

**Transport of Trapped-Ion Qubits within a Scalable
Quantum Processor**

by

R. Bradford Blakestad

B.S., California Institute of Technology, 2002

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

2010

This thesis entitled:
Transport of Trapped-Ion Qubits within a Scalable Quantum Processor
written by R. Bradford Blakestad
has been approved for the Department of Physics

Dr. David J. Wineland

Prof. Murray J. Holland

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.

Blakestad, R. Bradford (Ph.D., Physics)

Transport of Trapped-Ion Qubits within a Scalable Quantum Processor

Thesis directed by Dr. David J. Wineland

Key requirements for efficient large-scale quantum computing (QC) include reliable transport of information throughout the processor and the ability to perform gates between arbitrarily selected qubits. Trapped ions are a useful system for studying the elements of QC and can potentially satisfy these requirements. The internal atomic state of the ion can be used as a qubit. Quantum gate operations can be performed with electromagnetic fields and entangling operations are achieved by using those fields to map information between the qubits and a shared motional mode of the trapped ions. One approach to extending this to large-scale QC would be to distribute ions over separate zones in an array, where information would be shared between zones by physically transporting the ions.

Multidimensional arrays incorporating junctions would enable ions selected from arbitrary locations to be grouped together for multi-qubit gates. However, excitation of the harmonic motion of the ions during transport can negatively impact computational fidelity. Here I report a novel ion-trap design including an ‘X’ junction. Reliable transport of ${}^9\text{Be}^+$ ions through this array has been demonstrated with motional excitation corresponding to less than one quantum in the local confining potential. The process of designing and building the trap, as well as producing the control potentials that shuttle the ions will be described. In addition, sources of motional excitation will be discussed.

This work is an important step in demonstrating the scalability of this particular approach to trapped-ion QC. Such techniques should provide a path toward transfer of information in a larger-scale quantum processor and the potential to employ an increased number of qubits in quantum algorithm experiments.

Dedication

To my parents and grandparents,
for showing me what is important and for all the support over the years.

Acknowledgements

Throughout my dissertation work, I have had the honor and privilege of working with and learning from a lot of talented colleagues. First and foremost, I thank my adviser, Dave Wineland for his guidance. He has unparalleled knowledge and wisdom on everything related to trapped ions. Dave was quick to spot problems in my work and point me in the correct direction. In addition, Dave does an amazing job of managing one of the preeminent atomic physics groups in the world and I learned a lot as part of this efficient research machine.

Didi Leibfried also deserves recognition for his assistance. He was always there to answer those silly graduate-student questions and has an uncanny knack for pinpointing what is going wrong in the experiment and fixing it. Without him, this dissertation would have dragged on many more years.

I had the benefit of working closely with a number of excellent postdocs. Christian Ospelkaus has been a great friend and colleague. He made significant contributions to the experimental setup and was always willing to help with those details that were so important to the experiment's success. Aaron VanDevender was a big help, both with the lasers and in taking the actual data. And Jason Amini's provided invaluable assistance with the computer systems. I also appreciate Janus Wesenberg's involvement in the trap modeling and waveform software. I enjoyed working with Signe Seidelin while I fabricated the X-junction trap, and I am sure that she shares my contempt for electroplating.

Joe Britton, my office mate for the majority of my graduate career, had a significant impact on my project. He was always willing to help work through a problem or explain whatever it was

that I did not understand. Joe has also been great for bouncing ideas off of, no matter how crazy the ideas were.

In addition to those singled out above, I had the opportunity to learn from a number of other great colleagues in the NIST Ion Storage Group: Jim Bergquist, Sarah Bickman, Mike Biercuk, John Bollinger, Ryan Bowler, Kenton Brown, Anders Brusch, John Chiaverini, James Chou, Yves Colombe, Bob Drullinger, Ryan Epstein, David Hanneke, Jonathan Home, David Hume, Wayne Itano, Brana Jelenkovic, John Jost, Manny Knill, Jeroen Koelemeij, Chris Langer, Robert Maiwald, Roe Ozeri, Rainer Reichle, Till Rosenband, Piet Schmidt, Nobu Shiga, Mike Thorpe, Hermann Uys, and Ulrich Warring. It has truly been great to work with such a wonderful group of people, who have made my graduate-school years enjoyable and successful.

Finally, I would like to acknowledge my dissertation committee: Dave Wineland, Murray Holland (my titular adviser), Frank Barnes, Manny Knill, Scott Parker, and Jun Ye.

Contents

Chapter	
1 Introduction	1
2 ${}^9\text{Be}^+$ qubits - atom/laser interactions	12
2.1 ${}^9\text{Be}^+$ atomic structure	12
2.2 State detection	15
2.3 State initialization	17
2.4 Doppler cooling	20
2.5 Single-ion coherent operations	23
2.5.1 Single-photon operations	23
2.5.2 Two-photon stimulated-Raman operations	29
2.5.3 Ground state cooling	33
2.5.4 Shelving	34
2.6 Two-ion coherent operations	35
2.7 ${}^{24}\text{Mg}^+$ ions	37
3 Linear rf Paul ion traps	38
3.1 Static confining potential	39
3.2 rf Paul trap	41
3.2.1 Adiabatic approximation	43
3.2.2 Linear Paul trap	45

3.3	Mathieu equation	48
3.4	Inhomogenous Mathieu equation	54
3.4.1	Mathieu equation with monopole and dipole terms	55
3.4.2	Mathieu equation with nonlinear terms	58
3.5	Real trap geometries	63
3.5.1	Common linear trap architectures	66
3.5.2	Comparison of multipole terms in different trap geometries	71
4	Junction architectures and the X-junction array	77
4.1	Junction architectures	78
4.1.1	Two-layer junctions	78
4.1.2	Three-layer junctions	82
4.2	X-junction array	84
4.2.1	X-junction	87
4.2.2	Experiment zone	87
4.2.3	Load zone	89
5	Ion transport	92
5.1	Transport introduction	93
5.2	Modeling of the trap potentials	96
5.3	Solving for a single harmonic potential	98
5.3.1	Singular value decomposition and the pseudoinverse	99
5.3.2	Constraints	107
5.3.3	Issues	113
5.4	Solving for transport waveforms	117
5.4.1	Problems encountered solving for waveforms	120
5.4.2	Linear least-squares solution method	122
5.5	Other types of waveforms	125

5.6	Waveforms for the X-junction array	126
6	Apparatus	136
6.1	Ion trap	136
6.1.1	Trap fabrication	138
6.1.2	Assembly	143
6.2	Trapping potential	146
6.2.1	rf potential	146
6.2.2	Control potentials	150
6.3	Vacuum system	156
6.4	Loading	160
6.5	Lasers and optics	163
6.5.1	Doppler lasers	165
6.5.2	Raman lasers	168
6.5.3	Light detection	171
6.6	Controlling the experiment	172
7	Demonstration of transport through the X-junction	173
7.1	Experiment procedure	173
7.2	Transport success rates	175
7.3	Motional excitation rates	178
7.4	Qubit coherence	180
7.5	Transporting $^{24}\text{Mg}^+$	181
7.6	Mechanism that excite the secular motion	183
7.6.1	rf-noise heating mechanism	183
7.6.2	DAC update noise	186
7.6.3	Other mechanisms	188
7.7	Possible improvements	190

8	Conclusion	191
	Bibliography	193
	Appendix	
A	Glossary of acronyms and symbols	203
B	Methods for modeling trap potentials	206
C	Possible linear constraints	212
D	Maximum waveform step size	217
E	Measuring motional excitation	223

Figures

Figure

1.1 Ions in a linear array.	6
1.2 Possible multiplexed ion-trap array.	8
2.1 Energy level diagram for ${}^9\text{Be}^+$	14
2.2 Detection Histograms.	16
2.3 Energy level diagram with resonant lasers.	19
2.4 Rabi flopping of a single ion between the bright and dark qubit states.	27
2.5 Qubit energy level diagram that includes the Raman beam energies.	30
2.6 Raman beam geometry.	31
3.1 Diagram of a linear rf Paul trap.	42
3.2 Solutions to the Mathieu equation.	51
3.3 Two-wafer trap.	67
3.4 Surface-electrode trap geometries.	70
3.5 Various trap geometries.	74
3.6 Multipole components for a four-wire surface trap.	76
4.1 Two-layer junctions.	79
4.2 Axial pseudopotential near the X-junction.	81
4.3 Three-layer junctions.	82
4.4 X-junction array trap wafers.	84

4.5	Diagram of the X-junction array.	86
4.6	Diagram of the load zone and oven shield.	90
5.1	Sample BEM mesh for modeling the trap potentials.	96
5.2	Different transport regions in the X-junction array.	127
5.3	Sample waveform moving into the X-junction.	130
5.4	Secular frequencies during a transport into the X-junction.	132
5.5	Mathieu a_i and q_i parameters near the X-junction.	134
5.6	Micromotion and adiabaticity parameter near the X-junction.	135
6.1	Cross section of the trap wafers.	137
6.2	Photograph of the X-junction array.	139
6.3	Electroplating setup.	142
6.4	In-vacuum RC filter board.	145
6.5	Photograph of the connections between the X-junction array and the filter board.	146
6.6	rf setup.	149
6.7	Electronic filters for the control voltages.	151
6.8	Generic filter network.	152
6.9	Frequency response of different types of filter.	154
6.10	Comparison of the new and old control filters.	155
6.11	Photograph of the vacuum system.	157
6.12	Photograph of the X-junction array inside the rf resonator.	158
6.13	Glass vacuum envelope.	159
6.14	Picture of the trap inside of the glass envelope.	159
6.15	Building Mg ovens.	161
6.16	Oven assembly and detachable resonator piece.	162
6.17	Laser beam orientation relative to the trap.	164
6.18	Doppler beam setup.	166

6.19 Raman beam setup.	169
7.1 Four destinations for transport.	174
7.2 Reliability for transporting through the X-junction.	177
7.3 Test of decoherence while transporting through the X-junction.	182
7.4 Heating due to rf noise.	185
7.5 Plot showing motional excitation at various DAC update frequencies.	187
E.1 Fits to sideband flopping traces to extract the motional temperature.	225

Chapter 1

Introduction

Quantum computation (QC) has become a popular field in physics during the last two decades. Quantum algorithms have been shown to outperform known classical algorithms for several types of problems. One such quantum algorithm is Grover's search algorithm [Grover 97] which can be used to search an unstructured list and scales more efficiently with the size of the search-space than any known classical algorithm. Similarly, Shor's factoring algorithm [Shor 94] can outperform all known classical algorithms for factoring very large numbers, which has implications in public-key cryptography. Of particular interest to physicists is the ability of quantum computers to simulate other quantum systems, a task that is inefficient on classical computers [Feynman 82]. With a large number of researchers focused on quantum computing, our theoretical understanding of this field grows daily and the hope is that many new uses for quantum computers will be found in the coming years.

