. Nam, and D. J. Wineland. UV-sensitive superconducting nanowire single photon detectors for integration in an ion trap. 25(8):8705–8720, 2016.
- [4] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. Coherent diabatic ion transport and separation in a multizone trap array. Physical Review Letters, 109(8):1–4, 2012.
- [5] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger. Controlling fast transport of cold trapped ions. Physical Review Letters, 109(8):1–5, 2012.
- [6] Lov K Grover. A fast quantum mechanical algorithm for database search. In STOC '96 Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing, pages 212–219, 1996.
- [7] Peter W. Shor. Algorithms for Quantum Computation : Discrete Logarithms and Factoring. In Proceedings 35th Annual Symposium on Foundations of Computer Science, pages 124–134, 1994.
- [8] Richard P Feynman. Simulating Physics with Computers. R.P. Int J Theor Phys, 21:467–488, 1982.
- [9] Devoret M. H. and R. J. Schoelkopf. Superconducting Circuits for Quantum Information: An Outlook. Science, 339(6124):1169–1175, 2013.
- [10] Christoph Kloeffel and Daniel Loss. Prospects for Spin-Based Quantum Computing in Quantum Dots. Annual Rev. of Cond. Mat. Phys., 4:51–81, 2013.
- [11] Lilian Childress and Ronald Hanson. Diamond NV centers for quantum computing and quantum networks. MRS Bulletin, 38(2):134–138, 2013.

- [12] Pieter Kok, Jonathan P Dowling, and G J Milburn. Linear optical quantum computing with photonic qubits. Reviews of Modern Physics, 79(March):135–174, 2007.
- [13] Harry Levine, Alexander Keesling, Ahmed Omran, Hannes Bernien, Sylvain Schwartz, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. High-Fidelity Control and Entanglement of Rydberg-Atom Qubits. Physical Review Letters, 121(12):1–6, 2018.
- [14] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. Single-qubit-gate error below 10⁻⁴ in a trapped ion. Physical Review A - Atomic, Molecular, and Optical Physics, 84(3):1–4, 2011.
- [15] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits. Physical Review Letters, 117(6):1–6, 2016.
- [16] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-Fidelity Universal Gate Set for Be 9 + Ion Qubits. Physical Review Letters, 117(6):1–5, 2016.
- [17] T P Harty, M A Sepiol, D T C Allcock, C J Ballance, J E Tarlton, and D M Lucas. High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves. Physical Review Letters, 117(14):140501, 2016.
- [18] C Langer, R Ozeri, J D Jost, J Chiaverini, B Demarco, R B Blakestad, J Britton, and D B Hume. Long-Lived Qubit Memory Using Atomic Ions. Physical Review Letters, 95(6):060502, 2005.
- [19] K A Landsman, Y Wu, P H Leung, D Zhu, N M Linke, K R Brown, L Duan, and C Monroe. Two-qubit entangling gates within arbitrarily long chains of trapped ions. Physical Review A - Atomic, Molecular, and Optical Physics, 100(2):022332, 2019.
- [20] John Preskill. Quantum Computing in the NISQ era and beyond. Quantum, 2:79, 2018.
- [21] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G.S.L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. Nature, 574(7779):505–510, 2019.

- [22] G Pagano, A Bapat, P Becker, K. S. Collins, A. De, P W Hess, H W Kaplan, A Kypriandidis, W L Tan, C Baldwin, L T Brady, A Deshpande, F Liu, S Jordan, A V Gorshkov, and C Monroe. Quantum Approximate Optimization with a Trapped-Ion Quantum Simulator. [arXiv:1906.02700](https://arxiv.org/abs/1906.02700), 2019.
- [23] Cornelius Hempel, Christine Maier, Jonathan Romero, Jarrod McClean, Thomas Monz, Heng Shen, Petar Jurcevic, Ben P Lanyon, Peter Love, Ryan Babbush, Alán Aspuru-guzik, Rainer Blatt, and Christian F Roos. Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator. *Physical Review X*, 8(3):31022, 2018.
- [24] Yunseong Nam, Jwo-sy Chen, Neal C Pseni, Kenneth Wright, Conor Delaney, Dmitri Maslov, Kenneth R Brown, Stewart Allen, Jason M Amini, Joel Apisdorf, Kristin M Beck, Vandiver Chaplin, Mika Chmielewski, Coleman Collins, Shantanu Debnath, M Andrew, Kai M Hudek, Matthew Keesan, Sarah M Kreikemeier, Jonathan Mizrahi, Phil Solomon, Mike Williams, Jaime David Wong-campos, Christopher Monroe, and Jungsang Kim. Ground-state energy estimation of the water molecule on a trapped ion quantum computer. [arXiv:1902.10171](https://arxiv.org/abs/1902.10171), 2019.
- [25] G Pagano, P W Hess, H W Kaplan, W L Tan, P Richerme, Becker P, A Kypriandidis, J Zhang, E Birckelbaw, and M R Hernandez. Cryogenic trapped-ion system for large scale quantum simulation. *Quantum Science and Technology*, 4(4):014004, 2018.
- [26] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge University Press, New York, 1 edition, 2000.
- [27] Yong Wan, Daniel Kienzler, Stephen D Erickson, Karl H Mayer, Ting Rei Tan, Jenny J Wu, Hilma M Vasconcelos, Scott Glancy, Emanuel Knill, David J Wineland, Andrew C Wilson, and Dietrich Leibfried. Quantum gate teleportation between separated qubits in a trapped-ion processor. *Science*, 364(6443):875–878, 2019.
- [28] H Kaufmann, T Ruster, C T Schmiegelow, M A Luda, V Kaushal, J Schulz, D. von Lindenfels, F. Schmidt-Kaler, and U. G. Poschinger. Scalable Creation of Long-Lived Multipartite Entanglement. *Physical Review Letters*, 119(15):150504, 2017.
- [29] Sebastian Wolf, Dominik Studer, Klaus Wendt, and F. Schmidt-Kaler. Efficient and robust photo - ionization loading of beryllium ions. *Applied Physics B*, 124:30, 2018.
- [30] Colin D Bruzewicz, Robert McConnell, John Chiaverini, and Jeremy M Sage. Scalable loading of a two-dimensional trapped-ion array. *Nature Communications*, 7:13005, 2016.
- [31] C Ospelkaus, U Warring, Y Colombe, K R Brown, J M Amini, D Leibfried, and D J Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476:181, 2011.
- [32] Karan Mehta. *Integrated optical quantum manipulation and measurement of trapped ions by*. PhD thesis, Massachusetts Institute of Technology, 2017.
- [33] J. True Merrill, Curtis Volin, David Landgren, Jason M. Amini, Kenneth Wright, S. Charles Doret, C. S. Pai, Harley Hayden, Tyler Killian, Daniel Faircloth, Kenneth R. Brown, Alexa W. Harter, and Richard E. Slusher. Demonstration of integrated microscale optics in surface-electrode ion traps. *New Journal of Physics*, 13, 2011.