The excitement over quantum computing is somewhat muted by the current lack of a machine capable of running large-scale quantum algorithms. A number of physical systems have been proposed as potential platforms in which to build a quantum computer and there has been a great deal of promising research to this end in many of these systems. However, it has proven difficult to achieve the exacting control needed to operate on quantum bits (qubits). Furthermore, qubits are prone to decoherence which destroys the information stored in the qubits and can act on a very fast time-scale. In addition, a large number of qubits are needed to compete with modern classical computers on tasks such as factoring. Quickly and precisely controlling many (thousands) qubits

while subject to decoherence is a difficult technical challenge.

In an attempt to clarify which physical systems hold promise as a quantum computer, DiVincenzo specified a number of criteria that the system must satisfy [DiVincenzo 98]:

- (1) The system must contain two-state quantum subsystems that are well understood and characterized. These subsystems will constitute the qubits.
- (2) There must be an ability to initialize all qubits into a known initial state, such as $|000 \dots 00\rangle$.
- (3) The decoherence times of the qubits must be long compared to the gate operation times. This allows information stored in the qubits to persist long enough to be operated on.
- (4) A universal set of quantum gates must exist, which will allow any quantum algorithm to be realized by successive applications of these gates. This universal set includes arbitrary single-qubit rotations that change the state of a single qubit. The universal set also must include at least one two-qubit gate operation which can entangle the states of any two selected qubits.
- (5) There must be a method of reading out the state of any qubit.
- (6) Finally, the system must be scalable so that, if the previous criteria can be demonstrated on a small number of qubits, there is a straight-forward approach for scaling the system to a large number of qubits while maintaining the above criteria.

In 1995, Cirac and Zoller suggested that cold trapped ions could be used to realize a quantum computer [Cirac 95]. They took advantage of a significant body of previous research focused on controlling and measuring the quantum states of trapped ions with the goal of realizing an ion-based atomic frequency standard. Cirac and Zoller envisioned qubits encoded in the atomic states of ions confined within a harmonic potential provided by an rf Paul trap. Laser light would be used to interact with and control these qubits. Since the proposal, trapped ions have become one of the leading systems for studying the elements of quantum computing.

This proposal satisfies many of the DiVincenzo criteria. Requirement (1) is achieved by using a subset of the atomic states of ions to create qubits. For example, two hyperfine ground states could be used. Such states can have extremely long decoherence times, on the order of minutes [Bollinger 91], which satisfies requirement (3). State initialization would be performed using optical pumping. Readout of the atomic state is possible through state-dependent resonance fluorescence. Here, properly tuned light would resonantly scatter from an ion in one of the qubit states but would be off-resonant for an ion in the other qubit state, and thus not scatter. This allows a measurement of the ion's state based on whether or not the ion scatters photons. Thus, DiVincenzo criteria (1), (2), (3) and (5) are met.

The universal set of quantum gates, requirement (4), can also be achieved. Resonant light will induce Rabi flopping between the qubit states and can be used to realize single-qubit rotations. Two-qubit entangling operations are more difficult to achieve because the atomic states of different ions do not strongly interact. However, light-induced forces can be used to couple the atomic states to the motional state of the ion [Wineland 98]. The mutual repulsion of multiple charged ions, confined in the same harmonic trapping-potential, leads to shared motion. Cirac and Zoller outlined a method for mapping the qubit information between atomic states and a joint motional mode to produce an entangling operation between multiple ion qubits. This completes the required universal set of quantum gates. Subsequent proposals have suggested different implementations of entangling gates that are more robust [Mølmer 99, Milburn 00, Solano 01, Leibfried 03] but still operate on the same principal of a light-mediated interaction that allows the shared motion of the ions to be used as an information bus between ions that would otherwise not interact in the atomic-qubit basis.

These first five DiVincenzo criteria have been experimentally demonstrated. However, significant work is still needed to increase gate fidelity and further suppress decoherence. A review of the recent progress in this field can be found in [Blatt 08, Monroe 08, Häffner 08].

The scalability requirement is the last DiVincenzo criterion which remains to be demonstrated with trapped-ions due to the difficulty of performing the two-qubit entangling operations when a

large number of ions are present. Most of the previously demonstrated entangling gates require spectroscopically isolating a single mode of motion. If the laser interaction couples to more than the single intended mode, the qubit information can leak into the other modes where it would be lost. The result is a drop in the gate fidelity. Two ions will have six motional modes and the trapping parameters can be set so that these modes are easy to individually address. However, as the number of ions (and qubits) increases, the number of modes increases linearly while the separation between the frequencies of these modes decreases. At some point, around a few tens of ions, the mode frequencies become too tightly packed to be spectroscopically isolated [Hughes 96, Wineland 98]. Large-scale quantum computing will require far more than ten ions and, thus, cannot be achieved using the above technique.

To achieve scalability, an approach other than placing all the ions into the same harmonic trapping potential must be taken. There are a couple of proposals to address this. A non-harmonic potential could be used to handle more ions [Lin 09]. The approach pursued in this dissertation will involve multiple harmonic trapping regions distributed across a spatial array. By limiting the number of ions in each harmonic potential, high-fidelity gate operations could be performed, while the number of such regions could be increased to accommodate an arbitrary quantity of ion qubits.

Quantum gates may need to be performed between any arbitrarily selected qubits. If those qubits are separated by significant distances, as would be likely if they are distributed across different trapping regions in a large-scale array, two-qubit gate operations cannot be directly applied between the distant qubits. Instead, a method for transferring information between trapping regions must be constructed.

Transfer of quantum information within a quantum processor is a general problem faced by all quantum computing platforms. As any system grows in scale, there must be a fast and reliable means for moving qubits into locations where they can interact with each other. It is important that this process does not significantly slow down the processor. At the same time, this transport cannot introduce significant decoherence or errors into the qubit information.

It should be noted that proposals to use teleported gates and measurement-based quantum

computing [Gottesman 99] require distributed entanglement. The process of distributing such entanglement would involve the same types of operations as information transfer in the circuit-based model considered here.

One proposal for transporting information within the distributed array model is to map the qubit information from the ions onto some carrier system, such as a photon, that can quickly move between trapping regions [Duan 10]. This approach was demonstrated by probabilistically entangling two ions located in separate vacuum systems using photons coupled into fibers [Moerhing 07]. These experiments are currently limited by inefficient fiber coupling, but the use of cavities is expected to help.

The method of transporting information that will be pursued in this dissertation involves physically moving the ions, themselves, between trapping regions. This method was proposed in [Wineland 98, Kielpinski 02] and extends the simple linear rf Paul trap, that will be discussed later in Ch. 3, to form much more complex trapping structures. These intricate structures allow the potential to be changed and the ions moved within the trap.

The standard linear rf Paul trap (see Fig. 3.1) creates a nearly harmonic three-dimensional trapping potential using a combination of static and rf electric potentials. The rf potential provides strong confinement in two directions (the radial directions), while providing no confinement along the axial direction of the trap. In the absence of the static potentials, the ions would be free to travel along a line on the trap axis. The static potentials (which we will refer to as the ‘control’ potentials since they will later be varied for ion transfer) are introduced to create a harmonic trapping potential along the axis, which is relatively weak compared to the radial confinement. It is the axial motion that will be used to perform gate operations. Due to the trap anisotropy, all ions in this potential will line up along the trap axis. By shifting the control potentials, this trapping well¹ can be moved along the trap axis to transport the ions inside, as demonstrated in [Rowe 02].

In the simplest linear Paul trap, there is only one harmonic trapping well and all the ions would be confined in this well. By extending the trap in the axial direction, and providing more

¹ The term ‘trapping well’ is used to mean a local minimum of the potential.

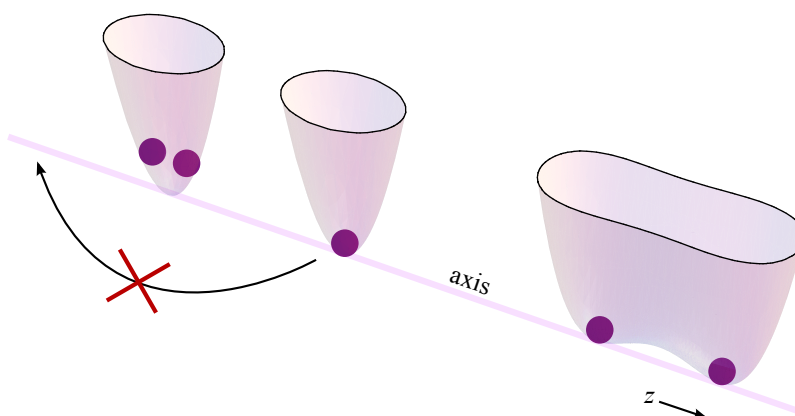


Figure 1.1: A linear array of trapping well. These wells, and the ions inside of them, are located on the axis of the trap array. Here four trapping wells are seen, one that holds a single ion, one that has a pair of ions, and one that is in the process of separating to become two wells. It is not possible to switch the order of the ions without moving off of the axis, though incorporating a junction into the linear array will allow this.

control electrodes, additional trapping wells along the trap axis can be created simultaneously. (These wells are deep enough and distant enough that tunneling between wells is ignored.) Continuing to extend the trap along the axial direction will allow for very large linear arrays of trapping wells.

To allow for two-qubit gates between arbitrarily chosen qubits, there needs to be a means for placing the corresponding ions into the same trapping well, even if the ions start in different wells. This requires the ability to combine ions from two separate wells into a single well and separate a single well into two wells. These types of procedures have been demonstrated in [Rowe 02, Barrett 04, Hensinger 06]. Joining sequences of transport along the axis with separation and combination of the wells will allow the ions to be dynamically grouped so that only ions that are to participate in the two-qubit gate are present in a given well during the gate operation.

These three types of transfer protocols are not sufficient for the arbitrary reordering needed to perform quantum computations. Since the trapping wells, and the ions inside of them, form a linear array along the trap axis (see Fig. 1.1), there is no way to switch the ion order without moving off of that axial line into a second dimension. If two ions need to interact, but are separated by 10,000 other ions on the axis, no sequence of moving the wells up and down the axial line, while

combining and separating, will result in these two ions ending up in the same trapping well without 10,000 other ions also in that well.

A recent experiment [Splatt 09] demonstrated order switching in a linear array by pushing a small chain of ions (two or three) off the trap axis using the control electrodes. By properly adjusting the voltages on these electrodes, the chain of ions could be made to rotate, changing their order. This method may be useful in some situations, but would require a large number of transfer operations (including a large number separating and combining operations) as the number of qubits increases.

The approach to solving the problem of order switching pursued here involves the use of a junction array. A junction array consists of two linear arrays that intersect to form a junction. Potential junction geometries include ‘X’, ‘T’, and ‘Y’ junctions, among others. With a junction array, ions can be moved along a linear array until they reach the junction, at which point they can either continue along their original linear path or turn into the second linear array. Using the example of two chosen ions separated by 10,000 ions, the first chosen ion could make the turn into the new array, then the next 10,000 ions could move straight through the junction, and the second chosen ion could also make the turn. At that point, the ions could be combined in the same trapping well and the gate operation could be performed. Using junctions would eliminate the need for most of the separating and combining operations required for the [Splatt 09] method. In addition, the two-dimensional architecture allowed by junctions would enable simultaneous transport of ions along different linear arrays to more efficiently reorder the ions.

Figure 1.2 depicts a possible architecture that follows [Wineland 98, Kielpinski 02] and includes four T-junctions joined by linear arrays. One region is designated as the interaction zone where lasers perform operations on the ions. The rest of the architecture allows the ions to be rearranged and transported into the interaction zone for the gate operations. In fact, while laser interactions proceed in that zone, the rest of the array can be used to rearrange the other ions in preparation for the next gate operation. The eventual goal would be to combine many such structures to create a larger-scale array that incorporates more junctions and interaction zones to

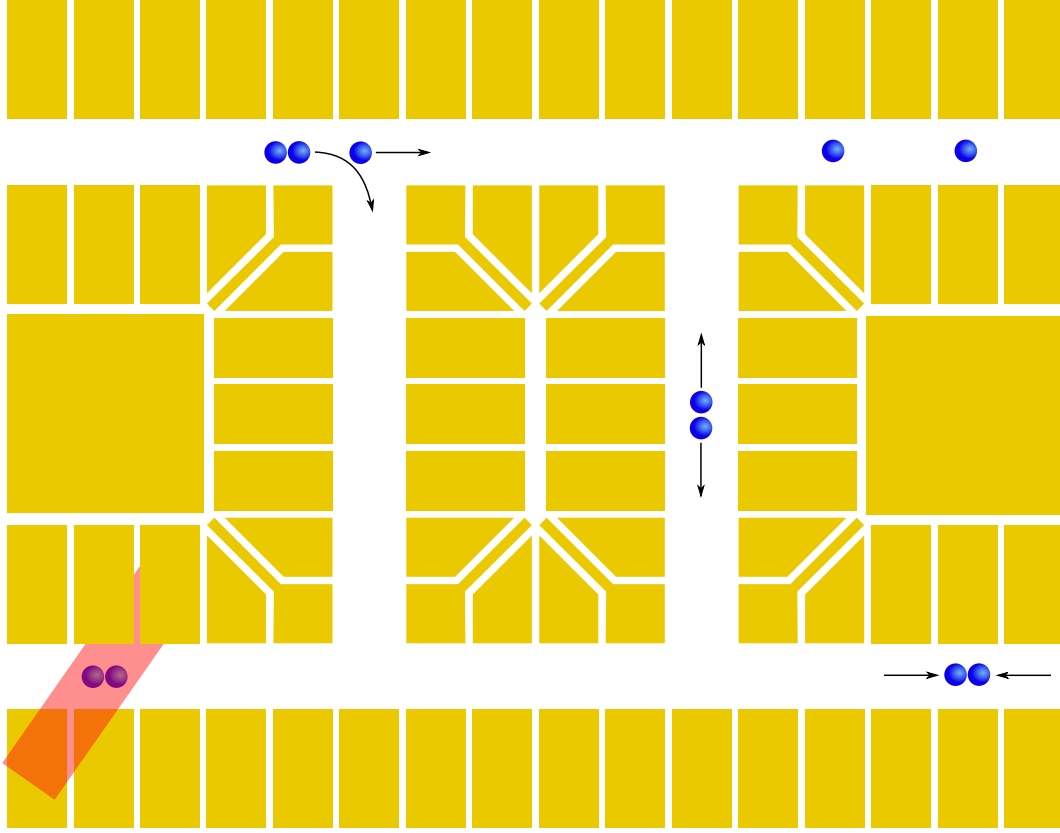


Figure 1.2: Cartoon depiction of a possible multiplexed array of trapping zones for ions. The rectangular shapes are meant to represent control electrodes (rf electrodes are not considered in this diagram, though would be needed in a real trap). The ions can be moved along the channels by adjusting the voltages on the electrodes. In the bottom left is a laser interaction zone where two ions are undergoing a gate operation. To perform gates between arbitrarily selected ions, those ions must be moved to this zone. This array incorporates four T-junctions with ions being transported through one junction in the top left. Other protocols, like separation and combination of trapping wells, are also depicted.

manipulate larger numbers of ions.

The four transport protocols required when using a multiplexed architecture like that suggested in [Kielinski 02] are: transporting ions in a linear array, separating ions from one trapping well into two wells, combining ions from multiple wells into a single well, and, finally, moving ions through junctions. All of these manipulations must be performed with high success probability so that qubits are not lost during transport. Furthermore, the transport must not directly decohere the atomic qubit state of the ion and fast transport will reduce the amount of time the qubits have

to decohere from other processes. In addition, it is important that the harmonic motion of the ions in their respective trapping wells is not excited during these transport procedures. Achieving high fidelity for the two-qubit quantum gates requires this harmonic motion be in a known starting state. To realize this, every experiment begins with the ions being laser cooled to the ground state of motion. If the transport excites this motion during the course of the experiment, the fidelity of subsequent gate operations will be reduced.

One may ask why an ion cannot simply be recooled if it becomes motionally excited during the experiment. Unfortunately, the laser cooling procedure involves a dissipative process that would destroy any information stored in the atomic states of the ion. A second ion species can be introduced and used to cool the qubit ions via sympathetic cooling [Barrett 03, Jost 09, Home 09]. Since the modes of motion are shared between ions in the same trapping well, cooling one species will cool the other. The lasers needed to interact with the sympathetic ion species will, in general, be so far off resonance with the qubit ion that they will not affect the qubit information. Sympathetic cooling is a tool that will almost certainly be needed for large-scale quantum computing to remove whatever motional excitation is produced from transport or other mechanisms. However, sympathetic cooling is slow and can take tens or hundreds of microseconds to remove a few quanta of motional excitation. During this cooling, the ion qubits will continue to decohere from unrelated processes (such as magnetic field noise). Thus it is important to minimize the use of sympathetic cooling through limiting the source of motional excitation. Therefore, the requirement that transport protocols cause minimal motional excitation still holds when using sympathetic cooling.

Returning to the four transport protocols needed for the multiplexed architecture, three have already been demonstrated. Transport through linear array was shown to be highly reliable, preserve the qubit coherence, and cause minimal motional excitation [Rowe 02]. Separating and combining of trapping wells has also been demonstrated [Rowe 02, Hensinger 06] and seen to be reliable while preserving qubit coherence [Barrett 04, Home 09], though, so far, has led to motional excitation on the order of a few quanta. This excitation is too large for good gate fidelities, but small enough that sympathetic cooling can be performed in a short period of time. These three

types of transport can be performed in tens of microseconds, while decoherence times can be as long as 15 s for a ${}^9\text{Be}^+$ qubit [Langer 05].

The final protocol required is the transport of ions through junctions. Unfortunately, this is also likely the most difficult transport protocol to perform. The trap potentials in the vicinity of a junction are much more complicated than those used in a linear array or for separation. This dissertation is focused on the challenge of junction transport to show that the multiplexed architecture of [Kielpinski 02] is a reasonable approach to achieving scalability in trapped-ion systems. An X-junction array was designed, fabricated, and tested with this goal [Blakestad 09].

A similar experiment using Cd^+ in a T-junction array was performed at the University of Michigan during roughly the same time period [Hensinger 06]. This was the first experiment to successfully demonstrate transport of ions through a two-dimensional array and implemented an ion order-switching protocol. However, this transport was accompanied by motional excitation of about 1 eV (10^8 quanta), which is significantly higher than acceptable for quantum computing. In addition, the round-trip failure rate for moving through the junction was 2%, also unacceptable for general computation.

The experiments in this dissertation demonstrate the successful transport of ${}^9\text{Be}^+$ (and ${}^{24}\text{Mg}^+$) in an X-junction array. The motional excitation for junction transport obtained with this trap was less than 1 quantum with a failure rate of $< 0.01\%$. This constitutes the first time that junction transport has been performed with excitation rates low enough that it could be used in a quantum processor. Since this is the last transport protocol needed for realizing the multiplexed architecture, and the multiplexed architecture is a way to demonstrate the last DiVincenzo requirement for trapped-ion quantum computing, it is an important step towards realizing scalable quantum computing.

This dissertation will explore this X-junction array and explain how these transport results were achieved. It will begin with a discussion of the ${}^9\text{Be}^+$ qubit and how quantum operations are performed on this ion. This is followed by a chapter on the theory of ion traps, while Ch. 4 will discuss the design of junctions along with the specific design of the X-junction array used in this

dissertation. Chapter 5 will be dedicated to the complicated series of potentials applied to the electrodes to transport ions through the array. The apparatus, including the ion trap, the vacuum system, and the laser systems will be explained in Ch. 6. The experimental results and a detailed review of the processes that caused excitation of the motion during transport through the junction will be covered in Ch. 7.

Chapter 2

${}^9\text{Be}^+$ qubits - atom/laser interactions

Quantum information processing requires information to be stored and manipulated in a quantum system. Most of this work, both theoretical and experimental, involves the use of simple two-level quantum systems, where information can be stored in the form of quantum bits, or qubits. In our system, we use the atomic energy states of a single trapped ${}^9\text{Be}^+$ ion to form a basis in which we store quantum information. Though the energy levels of ${}^9\text{Be}^+$ are numerous, through the judicious use of laser light, we can restrict the ion to just two states in the ground-state hyperfine manifold, thus approximating the two-level system of a qubit. Further use of laser light to couple to additional atomic states allows the initialization, readout, and manipulation of the qubit states. The shared vibrational motion of two or more ions, held in close proximity, can act as a bus for transferring information between the ions or entangling them. This requires that motion be initialized to a known starting state, which implies cooling the ions to the ground state for optimum performance. Again, laser light plays an important role in this cooling, to provide the coupling of atomic energy states to that motion. In this chapter, we will describe the atomic structure of ${}^9\text{Be}^+$ and the couplings and use of lasers to perform the procedures listed above.

2.1 ${}^9\text{Be}^+$ atomic structure

Alkali and alkali-like atoms have historically been a popular choice for atom cooling and trapping experiments. The single valence electron gives a simple and well-studied structure that is ideal for many experiments. Once singly ionized, the alkaline earth metals exhibit atomic structure

analogous to neutral alkali atoms, making them a popular choice for trapped ion experiments. The majority of work in this thesis involves the ${}^9\text{Be}^+$ ion (though some initial work was done with ${}^{24}\text{Mg}^+$, whose ion's atomic structure is discussed in Sec. 2.7). The relevant atomic structure is depicted in Fig. 2.1. The ground state level is a $2s^2S_{1/2}$ state (or ${}^2S_{1/2}$ for short) while the lowest excited states are P states. The nuclear spin of ${}^9\text{Be}^+$ ($I = \frac{3}{2}$) couples with the valence electron's angular momentum to form hyperfine structure. At low magnetic field, the ${}^2S_{1/2}$ states are split into two hyperfine manifolds, $F = 1$ and $F = 2$, spaced by ~ 1.25 GHz [Wineland 83], where the total angular momentum ($F = I + J$) is a good quantum number. Application of a magnetic field (~ 13 G) both defines a quantization axis and breaks the degeneracy of the m_F levels via the linear Zeeman effect. This splits the levels by an additional $\Delta E = g_F m_F \mu_B B$, where $g_1 = -\frac{1}{2}$ and $g_2 = +\frac{1}{2}$ are the Lande g-factors, $\mu_B \approx 1.4$ MHz/G is the Bohr magneton, and B is magnetic field magnitude. For $B = 13$ G, neighboring m_F levels are split by ~ 9.1 MHz. Two states are chosen, ${}^2S_{1/2} |F = 2, m_F = -2\rangle$ and ${}^2S_{1/2} |F = 1, m_F = -1\rangle$, and comprise the two level qubit system. These states are denoted $|\downarrow\rangle$ and $|\uparrow\rangle$, respectively. The use of arrows is to make an analogy to a spin- $\frac{1}{2}$ system. The total energy splitting between $|\downarrow\rangle$ and $|\uparrow\rangle$ is defined as $\hbar\omega_0$ where $\omega_0/2\pi \approx 1.28$ GHz for the experiments in this work.