- [34] R Srinivas, S C Burd, R T Sutherland, A C Wilson, D J Wineland, D Leibfried, D T C Allcock, and D H Slichter. Trapped-Ion Spin-Motion Coupling with Microwaves and a Near-Motional Oscillating Magnetic Field Gradient. Physical Review Letters, 122(16):163201, 2019.
- [35] J Stuart, R Panock, C D Bruzewicz, J A Sedlacek, R McConnell, I L Chuang, J M Sage, and J Chiaverini. Chip-Integrated Voltage Sources for Control of Trapped Ions. Physical Review Applied, 10(1):1, 2019.
- [36] A. P. Vandevender, Y. Colombe, J. Amini, D. Leibfried, and D. J. Wineland. Efficient fiber optic detection of trapped ion fluorescence. Physical Review Letters, 105:023001, 2010.
- [37] Peter F. Herskind, Shannon X. Wang, Molu Shi, Yufei Ge, Marko Cetina, and Isaac L. Chuang. Microfabricated surface ion trap on a high-finesse optical mirror. Optics Letters, 36(16):3045, 2011.
- [38] Craig R. Clark, Chin Wen Chou, A. R. Ellis, Jeff Hunker, Shanalyn A. Kemme, Peter Maunz, Boyan Tabakov, Chris Tigges, and Daniel L. Stick. Characterization of fluorescence collection optics integrated with a microfabricated surface electrode ion trap. Physical Review Applied, 1(2):1–5, 2014.
- [39] Moji Ghadimi, Valdis Blums, Benjamin G. Norton, Paul M. Fisher, Steven C. Connell, Jason M. Amini, Curtis Volin, Harley Hayden, Chien Shing Pai, David Kielpinski, Mirko Lobino, and Erik W. Streed. Scalable ion-photon quantum interface based on integrated diffractive mirrors. npj Quantum Information, 3(1):4–7, 2017.
- [40] Andre Van Rynbach, Peter Maunz, and Jungsang Kim. An integrated mirror and surface ion trap with a tunable trap location. Applied Physics Letters, 109(22), 2016.
- [41] Amira M Eltony, Shannon X Wang, Gleb M Askelrod, Peter F Herskind, and I L Chuang. Transparent ion trap with integrated photodetector Transparent ion trap with integrated photodetector. Appl. Phys Rev., 102(5):054106, 2013.
- [42] L J Stephenson, D P Nadlinger, B C Nichol, S An, P Drmota, T G Ballance, K Thirumalai, J F Goodwin, D M Lucas, and C J Ballance. High-rate, high-fidelity entanglement of qubits across an elementary quantum network. 2019.
- [43] V M Schäfer, C J Ballance, K Thirumalai, L J Stephenson, T G Ballance, A M Steane, and D M Lucas. Fast quantum logic gates with trapped-ion qubits. Nature, 555:75–78, 2018.
- [44] C Monroe and J Kim. Scaling the Ion Trap Quantum Processor. Science, 339(6124):1164–1170, 2013.
- [45] A. C. Wilson, Y. Colombe, K. R. Brown, E. Knill, D. Leibfried, and D. J. Wineland. Tunable spin-spin interactions and entanglement of ions in separate potential wells. Nature, 512(1):57–60, 2014.
- [46] Frederick Hakelberg, Philip Kiefer, Matthias Wittmer, Ulrich Warring, Tobias Schaetz, Albert-ludwigs-universitt Freiburg, and Physikalisches Institut. Interference in a Prototype of a two-dimensional Ion Trap Array Quantum Simulator. pages 1–11, 2018.
- [47] Katherine Casey McCormick. Control and measurement of a single-ion quantum harmonic oscillator. PhD thesis, University of Colorado, 2019.