A strength of hyperfine qubits is their long decoherence times. The relaxation time (T_1 time), over which the qubit decays from the higher energy state to the lower, is on the order of 10^{15} s for ${}^9\text{Be}^+$ [Wineland 98]. Phase decoherence is a different form of decoherence and is characterized by the time T_2 . This is the time scale over which the phase ϕ of a qubit in the superposition $|\uparrow\rangle + e^{i\phi} |\downarrow\rangle$ drifts from the intended phase. For ${}^9\text{Be}^+$ this is dominated by magnetic-field fluctuations that alter the energy splitting of the qubit states. Such changes cause the Larmor frequency to drift and the qubit to acquire a decoherence phase. Through careful stabilization of the currents in the magnetic-field coils and a mu-metal shield to attenuate ambient field, T_2 of 1.2 ms was observed for the qubit configuration mentioned above. Further improvement in T_2 could be achieved by using a magnetic-field insensitive qubit configuration [Langer 05]. At a higher magnetic field value, between the linear-Zeeman regime and the Paschen-Back regime, the hyperfine levels no longer shift linearly




Figure 2.1: The internal energy level diagram for ${}^9\text{Be}^+$. Hyperfine states ${}^2S_{1/2}|F=2, m_F=-2\rangle \equiv |\downarrow\rangle$ and ${}^2S_{1/2}|F=1, m_F=-1\rangle \equiv |\uparrow\rangle$ are selected to form the two-level qubit, with splitting $\omega_0/2\pi \approx 1.28$ GHz. The ${}^2P_{1/2}$ and ${}^2P_{3/2}$ manifolds contain 8 and 15 levels, respectively. The ${}^2P_{3/2}$ manifold is utilized during Doppler cooling, state detection, and state preparation. The ${}^2P_{1/2}$ manifold is utilized during ground state cooling and state preparation.

with the field. A number of values for B exist where two hyperfine states have the same slope of the magnetic field dependence. These two states can be used to define a qubit that is magnetic-field independent to first order. Using such a point at 119 G, $T_2 > 10$ s has been observed [Langer 05]. However, for the purposes of this dissertation, such long decoherence times were not required and did not warrant the added complication of creating a 119 G field.

In the excited P state, spin-orbit coupling gives rise to fine structure that splits it into $2p^2P_{1/2}$ and $2p^2P_{3/2}$ levels with a spacing of ~ 197 GHz [Bollinger 85]. (There is no fine structure in the $^2S_{1/2}$ state.) The hyperfine splitting of the $^2P_{1/2}$ states is ~ 237 MHz [Bollinger 85], while the $^2P_{3/2}$ hyperfine splitting is less than 1 MHz [Poulsen 75]. The linewidth of all these excited states is $\gamma = 2\pi \times 19.4$ MHz and the lifetime is 8.2 ns.

2.2 State detection

The $S \leftrightarrow P$ transition can be driven via electric dipole radiation at ~ 313 nm, which is in the ultraviolet (UV) region of the spectrum. In practice, this light is provided by a laser, where the frequency and polarization determine which hyperfine levels are coupled. The $^2S_{1/2} |2, -2\rangle$ to $^2P_{3/2} |3, -3\rangle$ transition, which is coupled by σ^- polarized light is of particular note, because this constitutes a cycling transition. Population in the $^2S_{1/2} |2, -2\rangle$ is driven to the $^2P_{3/2}$ level, but due to the loss of a quantum of angular momentum, selection rules require the population only be coupled to a $m_F = -3$ state. The $^2P_{3/2} |3, -3\rangle$ state is the only P state in both fine-structure manifolds that satisfies this criterion. Similarly, once population is in this excited state, selection rules dictate it must decay to an $m_F = -2, -3$ or -4 state. However, the $^2S_{1/2} |2, -2\rangle$ state is the only such S state. Therefore, the $^2S_{1/2} |2, -2\rangle$, when illuminated with σ^- light, will cycle between only these two states while continually scattering photons. A similar cycling transition exists between the $^2S_{1/2} |2, +2\rangle$ and $^2P_{3/2} |3, +3\rangle$ states, using σ^+ light. Any other combination of states or polarizations would allow the population to transfer to other hyperfine levels, eventually becoming non-resonant with the laser light at the intended transition, as the population transfers to levels with different energy splittings.

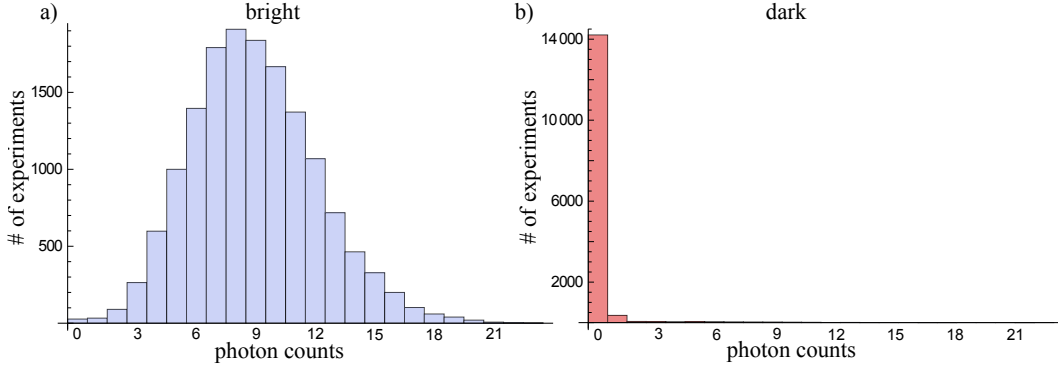


Figure 2.2: a) A sample histogram for an ion in the ‘bright’ ${}^2S_{1/2} |F = 2, m_F = -2\rangle$ state, fluorescing under BDD light. b) A histogram for an ion in the ‘dark’ ${}^2S_{1/2} |F = 1, m_F = -1\rangle$ state. Both histograms were taken using a $200 \mu\text{s}$ detection period. By comparing the histograms, the state of the ion can be determined.

The cycling transition is crucial for the detection scheme [Wineland 98], which allows discrimination between the $|\downarrow\rangle$ and $|\uparrow\rangle$ states. An ion in $|\downarrow\rangle$ (${}^2S_{1/2} |2, -2\rangle$) will scatter multiple photons when illuminated by resonant σ^- light from a laser. The steady-state scattering rate is given by [Metcalf 99]

$$\Gamma = \frac{\gamma}{2} \frac{s_0}{1 + s_0 + (2\delta/\gamma)^2}, \quad (2.1)$$

where γ is the transition linewidth and δ is the detuning of the light from resonance. The on-resonance (i.e., if $\delta = 0$) saturation parameter, $s_0 \equiv I/I_s$, is given for a specific light intensity I relative to the saturation intensity I_s where the transition saturates. This saturation intensity is given by

$$I_s \equiv \pi hc / 3\lambda^3 \tau, \quad (2.2)$$

where h is the Planck constant, λ is the wavelength, and τ is the time constant of the transition. The scattered photons can be collected and counted by a camera or a photomultiplier tube. For an ion in the $|\uparrow\rangle$ state, the same laser light will be $\omega_0/2\pi \approx 1.28$ GHz off resonance, significantly suppressing the fluorescence. Sample histograms of photon counts, for these two situations, using typical experimental parameters are shown in Fig. 2.2. By examining the overlap of these histograms, state discrimination approaching 100% can be achieved. The laser beam used for detection will be referred to as the Blue Doppler (BD) laser and derives its name from the second use for this beam:

Doppler cooling, which will be covered shortly. To accomplish the Doppler cooling, this beam is detuned 10 MHz to the red of the detection transition, still within the 19.4 MHz linewidth of the transition.

If a photon is scattered from the ${}^2S_{1/2} |1, -1\rangle$ state, it excites the atom to a $m_F = -2$ state, from which it could decay back to either ${}^2S_{1/2} |1, -1\rangle$, ${}^2S_{1/2} |2, -1\rangle$, or ${}^2S_{1/2} |2, -2\rangle$. In the latter case, the atom enters the cycling transition and begins scattering photons readily. When this occurs, it results in an ion initially in $|\uparrow\rangle$ switching to $|\downarrow\rangle$ and can lead to misidentification of the initial state. The typical detection period lasts for 200 μs , which was chosen to maximize the detection signal while minimizing the likelihood of this off-resonant scattering event. Analysis of this detection problem, along with timing-based maximum likelihood algorithms for combating it, are covered in detail in [Langer 06].

2.3 State initialization

One of the DiVincenzo criteria is the ability to initialize the qubit into a known starting state. In the present context, that involves optically pumping all population into a single energy level, the ${}^2S_{1/2} |F = 2, m_F = -2\rangle$ or $|\downarrow\rangle$ state. The short lifetime of the excited states implies that, in the absence of light, state population is distributed only among the ${}^2S_{1/2}$ ground states. Optical pumping among these states is achieved with laser light.

The BD beam weakly couples the $m_F \neq -2$ states and will remove angular momentum when it scatters. However, the BD beam is operated at a low intensity ($s_0 \approx 0.3$) below saturation. With such low intensity, and being farther off resonant for transitions starting from $m_F \neq -2$ states (especially those in the $F = 1$ manifold), this beam is inefficient at driving these transitions and performing state initialization. Errors in the σ^- polarization in the BD beam can also lead to pumping out of $|\downarrow\rangle$ faster than pumping into $|\downarrow\rangle$.

Therefore three other beams are used for initialization: two Red Doppler (RD1 and RD2) resonant beams and a Blue Doppler Detuned (BDD) beam (Fig. 2.3). All three beams have σ^- polarization, thus any scattering events will either maintain or reduce the orbital angular momentum

of the ion, due to selection rules.

The RD1 and RD2 beams are tuned to resonance with the ${}^2S_{1/2} |1, -1\rangle \leftrightarrow P_{1/2} |2, -2\rangle$ and the ${}^2S_{1/2} |2, -1\rangle \leftrightarrow P_{1/2} |2, -2\rangle$ transitions, respectively. Since these beams excite into an $m_F = -2$, the ion can return to one of three states via spontaneous emission: ${}^2S_{1/2} |1, -1\rangle$, ${}^2S_{1/2} |2, -1\rangle$, or ${}^2S_{1/2} |2, -2\rangle$. Only this last state is desired, but if the either of the first two states is reached, the ion will remain resonant with RD1 or RD2. It may take several scattering events for the ion to reach ${}^2S_{1/2} |2, -2\rangle$. However, by design, once there, neither beam will further interact with the state, since there is no $m_F = -3$ in the ${}^2P_{1/2}$ manifold to excite to. This is necessary so that the repumping process does not re-heat the ion's motion to the Doppler limit, if the ion has already been cooled to the ground state of motion, as will be discussed in Sec. 2.5.3. The RD beams are also operated at low intensity, making them inefficient at pumping the higher m_F levels.

The fourth beam, BDD, helps us with that. This beam is 602 MHz detuned to the red from BD. This large detuning helps it Doppler cool ions that have collided with background-gas molecules and become too hot to be efficiently cooled by BD. However, it can also be used for repumping. This beam typically has power 1.25 mW, which is well above saturation. This leads to a linewidth that is broad enough to address any level in the ${}^2S_{1/2}$ manifold and pump population to lower m_F states. In practice, a pulse of this beam is used initially to reduce m_F to point where the RD beams become effective (and also to cool ions if they become hot). Then RD pulses are used to complete the initialization.

In practice, the fidelity of the state initialization is limited by the purity of the σ^- polarization of the RD2 beam (the RD1 beam is too far off-resonant from the desired state). Any σ^+ or π light will cause the RD2 to continue acting on the ion once it reaches ${}^2S_{1/2} |2, -2\rangle$, pumping it back out of the intended state. This can transfer the ion to either the ${}^2S_{1/2} |2, -1\rangle$ or ${}^2S_{1/2} |1, -1\rangle$ states. In the latter case, the RD1 beam should pump back to the desired state. So the question is what is the ratio of population in the undesired ${}^2S_{1/2} |2, -1\rangle$ state versus ${}^2S_{1/2} |2, -2\rangle$ if some incorrect polarization is present?¹ The ratio of the optical pumping rates for all possible transitions is

¹ Polarization problems in the BD beam will cause a similar error in the state produced using the BD beam.




Figure 2.3: The three resonant lasers used in these experiments. Blue Doppler (BD) connects the ${}^2S_{1/2} |F=2, m_F=-2\rangle = |\downarrow\rangle$ to ${}^2P_{3/2} |F=3, m_F=-3\rangle$ (a cycling transition) for Doppler cooling and detection. Red Doppler 1 (RD1) connects ${}^2S_{1/2} |F=1, m_F=-1\rangle = |\uparrow\rangle$ to ${}^2P_{1/2} |F=2, m_F=-2\rangle$ for repumping and state initialization. Red Doppler 2 (RD2) connects ${}^2S_{1/2} |F=2, m_F=-1\rangle$ to ${}^2P_{1/2} |F=2, m_F=-2\rangle$ for repumping. All beams are σ^- polarized.

needed to answer the question. Two transitions dominate the pumping from the desired state, one involving σ^+ light via ${}^2P_{1/2}|2, -1\rangle$ and the other involving π light via ${}^2P_{1/2}|2, -2\rangle$. Other transitions involving the ${}^2P_{1/2}|1, -1\rangle$ and the ${}^2P_{3/2}$ are also possible, but much suppressed by the P state hyperfine and fine structure splitting, respectively. The optical pumping rate between any two states via a single excited state is given by [Langer 06]

$$\Gamma_{i \rightarrow f} = \frac{\gamma}{2} \frac{s_0}{1 + s_0 + (2\delta/\gamma)^2} c_{i \rightarrow f}, \quad (2.3)$$

where $c_{i \rightarrow f}$ is a coupling rate between levels. In our case, $c_{i \rightarrow f} = \frac{2}{9}$ for both undesired pumping routes, and also for the desired ${}^2S_{1/2}|2, -1\rangle \rightarrow {}^2S_{1/2}|2, -2\rangle$ route. Thus, the error in the intended state will be the sum ratios of the undesired pumping process to the desired process. We can define on-resonance saturation parameters for the four processes: s_0 for the desired polarization in RD2 along with s_{σ^+} , s_{π} , and s_{other} for the undesired processes. If the polarization errors are small, these final three parameters will be small relative to $(2\delta_i/\gamma)^2$, while $s_0 \gg (2\delta_0/\gamma)^2$. The total error can be written as

$$\mathcal{E}_{\text{init}} \approx \frac{\epsilon_{\sigma^+}(1 + s_0)}{1 + (2\delta_{\sigma^+}/\gamma)^2} + \frac{\epsilon_{\pi}(1 + s_0)}{1 + (2\delta_{\pi}/\gamma)^2} + \frac{\epsilon_{\text{other}}(1 + s_0)}{1 + (2\delta_{\text{other}}/\gamma)^2}, \quad (2.4)$$

where $\epsilon_{\sigma^+} = s_{\sigma^+}/s_0$, $\epsilon_{\pi} = s_{\pi}/s_0$, and $\epsilon_{\text{other}} = s_{\text{other}}/s_0$. The detuning for the different processes are $\delta_{\sigma^+} = 2\pi \times 12.3$ MHz, $\delta_{\pi} = 2\pi \times 9.1$ MHz, and $\delta_{\text{other}} > 2\pi \times 255$ MHz. If we assume $\epsilon_{\sigma^+} = 1\%$, this leads to an error of 0.8% for $s_0 = 0.5$. The ϵ_{π} is expected to be much less because the laser beam is basically parallel to the B -field. The contribution from all other routes is less than an order of magnitude below the dominant s_{σ^+} route. Notice in Eq. (2.4), that as the s_0 increases, the error becomes worse. For increasing intensities, the desired transition rates do not increase as quickly as the undesired transitions because the desired transitions are closer to saturation. All the resonant beams (BD, RD1, RD2) are operated below saturation to combat such errors.

2.4 Doppler cooling

For our purposes, it is important that the ion be well localized in space with minimal motion. The ion trap, itself, satisfies much of this requirement by providing a well localized harmonic

trapping potential. However, it is still necessary to remove excess ion motion in this trapping potential using laser cooling to the ground state of motion. The first step in this process is Doppler cooling.

The basic idea of Doppler cooling is to utilize the momentum of light to slow and cool an atom. An atom illuminated by a laser beam tuned near a resonant transition will absorb photons from the beam and then re-emit them in semi-random direction. The absorbed photons have a definite momentum, while the emitted photons will have an average momentum of zero due to their random direction. Therefore, each scattering event results in an average force along the beam, pushing the atom away from the light source. In Doppler cooling, the laser is tuned to the red of resonance. In this case, the light will appear blue-shifted in the frame of reference for an atom moving against the laser beam. If this Doppler shift is close to the laser detuning, the light will approximately be resonant. Thus, the detuning picks a velocity class of atoms with which it preferentially interacts. And, since the force opposes the motion into the beam, this slows the atoms in that velocity class, while having relatively less influence on atoms with other velocities. If the detuning were blue of resonance, the force would assist the motion and the ions would speed up. Placing an atom in two counter-propagating beams (with intensity below saturation), each tuned red of resonance, results in a force opposing motion in either direction. Three sets of opposing lasers along all three directions provides a force opposing all translational motion and is known as an optical molasses [Metcalf 99].

Simply reducing the atom's average velocity is not sufficient for cooling. There also needs to be a reduction in phase space. This is also achieved with the optical molasses because the Doppler force is velocity dependent and can be arranged so that it interacts less with atoms that have near-zero momentum. The ultimate limit on the cooled temperature of the atom, the Doppler limit, is determined by the cooling rate of the lasers and also the heating rate due to atom recoil when the photon is absorbed and emitted.

The first demonstrations of Doppler cooling were performed using trapped ions [Wineland 78, Neuhauser 78]. Only a single laser is necessary for Doppler cooling trapped ions. Though we have

not yet discussed the basics of the ion trap, at this point it is necessary to pull one result from our future discussion, namely that an ion confined in an ion trap behaves as though it was trapped in a three dimensional quadratic potential. This gives rise to separable harmonic motion along three perpendicular directions (\hat{x} , \hat{y} , and \hat{z}), each with its own motional frequency, ω_i . The trap acts to reverse the ion's velocity \mathbf{v} every half-cycle of the harmonic motion, so that a single beam can cool motion along any axis. This eliminates the need for two beams in every direction. Energy can be removed from all three modes of motion as long as the beam has a component of its \mathbf{k} -vector along all three axes. The cooling force on a given mode is proportional to the component of \mathbf{k} along that mode direction, while the recoil heating mechanism on re-emission proportional to the overlap of the ion's radiation pattern with the given mode. As long as the cooling is larger than the heating, the ion will cool.

For an isotropic angular distribution of emitted photons, minimum total motional energy is achieved when the \mathbf{k} is equally distributed along all three axes ($\hat{x} + \hat{y} + \hat{z}$). It is necessary that none of the harmonic modes are degenerate. In the case of degenerate modes (e.g., $\omega_x = \omega_y$), one can always make a change of basis directions (e.g., $\hat{x}' = \hat{x} + \hat{y}$, $\hat{y}' = \hat{x} - \hat{y}$) so that the harmonic motion is separated into motion perpendicular to and not-perpendicular to \mathbf{k} . The perpendicular motion will still incur the recoil heating mechanism, but will not have the cooling effect, leading to uncontrolled heating in that direction. Thus it is important for Doppler cooling that we are cognizant of the principal axes of the harmonic motion relative to the laser beam axis.

In the experiment, we operate in the Doppler limit where $(\mathbf{k} \cdot \mathbf{v})_{\text{rms}} \ll \gamma/2$. We also operate in the 'weak binding limit' where $\omega_i \ll \gamma$. In this case, the total energy of the atom along a single direction is [Itano 82]

$$E_i = \frac{\hbar}{2} \frac{\left(\frac{\gamma}{2}\right)^2 + (\omega_0 - \omega)^2}{(\omega_0 - \omega)}, \quad (2.5)$$

where $(\omega_0 - \omega)$ is the detuning of the Doppler beam from a resonant transition and the energy is minimized when $(\omega_0 - \omega) = \gamma/2$, resulting in $E_i = \hbar\gamma/2$. Since the motion is harmonic, the mean harmonic quantum number for each direction is $\bar{n}_i \simeq \gamma/2\omega_i$. For our trap parameters, $\bar{n}_i \gg 1$.

Doppler cooling produces a thermal distribution of harmonic oscillator states [Stenholm 86] where the equilibrium motional state is given by

$$|\mathcal{N}_i\rangle = \sum_{n=0}^{\infty} \frac{\bar{n}_i^n}{(1 + \bar{n}_i)^{n+1}} |n\rangle. \quad (2.6)$$

A temperature can be defined for this state by using $\bar{n}_i = 1/(e^{\hbar\omega_i/k_B T} - 1)$.

2.5 Single-ion coherent operations

In addition to the operations above, coherent operations on single qubits are needed for quantum computing. These include single qubit rotations that affect the state of a single qubit, as well as operations that coherently interact with the motional state of the ions. This interaction with the motional state is used both for ground-state cooling and coherent motional state manipulations.