- [48] Roman Schmied, Janus H. Wesenberg, and Dietrich Leibfried. Quantum simulation of the hexagonal Kitaev model with trapped ions. New Journal of Physics, 13, 2011.
- [49] J Chiaverini, R B Blakestad, J Britton, J D Jost, C Langer, D Leibfried, R Ozeri, and D J Wineland. Surface-electrode architecture for ion-trap quantum information processing. Quantum Information and Computation, 5(6):419–439, 2005.
- [50] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, D. P. Pappas, D. Leibfried, and D. J. Wineland. Reduction of anomalous heating in an in-situ-cleaned ion trap. 103001(September):1–5, 2011.
- [51] J A Sedlacek, A Greene, J Stuart, R McConnell, C D Bruzewicz, J M Sage, and J Chiaverini. Distance scaling of electric-field noise in a surface-electrode ion trap. Physical Review A - Atomic, Molecular, and Optical Physics, 97:020302(R), 2018.
- [52] Ivan A Boldin, Alexander Kraft, and Christof Wunderlich. Measuring Anomalous Heating in a Planar Ion Trap with Variable Ion-Surface Separation. Physical Review Letters, 120(2):23201, 2018.
- [53] Da An, Clemens Matthiesen, Erik Urban, and H Hartmut. Distance scaling and polarization of electric-field noise in a surface ion trap. [arXiv:1906.06489](https://arxiv.org/abs/1906.06489), 2019.
- [54] G Gabrielse, X Fei, L A Orozco, R L Tjoelker, J Haas, H Kalinowsky, A Trainor, and W Kells. Thousandfold improvement in the measured antiproton mass. Physical Review Letters, 65(11):1317–1320, 1990.
- [55] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. Scaling and suppression of anomalous heating in ion traps. Physical Review Letters, 97(10):1–4, 2006.
- [56] Jaroslaw Labaziewicz, Yufei Ge, Paul Antohi, David Leibbrandt, Kenneth R. Brown, and Isaac L. Chuang. Suppression of heating rates in cryogenic surface-electrode ion traps. Physical Review Letters, 100(1):1–4, 2008.
- [57] J A Sedlacek, J Stuart, D H Slichter, C D Bruzewicz, R McConnell, J M Sage, and J Chiaverini. Evidence for multiple mechanisms underlying surface electric-field noise in ion traps. Physical Review A - Atomic, Molecular, and Optical Physics, 98:063430, 2018.
- [58] Crystal Noel, Maya Berlin-Udi, Clemens Matthiesen, Jessica Yu, Yi Zhou, Vincenzo Lordi, and Hartmut Häffner. Electric-field noise from thermally activated fluctuators in a surface ion trap. Physical Review A, 99(6):1–11, 2019.
- [59] M Natarajan Chandra, G Tanner Michael, and H Hadfield Robert. Superconducting nanowire single-photon detectors: physics and applications. Superconductor Science and Technology, 25(6):63001, 2012.
- [60] Ting Rei Tan. High-Fidelity Entangling Gates with Trapped-Ions. PhD thesis, 2016.
- [61] C D Bruzewicz, R McConnell, J Stuart, J M Sage, and J Chiaverini. Dual-Species, Multi-Qubit Logic Primitives for Ca + /Sr + Trapped-Ion Crystals. [arXiv:1905.13122](https://arxiv.org/abs/1905.13122), 2019.
- [62] F. Marsili, V. B. Verma, J. A. Stern, S. Harrington, A. E. Lita, T. Gerrits, I. Vayshenker, B. Baek, M. D. Shaw, R. P. Mirin, and S. W. Nam. Detecting single infrared photons with 93% system efficiency. Nature Photonics, 7(3):210–214, 2013.

- [63] E. E. Wollman, V. B. Verma, A. D. Beyer, R. M. Briggs, B. Korzh, J. P. Allmaras, F. Marsili, A. E. Lita, R. P. Mirin, S. W. Nam, and M. D. Shaw. UV superconducting nanowire single-photon detectors with high efficiency, low noise, and 4 K operating temperature. Optics Express, 25(22):26792, 2017.
- [64] Ryan Bowler. Coherent Ion Transport in a Multi-electrode Trap Array. PhD thesis, University of Colorado, 2015.
- [65] David J Griffiths. Introduction to Electrodynamics. Prentice-Hall, 3rd edition, 1999.
- [66] John Townsend. A modern approach to quantum mechanics. University Science Books, 2000.
- [67] Claude Cohen-Tannoudji, Bernard Diu, and Franck Laloe. Quantum Mechanics. Wiley-VCH, 1977.
- [68] C. Flühmann, T. L. Nguyen, M. Marinelli, V. Negnevitsky, K. Mehta, and J. P. Home. Encoding a qubit in a trapped-ion mechanical oscillator. Nature, 566(7745):513–517, 2019.
- [69] S C Burd, R Srinivas, J J Bollinger, A C Wilson, D J Wineland, D Leibfried, D H Slichter, and D. T.C. Allcock. Quantum amplification of mechanical oscillator motion. Science, 364(6446):1163–1165, 2019.
- [70] Katherine C. McCormick, Jonas Keller, Shaun C. Burd, David J. Wineland, Andrew C. Wilson, and Dietrich Leibfried. Quantum-enhanced sensing of a single-ion mechanical oscillator. Nature, 572:86–90, 2019.
- [71] Hoi Kwan Lau and Daniel F V James. Decoherence and dephasing errors caused by the dc Stark effect in rapid ion transport. Physical Review A - Atomic, Molecular, and Optical Physics, 83:062330, 2011.
- [72] Peter Maunz. High Optical Access Trap 2 . 0. Technical report, Sandia National Laboratories, Albuquerque, NM, 2016.
- [73] Mário H. Oliveira and José A. Miranda. Biot-Savart-like law in electrostatics. European Journal of Physics, 22(1):31–38, 2001.
- [74] David T C Allcock. Surface-Electrode Ion Traps for Scalable Quantum Computing. PhD thesis, Oxford, 2011.
- [75] David Hucul, Justin E. Christensen, Eric R. Hudson, and Wesley C. Campbell. Spectroscopy of a Synthetic Trapped Ion Qubit. Physical Review Letters, 119:100501, 2017.
- [76] C Figgatt, A Ostrander, N M Linke, K A Landsman, D Zhu, D Maslov, and C Monroe. Parallel entangling operations on a universal ion-trap quantum computer. Nature, 572:368–372, 2019.
- [77] A Kramida and W C Martin. A Compilation of Energy Levels and Wavelengths for the Spectrum of Neutral Beryllium (Be I). Journal of Physical and Chemical Reference Data, 26:1185, 1997.
- [78] A E Kramida and W C Martin. No Title, 2000.