For ${}^9\text{Be}^+$, the $|\uparrow\rangle$ and $|\downarrow\rangle$ states are differentiated by a magnetic-spin and can be directly coupled using a magnetic-dipole interaction involving a single photon at ω_0 . In theory, a similar interaction can also be used to couple to the motional states. However, in practice, two-photon interactions are used. The dynamics induced by the two-photon operations are qualitatively the same as those of the single-photon operations, so it is instructive to start with a description of the single-photon case for simplicity. The two-photon process will then be described.

2.5.1 Single-photon operations

The Hamiltonian for the spin- $\frac{1}{2}$ qubit system, including a three-dimensional harmonic confining potential and an applied magnetic field \mathbf{B} that interacts with the qubit via a dipole interaction, is

$$\hat{H} = \hat{H}_{\uparrow\downarrow} + \hat{H}_M + \hat{H}_{\text{int}}. \quad (2.7)$$

Here the internal atomic-state Hamiltonian is given by $\hat{H}_{\uparrow\downarrow} = \omega_0 \hat{S}_3$, where \hat{S}_3 is the analog to a spin operator that acts upon the atom's two-level internal state, and ω_0 is the splitting of the two levels. The motional energy is given by the harmonic Hamiltonian $\hat{H}_M = \hbar\omega_x \hat{a}_x^\dagger \hat{a}_x + \hbar\omega_y \hat{a}_y^\dagger \hat{a}_y + \hbar\omega_z \hat{a}_z^\dagger \hat{a}_z$, where \hat{a}_i and \hat{a}_i^\dagger are the annihilation and creation operators in a given direction and ω_i is the

frequency of the harmonic motion along the same direction. The harmonic zero-point energy is subtracted off. It is important to note that the coordinates in which $\hat{H}_{\uparrow\downarrow}$ acts ($\hat{e}_1, \hat{e}_2, \hat{e}_3$) are those of an internal atomic coordinate system (typically defined by an applied static magnetic field) while the coordinates in which \hat{H}_M acts ($\hat{x}, \hat{y}, \hat{z}$) are relative to an external direction (typically defined by the trap mode axes). These two sets of coordinate systems are typically not aligned to each other.

The magnetic-dipole interaction is given by

$$\hat{H}_{\text{int}} = -\boldsymbol{\mu} \cdot \mathbf{B}(\mathbf{r}, t), \quad (2.8)$$

where $\boldsymbol{\mu}$ is the magnetic dipole moment of the atom and $\mathbf{B}(\mathbf{r}, t)$ is the magnetic field at the position of the atom, $\mathbf{r} = (x, y, z)$.

Assume $\mathbf{B}(\mathbf{r}, t)$ is a plane wave of the form $\mathbf{B}(\mathbf{r}, t) = B_0 \hat{e}_1 \cos(\mathbf{k} \cdot \mathbf{r} - \omega_L t + \phi)$, where \mathbf{k} is the wave vector, ω_L is the frequency of the wave, and the field is linearly polarized in a direction \hat{e}_1 . We can rewrite Eq. (2.8) as

$$\hat{H}_{\text{int}} = \hbar\Omega(\hat{S}_+ + \hat{S}_-) \left(e^{i(k_z z - \omega_L t + \phi)} + e^{-i(k_z z - \omega_L t + \phi)} \right), \quad (2.9)$$

where $\Omega \equiv -\mu B_0/4$ is the Rabi frequency, and $\hat{S}_{\pm} = \hat{S}_1 \pm i\hat{S}_2$ are the ladder operators that act upon $|\downarrow\rangle$ and $|\uparrow\rangle$. Only a single motional direction \hat{z} will be considered for simplicity. However, as will see shortly, only a single motional direction can contribute to the state dynamics for a given ω_L , so this simplification will not affect the qualitative results.

It is customary to switch to the interaction picture, where $\hat{H}_{\text{int}} \rightarrow \hat{H}'_{\text{int}} \equiv \hat{U}^\dagger(t) \hat{H}_{\text{int}} \hat{U}(t)$ and $\hat{U}(t) \equiv e^{-i(\hat{H}_{\uparrow\downarrow} + \hat{H}_M)t/\hbar}$ is time evolution operator. It is also helpful to substitute $z = z_0(\hat{a}^\dagger + \hat{a})$, where $z_0 \equiv \sqrt{\frac{\hbar}{2m\omega_z}}$ is the spatial extent of the ground state wave function of the harmonic oscillator. The Lamb-Dicke parameter, $\eta \equiv k_z z_0$, which is a measure of the size of the atom's wave function to the wave length of light, can also be introduced. Using these substitutions leads to [Wineland 79, Itano 82, Wineland 98]

$$\hat{H}'_{\text{int}} = \hbar\Omega \hat{S}_+ e^{-i[\delta t - \phi]} \sum_{j,l=0}^{\infty} e^{i(j-l)\omega_z t} e^{-\eta^2/2} \frac{(i\eta \hat{a}^\dagger)^j}{j!} \frac{(i\eta \hat{a})^l}{l!} + h.c. \quad (2.10)$$

where $h.c.$ is the Hermitian conjugate and $\delta \equiv \omega_L - \omega_0$ is the detuning of the field from the qubit frequency. The rotating wave approximation was used, assuming $\delta \approx 0$, to ignore all terms that oscillate rapidly compared to δ .

Define Δn to be an integer and choose the magnetic-field detuning to be $\delta = (\Delta n + \epsilon)\omega_z$ and $\epsilon \ll 1$. In this case, a second rotating wave approximation can be performed that will pick out only the terms in the sum where $\Delta n = j - l$. The near-resonant detuning couples the motional states, $|n\rangle \leftrightarrow |n + \Delta n\rangle$ and allows for a coupling of the internal atomic state to the motional state of the atom.

Assume that the ion starts in a single atomic and motional state: $|\downarrow, n\rangle$. The interaction \hat{H}'_{int} will couple this single state to $|\uparrow, n'\rangle$, where $n' \equiv n + \Delta n$. The state of the ion under the influence of this coupling is $|\Psi(t)\rangle = C_{\downarrow, n}(t)|\downarrow, n\rangle + C_{\uparrow, n'}(t)|\uparrow, n'\rangle$ and the evolution of the coefficients can be found using the Schrödinger equation, $i\hbar \frac{\partial}{\partial t} |\Psi(t)\rangle = \hat{H}'_{\text{int}} |\Psi(t)\rangle$. There will be coupling to other motional states, but because the field detuning was chosen to be resonant with one transition, all other couplings will be small and can be ignored to a good approximation. The state, written in matrix form is

$$|\Psi(t)\rangle = \begin{bmatrix} C_{\uparrow, n'}(t) \\ C_{\downarrow, n}(t) \end{bmatrix}. \quad (2.11)$$

Solving the differential equations that result from the Schrödinger equation gives [Wineland 79, Wineland 98]:

$$|\Psi(t)\rangle = \begin{bmatrix} e^{-i\epsilon t/2} \left[\cos\left(\frac{X_{n, n'} t}{2}\right) + \frac{i\epsilon}{X_{n, n'}} \sin\left(\frac{X_{n, n'} t}{2}\right) \right] & -2ie^{-i\left(\frac{\epsilon}{2}t - \frac{\pi}{2}|\Delta n| - \phi\right)} \frac{\Omega_{n, n'}}{X_{n, n'}} \sin\left(\frac{X_{n, n'} t}{2}\right) \\ -2ie^{i\left(\frac{\epsilon}{2}t - \frac{\pi}{2}|\Delta n| - \phi\right)} \frac{\Omega_{n, n'}}{X_{n, n'}} \sin\left(\frac{X_{n, n'} t}{2}\right) & e^{i\epsilon t/2} \left[\cos\left(\frac{X_{n, n'} t}{2}\right) - \frac{i\epsilon}{X_{n, n'}} \sin\left(\frac{X_{n, n'} t}{2}\right) \right] \end{bmatrix} \times |\Psi(0)\rangle \quad (2.12)$$

where $X_{n, n'} \equiv \sqrt{4\Omega_{n, n'}^2 + \epsilon}$ and $\Omega_{n, n'}$ is an adjusted Rabi frequency given by [Wineland 79, Wineland 98]:

$$\Omega_{n, n'} \equiv \Omega \left| \langle n | e^{i\eta(\hat{a}^\dagger + \hat{a})} | n' \rangle \right|. \quad (2.13)$$

This can be expressed in a different form as

$$\Omega_{n,n'} = \Omega e^{-\frac{\eta^2}{2}} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta^{|n-n'|} L_{n_{<}}^{|n-n'|}(\eta^2), \quad (2.14)$$

where $n_{<}(n_{>})$ is the lesser (greater) of n and n' , and $L_n^\alpha(X)$ is the generalized Laguerre polynomial,

$$L_n^\alpha(X) = \sum_{m=0}^n (-1)^m \binom{n+\alpha}{n-m} \frac{X^m}{m!}. \quad (2.15)$$

In the case of on-resonant radiation ($\epsilon = 0$), Eq. (2.12) reduces to

$$|\Psi(t)\rangle = \begin{bmatrix} \cos(\Omega_{n,n'}t) & -ie^{i(\frac{\pi}{2}|\Delta n|+\phi)} \sin(\Omega_{n,n'}t) \\ -ie^{-i(\frac{\pi}{2}|\Delta n|+\phi)} \sin(\Omega_{n,n'}t) & \cos(\Omega_{n,n'}t) \end{bmatrix} |\Psi(0)\rangle. \quad (2.16)$$

This leads to familiar Rabi flopping behavior with oscillations between two states (Fig. 2.4 shows Rabi flopping driven by the two-photon Raman process discussed later), $|\downarrow, n\rangle \leftrightarrow |\uparrow, n'\rangle$ once the field is switched on, where n' is chosen by the frequency of the field. If $\omega_L = \omega_0$, then the motional state is not changed and this is called a ‘carrier’ transition. Another transition, called the ‘red-sideband’ transition, occurs if $\omega_L = \omega_0 - \omega_z$, which gives rise to the coupling of $|\downarrow, n\rangle \leftrightarrow |\uparrow, n-1\rangle$. Similarly, the ‘blue-sideband’ transition occurs at $\omega_L = \omega_0 + \omega_z$ and $|\downarrow, n\rangle \leftrightarrow |\uparrow, n+1\rangle$. The red-sideband gives rise to a transition of the form $|\text{ground}, n\rangle \leftrightarrow |\text{excited}, n-1\rangle$ which is formally equivalent to the Jaynes-Cummings hamiltonian used in quantum optics [Haroche 06]. The blue-sideband is sometimes called the anti-Jaynes-Cummings hamiltonian. Higher-order sidebands can also be driven to couple n states with greater Δn . However, the carrier and first red- and blue-sideband transitions are the only three transitions used in this dissertation.

Now consider three-dimensional harmonic motion with three distinct frequencies: ω_x , ω_y , and ω_z . For a transition to occur, $\omega_L - \omega_0 = \Delta n \omega_z$. As long as the other two motional frequencies are not degenerate with $\Delta n \omega_z$, they will not be excited at the same time. Typically this is the case and the motional modes can all be addressed individually.