- [79] J J Bollinger, J S Wells, D J Wineland, and Wayne M Itano. Hyperfine structure of the 2p 2P_{1/2} state in Be⁺. Physical Review A - Atomic, Molecular, and Optical Physics, 31(4):2711–2714, 1985.
- [80] C. E. Langer. High Fidelity Quantum Information Processing with Trapped Ions. PhD thesis, 2006.
- [81] D. J. Wineland, J. J. Bollinger, and Wayne M Itano. Laser-Fluorescence Mass Spectroscopy. Physical Review Letters, 50(9):628–631, 1983.
- [82] O Poulsen, T Andersen, and N J Skouboe. Fast-beam, zero-field level-crossing measurements of radiative lifetimes , fine and hyperfine structures in excited states of ionic and neutral beryllium. J Phys B, 8(9):1393–1405, 1975.
- [83] B H Bransden and C J Joachain. Physics of atoms and molecules. Prentice-Hall, 2nd edition, 2003.
- [84] Harold J Metcalf and Peter van der Straten. Laser Cooling and Trapping. Springer, New York, 1999.
- [85] Christopher J Foot. Atomic Physics. Oxford University Press, 2005.
- [86] Q A Turchette, D Kielpinski, B E King, D Leibfried, D M Meekhof, C J Myatt, M A Rowe, C A Sackett, C S Wood, W M Itano, C Monroe, and D J Wineland. Heating of trapped ions from the quantum ground state. Physical Review A - Atomic, Molecular, and Optical Physics, 61:063418, 2000.
- [87] Anders Sørensen and Klaus Mølmer. Quantum computation with ions in thermal motion. Physical Review Letters, 82(9):1971–1974, 1999.
- [88] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jenlenković, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. Nature, 422(6930):412–415, 2003.
- [89] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland. Generation of nonclassical motional states of a trapped atom. Physical Review Letters, 76(11):1796–1799, 1996.
- [90] M E Poitzch, J C Berquist, W M Itano, and D J Wineland. Cryogenic linear ion trap for accurate spectroscopy. Review of Scientific Instruments, 67:129, 2016.
- [91] P B Antohi, D Schuster, G M Akselrod, J Labaziewicz, Y Ge, Z Lin, W S Bakr, and I L Chuang. Cryogenic ion trapping systems with surface- electrode traps. Review of Scientific Instruments, page 013103, 2009.
- [92] M F Brandl, M W Van Mourik, L Postler, A Nolf, K Lakhmanskiy, R R Paiva, S Möller, N Daniilidis, H Häffner, V Kaushal, T Ruster, C Warschburger, H Kaufmann, U G Poschinger, P Schindler, T Monz, and R Blatt. Cryogenic setup for trapped ion quantum computing. Review of Scientific Instruments, 87:113103, 2016.
- [93] Jack W Ekin. Experimental techniques for low-temperature measurements. Oxford University Press, 2006.

- [94] Matteo Fadel. Cryogenic setup for the fast manipulation of the quantum motional states of trapped ions. PhD thesis, ETH Zurich, 2013.
- [95] M Schwarz, O O Versolato, A Windeberger, F R Brunner, T Ballance, S N Eberle, J Ullrich, P O Schmidt, A K Hansen, A D Gingell, M Drewsen, and J R Crespo López-urrutia. Cryogenic linear Paul trap for cold highly charged ion experiments. Review of Scientific Instruments, 83:083115, 2016.
- [96] Jeremy M Sage, Andrew J Kerman, and John Chiaverini. Loading of a surface-electrode ion trap from a remote, precooled source. Physical Review A - Atomic, Molecular, and Optical Physics, page 013417, 2012.
- [97] A T A M De Waele. Basic Operation of Cryocoolers and Related Thermal Machines. J Low Temp Phys, 164:179–236, 2011.
- [98] Anthony E Siegman. Lasers. University Science Books, 1986.
- [99] Christopher J Ballance. High-Fidelity Quantum Logic in $^{40}\text{Ca}^+$. PhD thesis, Oxford University, 2014.
- [100] A. Bermudez, P. O. Schmidt, M. B. Plenio, and A. Retzker. Robust trapped-ion quantum logic gates by continuous dynamical decoupling. Physical Review A - Atomic, Molecular, and Optical Physics, 85(4):040302, 2012.
- [101] R J Niffenegger, J Stuart, D Kharas, S Bramhavar, C D Bruzewicz, W Loh, R McConnell, D Reens, G N West, J M Sage, and J Chiaverini. Integrated optical control and enhanced coherence of ion qubits via multi-wavelength photonics. pages 1–9, 2020.
- [102] G. L. Weissler and W. R. Carlson. Vacuum Physics and Technology. Academic Press, 1980.
- [103] C L Yu, S S Wang, and T H Chuang. Intermetallic compounds formed at the interface between liquid indium and copper substrates. Journal of Electronic Materials, 31(5):488–493, 2002.
- [104] A. Duckham, J. E. Newson, M. V. Brown, T. R. Rude, O. M. Knio, E. M. Heian, and J. S. Subramanian. Method for fabricating large dimension bonds using reactive multilayer joining, 2010.
- [105] John Rumble, editor. CRC Handbook of Chemistry and Physics. CRC Press, 100th edition, 2019.
- [106] R. D. Knight. Storage of ions from laser-produced plasmas. Applied Physics Letters, 38(4):221–223, 1981.
- [107] R. J. Hendricks, D. M. Grant, P. F. Herskind, A. Dantan, and M. Drewsen. An all-optical ion-loading technique for scalable microtrap architectures. Applied Physics B: Lasers and Optics, 88(4):507–513, 2007.
- [108] K. Zimmermann, M. V. Okhapkin, O. A. Herrera-Sancho, and E. Peik. Laser ablation loading of a radiofrequency ion trap. Applied Physics B: Lasers and Optics, 107(4):883–889, 2012.
- [109] K. Sheridan, W. Lange, and M. Keller. All-optical ion generation for ion trap loading. Applied Physics B: Lasers and Optics, 104(4):755–761, 2011.

- [110] S. Olmschenk and P. Becker. Laser ablation production of Ba, Ca, Dy, Er, La, Lu, and Yb ions. Applied Physics B: Lasers and Optics, 123(4):1–6, 2017.
- [111] H Shao, M Wang, M Zeng, H Guan, and K Gao. Laser ablation and two-step photo-ionization for the generation of 40 Ca^+ . Journal of Physics Communications, 2(9):095019, 2018.
- [112] David R Leibbrandt, Robert J Clark, Jaroslaw Labaziewicz, Paul Antohi, Waseem Bakr, Kenneth R Brown, and Isaac L Chuang. Laser ablation loading of a surface-electrode ion trap. Physical Review A - Atomic, Molecular, and Optical Physics, 76(5), 2007.
- [113] Geert Vrijsen, Yuhi Aikyo, Robert F. Spivey, I. Volkan Inlek, and Jungsang Kim. Efficient isotope-selective pulsed laser ablation loading of 174 Yb^+ ions in a surface electrode trap. Optics Express, 27(23):33907, 2019.
- [114] Peng Zhou. Motional State Control and Ablation Loading in a Cryogenic Surface-Electrode Trap. PhD thesis, 2017.
- [115] Claude Phipps. Laser Ablation and its Applications. Springer, 2007.
- [116] R. Bowler, U. Warring, J. W. Britton, B. C. Sawyer, and J. Amini. Arbitrary waveform generator for quantum information processing with trapped ions. Review of Scientific Instruments, 84(3):1–6, 2013.
- [117] M. F. Brandl, P. Schindler, T. Monz, and R. Blatt. Cryogenic resonator design for trapped ion experiments in Paul traps. Applied Physics B: Lasers and Optics, 122(6):1–13, 2016.
- [118] David M Pozar. Microwave Engineering. Wiley, 4th edition, 2011.
- [119] W. W. Macalpine and R. O. Schildknecht. Coaxial Resonators with Helical Inner Conductor. Proceedings of the IRE, 47(12):2099–2105, 1959.
- [120] Yves Colombe, Daniel H. Slichter, Andrew C. Wilson, Dietrich Leibfried, and David J. Wineland. Single-mode optical fiber for high-power, low-loss UV transmission. Optics Express, 22(16):19783, 2014.
- [121] A. C. Wilson, C. Ospelkaus, A. P. Van Devender, J. A. Mlynek, K. R. Brown, D. Leibfried, and D. J. Wineland. A 750-mW, continuous-wave, solid-state laser source at 313 nm for cooling and manipulating trapped 9Be^+ ions. Applied Physics B: Lasers and Optics, 105(4):741–748, 2011.
- [122] T. W. Hansch, I. S. Shahin, and A. L. Schawlow. Optical Resolution of the Lamb Shift in Atomic Hydrogen by Laser Saturation Spectroscopy. Nature Physical Science, 235:63–65, 1972.
- [123] T W Hansch and B Couillaud. Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity. Optics Communications, 35(3):441–444, 1980.
- [124] Sebastian Bourdeauducq. m-labs/artiq: 3.6.
- [125] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. Ion-trap measurements of electric-field noise near surfaces. Reviews of Modern Physics, 87(4), 2015.

- [126] Adam C. Keith, Charles H. Baldwin, Scott Glancy, and E. Knill. Joint quantum-state and measurement tomography with incomplete measurements. Physical Review A, 98(4):1–12, 2018.
- [127] A. H. Burrell. High Fidelity Readout of Trapped Ion Qubits. PhD thesis, 2010.
- [128] Neil W Ashcroft and N David Mermin. Solid State Physics. Brooks/Cole, 1976.
- [129] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. Journal of Applied Physics, 83(10):5025–5033, 1998.
- [130] U Warring, C Ospelkaus, Y Colombe, K R Brown, J M Amini, M Carsjens, D Leibfried, and D J Wineland. Techniques for microwave near-field quantum control of trapped ions. Physical Review A, 013437:1–11, 2013.
- [131] Stephen Crain, Clinton Cahall, Geert Vrijsen, Emma E. Wollman, Matthew D. Shaw, Varun B. Verma, Sae Woo Nam, and Jungsang Kim. High-speed low-crosstalk detection of a $^{171}\text{Yb}^+$ qubit using superconducting nanowire single photon detectors. Communications Physics, 2(1):8–13, 2019.
- [132] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T.C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. Physical Review Letters, 100(20):2–5, 2008.
- [133] D. V. Reddy, A. E. Lita, S. W. Nam, R. P. Mirin, and V. B. Verma. Achieving 98% system efficiency at 1550nm in superconducting nanowire single photon detectors. In Rochester Conf. Coherence Quantum Opt., page W2B.2, Washington, DC, 2019.
- [134] F. Marsili, V. B. Verma, J. A. Stern, S. Harrington, A. E. Lita, T. Gerrits, I. Vayshenker, B. Baek, M. D. Shaw, R. P. Mirin, and S. W. Nam. Detecting single infrared photons with 93% system efficiency. Nature Photonics, 7(3):210–214, 2013.
- [135] Risheng Cheng, Sihao Wang, and Hong X. Tang. Superconducting nanowire single-photon detectors fabricated from atomic-layer-deposited NbN. Applied Physics Letters, 115(24), 2019.
- [136] Yonit Hochberg, Ilya Charaev, Sae Woo Nam, Varun Verma, Marco Colangelo, and Karl K. Berggren. Detecting Sub-GeV Dark Matter with Superconducting Nanowires. Physical Review Letters, 123(15):1–7, 2019.
- [137] B. A. Korzh, Q-Y. Zhao, S. Frasca, J. P. Allmaras, T. M. Autry, E. A. Bersin, M. Colangelo, G. M. Crouch, A. E. Dane, T. Gerrits, F. Marsili, G. Moody, E. Ramirez, J. D. Rezac, M. J. Stevens, E. E. Wollman, D. Zhu, P. D. Hale, K. L. Silverman, R. P. Mirin, S. W. Nam, M. D. Shaw, and K. K. Berggren. Demonstrating sub-3 ps temporal resolution in a superconducting nanowire single-photon detector. pages 1–26, 2018.
- [138] Erik W. Streed, Benjamin G. Norton, Andreas Jechow, Till J. Weinhold, and David Kielpinski. Imaging of Trapped Ions with a Microfabricated Optic for Quantum Information Processing. Physical Review Letters, 106(1):2–5, 2011.
- [139] John Chiaverini and Jeremy M Sage. private communication, 2019.