If the magnetic field is pulsed on for $t = \pi/\Omega_{n,n'}$, the internal qubit state will flip, which corresponds to a π -rotation about the vertical axis of the Bloch sphere. Appropriately, this is




Figure 2.4: A ‘Rabi flopping’ trace showing $|\downarrow\rangle \leftrightarrow |\uparrow\rangle$ evolution of the internal atomic state as ‘carrier’ ($\epsilon = 0$) light is applied. This flopping used the coherent two-photon Raman process discussed in Sec. 2.5.2 as opposed to single microwave photons. After the ‘carrier’ light is turned off, a $200 \mu\text{s}$ BD detection pulse is applied and the number of scattered photons is counted. An ion in the $|\downarrow\rangle$ state during detection scatters photons while the $|\uparrow\rangle$ state is dark. Each point represents an average of 200 consecutive experiments. (a) and (b) show flopping over different time scales. The noise and loss of contrast at longer times is due to technical issues and is primarily caused by magnetic field fluctuations and laser intensity noise, both of which affect the instantaneous Rabi frequency.

called a π -pulse. Similarly, a laser pulse for $t = \pi/2\Omega_{n,n'}$ would rotate a qubit from the poles of the Bloch sphere to an equal superposition $|\downarrow\rangle \rightarrow \frac{1}{\sqrt{2}}(|\downarrow\rangle + |\uparrow\rangle)$ and is called a $\frac{\pi}{2}$ -pulse. Quantum computing requires this capability to rotate a single qubit. The carrier transition is typically used for changing the qubit so that the internal and motional states are not entangled in the process. (The motional states have a shorter coherence time, so it is advantageous to keep the motional and qubit states separate.) The blue and red-sidebands are more typically used in ground state cooling and measurement of temperature, as will be discussed later, as well as in the original two-ion gate proposal in [Cirac 95].

Examining the coupling between states yields a couple of insights. If the intent is to drive carrier transitions as efficiently as possible (i.e., maximize $\Omega_{n,n'}$), this occurs when the expectation value in Eq. (2.13) is maximized. This condition is met when $\eta \rightarrow 0$ and the wave length of the field is much greater than the ion's wave packet. In other words, when the magnetic field has no spatial dependence. However, if the intent is to drive sideband transitions (of any order), small η leads to a small \hat{H}'_{int} as seen in Eq. (2.10), resulting in small coupling. A larger η is needed to maximize the sideband interaction. To achieve coupling to the motion, a spatially dependent force on the ion is required. Such a force occurs when the magnetic field has a spatial gradient. This implies either a large wave packet or a large \mathbf{k} . Therefore, the carrier and sideband transitions place different requirements on the field.

To couple directly to the ${}^9\text{Be}^+$ qubit with a single photon, the frequency of the photon must be 1.28 GHz, which is a microwave field with small \mathbf{k} . This makes these photons ideal for carrier transitions. However, the large wavelength of microwaves (~ 23 cm) makes it difficult to selectively address a single ion if more than one ion is present in the trap ($\ll 1$ cm). Another problem is that the small \mathbf{k} makes the coupling to the motional state extremely weak. Nonetheless, a number of interesting experiments have been performed using microwaves. These include an ion-based atomic clock [Bollinger 91] and the $g - 2$ experiments of Dehmelt [Dehmelt 90].

The problems with microwaves can be circumvented if magnetic-field inducing structures in the vicinity of the ion can be made small enough to achieve a high effective Lamb-Dicke

parameter using the geometric structure (instead of the wavelength) to create a large magnetic-field gradient [Wineland 98]. Some proposals have been made to do this with small current carrying wires in a surface-electrode trap [Chiaverini 08, Ospelkaus 08]. Similar results could be achieved by transporting ions above static magnetic domains [Leibfried 07].

However, the approach followed in this dissertation is to replace single photon interactions with a coherent two-photon process with each photon having much shorter wavelength (313 nm), and thus addressability (because lasers can be used) and large η .

2.5.2 Two-photon stimulated-Raman operations

The single-photon operations are often used in ion species that have large qubit splittings that correspond to optical transitions accessible via electric-dipole transitions, such as $^{40}\text{Ca}^+$ [Nägerl 99]. This allows the use of lasers to address the problems above. However, to use lasers on $\omega_0/2\pi \approx 1.28$ GHz qubit transition in $^9\text{Be}^+$, we must use two-photon stimulated Raman transitions. The basic idea is instead of directly coupling population between $|\downarrow\rangle$ and $|\uparrow\rangle$, these levels are connected via a third, higher energy level. This is depicted in Fig. 2.5. The population in the two lower qubit levels can access the higher level via an electric-dipole interaction using ~ 313 nm light. However, the qubit transition needs two photons, one that connects $|\downarrow\rangle$ to the higher state and one that connects $|\uparrow\rangle$ to the higher state, where the difference frequency between the beams is $\omega_0/2\pi \approx 1.28$ GHz.

The theory of this three-level, two-photon interaction is covered in [Wineland 98, Wineland 03] and will not be covered here. However, the conclusion is that this more complicated three level system can be simplified to a two-level analog when the two lasers are applied simultaneously, using an adiabatic elimination argument. The qualitative behavior is identical that of the single-photon system described above, with three replacements: $\mathbf{k} \rightarrow \mathbf{k}_1 - \mathbf{k}_2$, $\omega_L \rightarrow \omega_1 - \omega_2$, and $\phi \rightarrow \phi_1 - \phi_2$. Here, the indices denote which laser the quantity corresponds to.

It is advantageous to use a virtual level rather than an actual atomic level (such as the P levels) as the third level. This is done at the expense of decreased Rabi frequency for the transition.




Figure 2.5: The qubit energy levels and Raman beam frequencies. The $|\downarrow\rangle$ and $|\uparrow\rangle$ qubit states are shown with a ladder of equally-spaced harmonic-oscillator states $|n\rangle$ superimposed on top of them. The two Raman beams can connect $|\downarrow\rangle$ and $|\uparrow\rangle$ via a third level near the P levels. This third level is actually a virtual level detuned by Δ_P from the ${}^2P_{1/2}$ states and $\Delta_{P'}$ from the ${}^2P_{3/2}$ states. Each Raman beam has its own wave vector and frequency. By adjusting $\omega_1 - \omega_2$ appropriately, the $|n\rangle$ state can be changed simultaneously with the qubit state. The figure depicts a frequency configuration that would lead to a blue-sideband transition.

However, the virtual level suppresses spontaneous emission from the short lived excited state, which would lead to a reduction in the fidelity of the operation. The errors decrease as the detuning from the closest atomic level, Δ , increases. The Raman beam intensities must be increased by Δ to maintain the same Rabi frequency. A detailed discussion of the fidelity of such operations as function of the detunings can be found in [Ozeri 07]. In practice, the lasers are tuned so that the virtual level falls between the ${}^2P_{1/2}$ and ${}^2P_{3/2}$ levels, detuned by $\Delta_P = 80$ GHz from the ${}^2P_{1/2}$ level. For ~ 10 mW of Raman power at this detuning, focused to a waist $w_0 \approx 30 \mu\text{m}$, we achieve a π time of $\sim 2 \mu\text{s}$ for a carrier transition.

As in the single-photon case, carrier transitions can be driven using light frequencies of $\omega_1 - \omega_2 = \omega_0$. Sidebands can be driven if $\omega_1 - \omega_2 = \omega_0 + \Delta n \omega_z$.

The polarizations of these beams are chosen so that one is π light while the other is σ^+/σ^- . One of the circular polarizations provides the change in angular momentum between $|\uparrow\rangle$ and $|\downarrow\rangle$, while the other circular polarization is included to balance Stark shifts [Wineland 03].




Figure 2.6: The geometric arrangement for the two Raman beams. a) The copropagating arrangement used for co-carrier transitions. The two Raman beams follow the same path onto the ion, represented by the ball. The resulting $\Delta\mathbf{k}$ is very small. In practice, these beams are arranged at 45° to the primary ion motion along the trap's \hat{z} axis, though this is not necessary. b) The 90° arrangement used for sideband transitions. The beams are arranged perpendicular each other in such a way that the resulting $\Delta\mathbf{k}$ is large and points along the ion's z motion, while not pointing along the other axes of motion.

As stated above, the requirements on η for good coupling to the carrier transitions are different than those for good coupling to sideband transitions. A benefit of the Raman interaction is that $\mathbf{k} = \mathbf{k}_1 - \mathbf{k}_2$, and thus the effective η , can be adjusted using the geometry of the beams. If the two beams are copropagating (Fig. 2.6(a)), $\mathbf{k}_1 - \mathbf{k}_2 \approx 0.08/\text{m}$ which is small and gives the same η that would have existed for the single-microwave-photon drive. In this case the carrier Rabi flopping rate is maximized. Such transitions are called copropagating carrier (or ‘co-carrier’) transitions.

On the other hand, to maximize the coupling to the motion and thus the sideband transitions, the two beams can be aligned counterpropagating to maximize $|\mathbf{k}_1 - \mathbf{k}_2| \approx 4 \times 10^7/\text{m}$. This gives an effective \mathbf{k} and η that corresponds to a UV wavelength instead of a microwave. For a 3.6 MHz motional frequency, $\eta = 0.48$. Though counterpropagating beams gives the largest η , aligning the beams at 90° to each other, as in Fig. 2.6(b), yields $|\mathbf{k}_1 - \mathbf{k}_2| \approx 2.8 \times 10^7/\text{m}$ (down by a factor of $\frac{1}{\sqrt{2}}$) and $\eta = 0.34$. By aligning the 90° beams so that $\mathbf{k}_1 - \mathbf{k}_2$ points along the direction of the harmonic motion that we intend to interact with, the beams provide significant coupling to that mode of motion while giving minimal coupling to the other modes of motion.

So far, we have assumed the ion is in a single motional state $|n\rangle$. However, the ion can be in a distribution of $|n\rangle$ states. For example, an ion that has just undergone Doppler cooling will exist in a thermal distribution given by Eq. (2.6). If the light interaction is sensitive to the motion,

the atomic qubit state for each motional state will oscillate with a different Rabi frequency $\Omega_{n,n'}$, leading to a complicated time-dependent detection signal when averaged over many experiments. Since detection is only sensitive to the qubit state and not to $|n\rangle$, the observed Rabi flopping will be a convolution of all the individual flopping terms for different $|n\rangle$.

The qubit will generally become entangled with the motion during the flopping, leading to poor qubit rotations and low-fidelity single-qubit gates. This is true even for the carrier transition where $\Delta n = 0$. One option to increase the fidelity of the qubit rotations is to remove the $|n\rangle$ dependence of the coupling. By minimizing the overlap of \mathbf{k} and the motional axis \hat{z} , which removes the $|n\rangle$ state dependence in Eq. (2.9). This is exactly the case of the co-carrier beam geometry, so co-carrier interactions will not be highly sensitive to this issue. In fact, co-carrier spin flips can be done with high fidelity $> 99\%$ on ions that have not been ground-state cooled and only been Doppler cooled.

Alternatively, the atom can be prepared in a single $|n\rangle$ state, called a number state or a Fock state. The easiest such state to achieve is the ground state $|0\rangle$ through ground state laser cooling.

The co-carrier beams can only be used for carrier transitions while the 90° beams can be used for either sideband or carrier transitions. However, the 90° geometry's sensitivity to motion makes the co-carrier beams much more desirable for high-fidelity carrier operations. The optimal solution is having both geometries at our disposal.