- [140] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. A quantum engineer's guide to superconducting qubits. Applied Physics Reviews, 6(2):1–66, 2019.
- [141] M. Palmero, R. Bowler, J. P. Gaebler, D. Leibfried, and J. G. Muga. Fast transport of mixed-species ion chains within a Paul trap. Physical Review A - Atomic, Molecular, and Optical Physics, 90(5):1–5, 2014.
- [142] Roy J. Glauber. Coherent and incoherent states of the radiation field. Physical Review, 131(6):2766–2788, 1963.
- [143] T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland. Multi-element logic gates for trapped-ion qubits. Nature, 528(7582):380–383, 2015.
- [144] T Ruster, C Warschburger, H Kaufmann, C T Schmiegelow, A Walther, M Hettrich, A Pfister, V Kaushal, and U G Poschinger. Experimental realization of fast ion separation in segmented Paul traps. Physical Review A - Atomic, Molecular, and Optical Physics, 90(3):1–10, 2014.

Appendix A

Heat load estimation

A.0.0.1 Thermal links

Operation of the trap described in Section 4.2.1 requires 30 electrical connections in addition to the ground reference:

- 22 DC connections to provide DC trapping potentials
- One coaxial connection to drive the RF trapping potential
- One coaxial cable to drive the microwave antenna providing global qubit spin flips
- One coaxial cable for detector bias current and readout
- One DC connection to drive oven current
- Four connections for a four-wire temperature measurement

In general, we wish to be able to test many devices with this apparatus with quick exchange of traps, taking full advantage of the reduced requirements for vacuum hygiene that a cryogenic system enables; therefore, we wish to have many electrical connections, in order to enable functionality with increasingly complex traps. As an example, consider the “triangle trap” described in [47], which has two independent trapping regions with 30 and 36 electrodes each; we would require 66 DC connections and two RF lines in order to provide all trapping voltages. We would then require even more connections to control hardware like oven connections, microwave antennae, temperature measurement, etc.

However, one runs into a trade-off when considering large numbers of electrical connections for cryogenic operation. In normal (non-superconducting) metals, electrical and thermal resistivity are directly related [128]. So the low electrical resistivity we would typically desire for electrical connections (particularly for high-current operations) would also provide low thermal resistivity between the 4 K experimental stage and the 295 K laboratory. We find a solution by considering the specific needs of our electrical connections. The 22 lines driving our DC trapping potential generally do not need high current, since these voltages are generally static; thus we can accept a higher electrical resistance in these lines in order to provide a greater thermal resistance. (This higher resistance *will* provide a filtering effect on any time-varying waveforms applied to the trap electrodes, but that can be compensated for, as I will discuss in Chapter 7) However, the oven lines, which typically need 1 – 1.5 A to generate sufficient neutral beryllium for loading, must be much lower resistance; higher resistance lines would not only heat the cryostat substantially but risk damage with such high currents.

The full wiring configuration in this apparatus is:

- 100 × “high resistance” lines, configured as 4 × 25 wire “looms” (so-called because the wires are woven together, for robustness). The individual wires are 38-gauge constantan, which has thermal resistivity of [xxxx] at room temperature and [xxxx] at 4 K (compared to high purity copper’s [xxxxx], respectively). The looms are configured as 12 twisted pairs with 1 lone wire. The looms are connectorized with micro-D connectors.
- 8 × “low resistance” lines, for ovens and any other applications which require a lower-resistance connection to the cryogenic stage. These low resistance wires are kapton-wrapped 30-gauge copper, connected on either end to a PEEK DSUB connector.
- 4 × semirigid coaxial cables. We use [xxx] coax from [xxx], which has a [xxxx] BeCu core and [xxxx] stainless steel shielding. These cables provide good thermal resistance but are challenging to work with, due to the stiff stainless steel outer conductor. At the room temperature and 4 K stages, the stainless steel cables are connected to short lengths of hand

formable coax, which introduce negligible thermal resistance but enable simpler mechanical connections.

All cables are thermally and mechanically anchored at the 30 K stage and the 4 K stage. The looms and the low resistance lines are wrapped around copper bobbins and potted in Stycast 2850FT, a low-outgassing, thermally conductive but electrically resistive epoxy commonly used for cryogenic potting applications. The four coax cables are thermally anchored by mechanical connections to the shielding of the SMA connectors.

We can estimate the heat load on the 4 K stage due to the cable connections by assuming that the thermal anchors are perfectly efficient. In this case, we can treat each wire or cable as a uniform rod thermally anchored on either end, with the small concession that we must consider the temperature-dependent thermal conductivities $k(T)$ of each material:

$$q = -k(T) \frac{\partial T}{\partial x} \quad (\text{A.1})$$

Although the thermal conductivity $k(T)$ is not an analytical function of temperature, this is a numerically solvable differential equation. There are four configurations we much consider:

- 100x flow from 30 K to 4 K, 38 gauge constantan, 0.5 m: xxx
- 8x flow from 30 K to 4 K, 30 gauge copper, 0.5 m: xxx
- 4x flow from 30 K to 4 K, [xxxx] area stainless steel, 0.5 m: xxx
- Noting that the core of the coaxial cables are not thermally anchored at the intermediate stage, 4x flow from 295 K to 4 K, xxxx gauge BeCu, 1 m: xxx

This gives us in total 100 mW of heat load from the cabling.

A.0.0.2 Black-body radiation

We know from the Stefan-Boltzmann Law that an ideal black-body emitter of surface area A will emit

$$P = \sigma AT^4 \quad (\text{A.2})$$

where $\sigma = 5.67 \times 10^{-8} \text{ Wm}^{-2}\text{K}^{-4}$ is the Stefan-Boltzmann constant. Estimating the cryostat as having a surface area of 1 m^2 , the cryostat would then absorb 430 W from the ambient room temperature black-body radiation ($T = 295 \text{ K}$), unacceptable for a cryostat with cooling power 1.5 W. Recall that this cryostat has an intermediate stage at approximately 30 K, as illustrated in Figure 4.1. This intermediate stage has substantially higher cooling power. We anchor a radiation shield on this intermediate stage, shown in Figure A.1(a) to reflect the bulk of this radiation. Assuming this radiation shield perfectly thermalizes with the 30 K stage, the 4 K cold stage would instead absorb 46 mW of black-body radiation, again estimating 1 m^2 surface area.

A perfect radiation shield is not appropriate for atomic physics experiments, since we require optical access for laser addressing and fluorescence imaging. As seen in Figure A.1(a), the radiation shield has eight 1" viewports for laser access, as well as a 2" viewport on the bottom flange for fluorescence imaging; we cover two with aluminum shields and install windows in the other six. These windows are UV fused silica (UVFS), which is transparent at all relevant laser wavelengths. The UVFS transmission spectrum (shown in blue in Figure A.2) mostly blocks room temperature blackbody radiation; however, there is a small but non-negligible overlap around $4 \mu\text{m}$ wavelength. Integrating over this overlap, the windows transmit 0.07% of incident blackbody radiation, or 1.4 mW.

A.0.0.3 Dissipated RF drive

All Paul traps require an RF drive to trap ions. ${}^9\text{Be}^+$ is a light ion, so we are able to trap with relatively low rf trapping voltages; we further reduce the trapping voltages needed by applying an RF drive with a relatively low (for ${}^9\text{Be}^+$) drive frequency of $\approx 67 \text{ MHz}$. Recall that the trap radial frequencies are proportional to V_0/Ω_{RF} [1], so we can achieve comparable trap stability with lower RF amplitude by lowering amplitude and frequency simultaneously. This comes with the trade-off of reduced trap depths; however, with the very low rate of background gas collisions provided by a cryogenic environment, low trap depths are less of a concern.




Figure A.1: Radiation shields at (a) 30 K and (b) 4 K. Both are vented to the larger cryostat vacuum space, with no line-of-sight from the vents to the ions. Laser access is available through eight viewports; we use six of them, and cover the final two with aluminum shields. Fused silica windows are installed into the other six.




Figure A.2: In pink: the radiation spectrum for a 293K black-body emitter, normalized to arbitrary units. In blue: the transmission spectrum for a 10mm uncoated fused silica window. Note that there is a small but non-negligible overlap around 4 μm .

Appendix B

Trapping Potentials

Table B.1 lists voltages that generate a trap with 2.0 MHz axial frequency in all seven zones of the detector trap according to simulations using the Biot-Savart method for simulating trap potentials. These can be adapted to trapping in the transport trap by grounding the ‘det’ electrode.

Electrode	Zone 2	Zone 3	Zone 4	Zone 5	Zone 6	Zone 7	Zone 8
DC1	0.65V	0.67V	0.61V	0.23V	0.20V	0.23V	0.2V
DC2	-1.38	0.13	0.29	0.41	0.24	0.20	0.14
DC3	0.11	-1.07	-0.23	0.56	0.44	0.36	0.23
DC4	1.26	-0.22	1.19	-0.35	0.52	0.62	0.38
DC5	0.55	1.13	-0.35	-1.21	-0.39	0.65	0.45
DC6	0.12	0.63	0.47	-0.18	-1.17	-0.25	0.39
DC7	-0.04	0.11	0.35	0.31	-0.13	-1.06	-0.45
DC8	-0.08	-0.09	0.15	0.30	0.28	-0.08	-0.96
DC9	-0.34	-0.56	0.08	0.61	0.58	0.57	0.28
DC10	0.54	0.61	0.58	0.21	0.19	0.23	0.20
DC11	-1.54	0.04	0.24	0.38	0.22	0.19	0.13
DC12	-0.07	-1.22	-0.33	0.49	0.40	0.35	0.23
DC13	1.15	-0.35	-1.33	-0.48	0.44	0.58	0.37
DC14	0.50	1.06	-0.46	-1.37	-0.54	0.54	0.42
DC15	0.10	0.60	0.41	-0.29	-1.32	-0.43	0.34
DC16	-0.05	0.10	0.33	0.26	-0.21	-1.24	-0.53
DC17	-0.09	-0.09	0.14	0.28	0.25	-0.18	-1.05
DC18	-0.34	-0.57	0.08	0.59	0.56	0.50	0.17
CUL	0.04	-0.01	-0.08	-0.07	-0.07	-0.01	0.03
CLL	0.15	0.08	-0.01	0.01	0.01	0.09	0.01
CUR	-0.01	0.14	0.32	-0.06	-0.07	-0.01	-0.06
CLR	-0.01	0.14	0.32	-0.06	-0.06	0.13	0.05

Table B.1: Trap voltages that generate a trap with 2.0 MHz axial frequency in all seven trap zones.

Appendix C

Code

Appendix D

Derivation of transport as a displacement operator

In a quantum mechanical picture, the ion motion is described by the Hamiltonian

$$\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega^2(\hat{x} - X(t))^2 \quad (\text{D.1})$$

where, again, we are considering a rigid well with unchanging trap frequency ω with time-dependent well minimum position $X(t)$. Following the derivation in [71], we can shift into the moving frame of the well using the translation operator

$$\hat{T}(X(t)) = \exp\left(-\frac{iX(t)\hat{p}}{\hbar}\right) \quad (\text{D.2})$$

Consider the Schrödinger Equation acting on the translated wavefunction $|\chi\rangle \equiv \hat{T}(X(t))|\psi\rangle$:

$$\begin{aligned} i\hbar\frac{\partial}{\partial t}|\chi\rangle &= i\hbar\frac{\partial}{\partial t}\left(\hat{T}(X(t))|\psi\rangle\right) \\ &= i\hbar\frac{\partial}{\partial t}\left(\hat{T}\right)|\psi\rangle + i\hbar\hat{T}\frac{\partial}{\partial t}|\psi\rangle \\ &= i\hbar\frac{\partial}{\partial t}\left(\exp(-iX(t)\hat{p}/\hbar)\right)|\psi\rangle + i\hbar\hat{T}\hat{H}|\psi\rangle \\ &= \dot{X}(t)\hat{p}\hat{T}|\psi\rangle + \left(\hat{T}\hat{H}\hat{T}^\dagger\right)\hat{T}|\psi\rangle \\ &= \left(\hat{H}_0 + \dot{X}(t)\hat{p}\right)|\chi\rangle \end{aligned} \quad (\text{D.3})$$

Move into the interaction picture $|\tilde{\chi}\rangle \equiv \exp(-i\hat{H}_0t/\hbar)|\chi\rangle$, so

$$\begin{aligned} \frac{\partial}{\partial t}|\tilde{\chi}\rangle &= i\hbar\frac{\partial}{\partial t}\left(\exp(i\hat{H}_0t/\hbar)|\chi\rangle\right) \\ &= -\hat{H}_0\exp(i\hat{H}_0t/\hbar)|\chi\rangle + i\hbar\exp(i\hat{H}_0t/\hbar)\frac{\partial}{\partial t}|\chi\rangle \\ &= -\hat{H}_0\exp(i\hat{H}_0t/\hbar)|\chi\rangle + \exp(i\hat{H}_0t/\hbar)\left(\hat{H}_0 + \dot{X}(t)\hat{p}\right)|\chi\rangle \end{aligned} \quad (\text{D.4})$$

Noting that $[\hat{A}, f(\hat{A})] = 0$, this simplifies to

$$\begin{aligned} \frac{\partial}{\partial t} |\tilde{\chi}\rangle &= \exp(i\hat{H}_0 t/\hbar) \dot{X}(t) \hat{p} \left(\exp(-i\hat{H}_0 t/\hbar) \exp(i\hat{H}_0 t/\hbar) \right) |\chi\rangle \\ &= \dot{X}(t) \left[\exp(i\hat{H}_0 t/\hbar) \hat{p} \exp(-i\hat{H}_0 t/\hbar) \right] |\tilde{\chi}\rangle \end{aligned} \quad (\text{D.5})$$

Note that the operator $\hat{p} \equiv \exp(i\hat{H}_0 t/\hbar) \hat{p} \exp(-i\hat{H}_0 t/\hbar)$ is a momentum operator in the frame rotating with \hat{H}_0 relative to the moving frame.

Up until, this derivation has been entirely general to any rigid moving potential. If the ion is in a harmonic oscillator potential, then, in a frame rotating with angular frequency ω

$$\begin{aligned} \hat{H}_0 &= \hbar\omega \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) \\ \hat{p} &= i\sqrt{\frac{\hbar m\omega}{2}} \left(\hat{a}^\dagger e^{i\omega t} - \hat{a} e^{-i\omega t} \right) \\ \hat{x} &= \sqrt{\frac{\hbar}{2m\omega}} \left(\hat{a}^\dagger e^{i\omega t} + \hat{a} e^{-i\omega t} \right) \end{aligned} \quad (\text{D.6})$$

where \hat{a}^\dagger , \hat{a} are the creation and annihilation operators, respectively. So

$$\begin{aligned} \frac{\partial}{\partial t} |\tilde{\chi}\rangle &= i\hbar\sqrt{\frac{m\omega}{2\hbar}} \dot{X}(t) \left(\hat{a}^\dagger e^{i\omega t} - \hat{a} e^{-i\omega t} \right) \\ &= -i\hbar \left(\dot{\alpha}(t) \hat{a} - \dot{\alpha}^*(t) \hat{a}^\dagger \right) \end{aligned} \quad (\text{D.7})$$

where

$$\dot{\alpha}(t) \equiv \dot{X}(t) e^{-i\omega t} \sqrt{\frac{m\omega}{2\hbar}} \quad (\text{D.8})$$