Though the state evolution is qualitatively the same for the Raman process and a single-photon process, there are some important differences. First, the single photon process must act through a magnetic dipole process to induce changes between hyperfine states. The Raman beams interact individually via an electric dipole coupling to an excited state, but the simultaneous interactions change the hyperfine state. Second, the issue of spontaneous emission resulting in errors is a problem with the optical Raman process and could be avoided with a single-photon microwave transition. Finally, the Raman arrangement uses non-resonant fields that lead to AC Stark shifts [Wineland 03]. These shifts slightly alter ω_0 while the beams are applied and lead to a small change in the qubit precession frequency. This results in qubit phase accumulation that must be accounted

for when performing experiments that are phase-sensitive (e.g., Ramsey spectroscopy) or when using multiple ions where the relative phase is important (e.g., large scale quantum computing). This shift can be minimized by using near-linear polarizations for the two beams, but is difficult to cancel completely.

2.5.3 Ground state cooling

Multiqubit gates entangle ions via an interaction between light and the harmonic motion of the ions. To achieve maximum fidelity, that motion must be initialized in a known Fock state. Doppler cooling produces a thermal state that contains too many Fock states to obtain optimal fidelity from the gate operations. Using a method similar to that in [Monroe 95], it is possible to cool the motion of the ion to the ground state $|0\rangle$, which is a convenient Fock state for gate operations.

As shown above, the 90° red-sideband transition couples the $|\downarrow, n\rangle \leftrightarrow |\uparrow, n-1\rangle$, which can be used to remove a single quantum of motion. However, continued application of this sideband light will return that quantum of motion to the ion and cannot be used, alone, to cool the ion to the ground state .

In addition to red-sideband light, pulses of RD1 and RD2 light are also needed to realize ground state cooling. The process begins by optically pumping to the $|\downarrow\rangle$ state, after which a $\frac{\pi}{2}$ -pulse is applied using the red sideband to transition from $|\downarrow, n\rangle \rightarrow |\uparrow, n-1\rangle$. Then a pulse of RD1 and RD2 light is applied. This optical-pumping light will efficiently pump $|\uparrow\rangle$ back to $|\downarrow\rangle$ using a minimal number of scattering events, such that the recoil heating is $\ll \hbar\omega_z$. Therefore the red-sideband and RD repumping pulses, in succession, change the ion's state from $|\downarrow, n\rangle \rightarrow |\downarrow, n-1\rangle$. By iterating between these two types of pulses, the motional occupation ratchets down one quantum at a time, until it reaches $|\downarrow, 0\rangle$. Once the ground state is reached, the red sideband pulse can no longer remove a quantum of motional energy and stops interacting with the ion's state. This causes the RD light to also stop interacting. Any further pulses have little effect on the ion. Some residual energy due to the RD recoil and other inefficiencies will keep the ion from reaching the true ground

state, but $\langle n \rangle \approx 0.01$ can be achieved.

Since the ion starts in a thermal state after Doppler cooling, there will be a distribution of motional states prior to ground state cooling. Achieving a perfect π -pulse on the red sideband for all these states simultaneously is not possible because the π -time (and $\Omega_{n,n-1}$) depends on $|n\rangle$. However, a perfect π -pulse is not needed. During a single pulse, any population that fails to make the intended red-sideband transition due to an incorrect pulse duration, will simply miss that single round of cooling. That population can still be cooled by the next pulse. If enough pulses are performed, all the population will eventually arrive in $|\downarrow, 0\rangle$. On the order of 20 cooling pulses are typically sufficient to go from the Doppler limit of $\langle n \rangle \approx 3$ to the ground state of a single mode of motion.

2.5.4 Shelving

One problem with the fidelity of state detection is that the ‘dark’ ${}^2S_{1/2} |F = 1, m_F = -1\rangle$ state, which ideally would not scatter any photons, still has some probability of off-resonantly scattering a photon from the BD beam during the detection period. If this causes the ion to be pumped to the ${}^2S_{1/2} |2, -2\rangle$ state, the ion can start scattering on the cycling transition for the remainder of the detection period. This results in reduced state discrimination, as the would-be dark $|\uparrow\rangle$ state masquerades as a bright $|\downarrow\rangle$ state.

However, in a process known as ‘shelving’, Raman transitions can be used to coherently transfer from the ‘dark’ ${}^2S_{1/2} |F = 1, m_F = -1\rangle$ state to an even darker state outside of the qubit manifold. The Raman-beam frequencies are adjusted so that $\omega_1 - \omega_2$ corresponds to the energy splitting between $S_{1/2} |1, -1\rangle$ and $S_{1/2} |2, 0\rangle$, and then a π -pulse is applied to transfer any population in the $|\uparrow\rangle$ state to the $S_{1/2} |2, 0\rangle$ state. A second shelving pulse can be performed to transfer that population even farther into $S_{1/2} |1, +1\rangle$. After this, when the BD beam light is applied for detection, this shelved population would have to scatter multiple photons before it could reach ${}^2S_{1/2} |2, -2\rangle$ and the cycling transition. Since all of these scattering events would be off-resonant, the likelihood of this is much reduced over the non-shelved case and the state discrimination is

increased.

Standard optical pumping with the use of either the high-intensity BDD beam or the RD beams will return the shelved population to the qubit manifold to prepare for the next experiment.

2.6 Two-ion coherent operations

The techniques for interacting with the internal and motional states of single ions can be extended to N ions. For ions located together in a single trapping well, their mutual Coulombic repulsion will keep the ions separated by several micrometers. This repulsion also acts to couple the motion of the ions. Thus each ion is subject to a combination of Coulombic repulsion and harmonic trapping potential. The motion of the ions can be described in terms of $3N$ modes of motion. In a single direction, the motion of two ions is described by a center-of-mass (COM) mode, in which both the ions' motion is in phase, and a stretch mode, where the motion is out of phase. These modes can be individually addressed to perform carrier and sideband transitions [Wineland 98], as well as to cool to the ground state [King 98]. Further discussion of this can be found in [Wineland 98, King 99].

One of the requirements of quantum computing is the ability to perform some type of two-qubit entangling operation. The joint motion of multiple ions can be used as an information bus for interacting and entangling the atomic states of those ions and realizing this requirement. The first proposal for a two-ion entangling gate is outlined in [Cirac 95]. In this proposal, the qubit information from one ion can be transferred from the atomic state to the shared motional state of both ions using a Raman sideband pulse. Then additional pulses can be used to interact the motion with the second ion's qubit state to produce the desired entanglement. This requires that the motional state be in a Fock state prior to the first pulse. If this is not the case, the fidelity of the operation would be significantly reduced. More recent entangling gate operations that are less sensitive to the exact motional state have been proposed [Mølmer 99, Milburn 00]. These operations still use the motion as a bus, but they are devised to interact with all motional states in the same way to provide the motional insensitivity. However, this is true only to the extent

to which the ion's harmonic wave packet experiences the same light force across its spatial extent [Wineland 98], which is true to the extent that η is small. For small η , the motion would have to be in a high $|n\rangle$ for the wave packet to be large enough to feel an inhomogeneous force due to the finite wavelength of the light. On the other hand, if η is small, the force (and the coupling to the motion) will also be small which will result in a slow gate operation. Therefore, there is a trade off between slow gates and harmonic-state sensitivity. For typical parameters in this dissertation, $\eta \approx 0.3$ and the gate fidelities of these more robust entangling operations will remain high for an average motional occupation of $\bar{n} < 1$. Thus, despite the increased motional insensitivity, ground state cooling is still needed.

The entangling phase-gate demonstrated in [Leibfried 03] achieved a fidelity of 97% with a gate duration of 39 μs using ${}^9\text{Be}^+$ ions under nominally identical operating parameters to those used in this dissertation. In that experiment, the fidelity was limited by laser intensity fluctuations (1% error), drifts in the detuning of the lasers relative to the ω_z (1% error), and spontaneous emission from the virtual level (2.2% error). The spontaneous emission can be suppressed by using other ion species or by detuning farther from the excited P levels while using higher laser intensity [Ozeri 07]. If this is done, and the first two sources of error are stabilized to the level of 10^{-3} , a total gate error of 10^{-4} should be achievable.

For this to be possible, it is important that the motional state of the ions be maintained very near $|0\rangle$. The typical experiment is preceded by a series of Raman cooling pulses to put all the ions into the motional ground state. The experiment itself will consist of a series of laser interactions and transport operations that perform the desired qubit operations. However, if the ions gain motional energy during these processes, the fidelity of the gate operations will be adversely affected. Once the experiment has commenced, ground state cooling can no longer be performed directly on the qubit ions because that operation destroys the atomic state information that is needed for the computation. A potentially significant cause of heating during large-scale quantum computing experiments is the transport of ions through the multiplexed trap architecture. This potential loss of gate fidelity is the reason ion transport must excite the motion of the ion as little

as possible.

2.7 $^{24}\text{Mg}^+$ ions

Most of the work described in this dissertation involves $^9\text{Be}^+$ ions. However, the initial demonstrations of transport through the X-junction array were performed with $^{24}\text{Mg}^+$. The $^{24}\text{Mg}^+$ ion does not have hyperfine structure allowing it to be manipulated with a smaller number of laser frequencies but makes it a poor candidate for a qubit. The NIST Ion Storage Group has a $^{24}\text{Mg}^+$ apparatus devoted to testing the non-quantum aspects of new traps and it was in that facility that the first transports through the X-junction were performed.

Two electronic states were used when working with $^{24}\text{Mg}^+$: the $3s^2S_{1/2}$ and the $3p^2P_{3/2}$ states. A 280 nm laser beam with σ^- polarization created a cycling transition between $^2S_{1/2} |J = 1/2, m_J = -1/2\rangle \leftrightarrow ^2P_{3/2} |J = 3/2, m_J = -3/2\rangle$ for performing both Doppler cooling and resonant detection. The linewidth of this transition is 40 MHz and the beam was tuned 20 MHz red of resonance. A second beam tuned 400 MHz to the red was also used, much like BDD, to cool the ions during loading and when they were excessively hot.

In addition to cooling and detecting, the near-resonance beam was also used to measure the excitation of the motion of the ion from various sources following the proposal in [Wesenberg 07].

With $^{24}\text{Mg}^+$, no magnetic field was applied and laser polarization determined the quantization axis for the atomic state. This was sufficient for demonstrating ion transport.

Chapter 3

Linear rf Paul ion traps

The ion trap is an excellent platform for studying a broad range of physics. Traps can confine anywhere from a single ion to a collection of millions of ions and do so for periods that range from seconds to months, depending on the system. The trapped ions often reside hundreds of micrometers or more from the nearest electrodes, allowing them to be well isolated from unwanted external forces. The influence of the trapping potentials on the internal atomic states of an ion is often well known and understood. In fact, the very nature of the trapping mechanism leads an ion to seek out a location in the trap where these influences are minimized. Due to these properties, ion traps have been used in a number of impressive experiments. Laser cooling, which has become a requirement for most modern atomic physics research, was first demonstrated on trapped ions [Wineland 78, Neuhauser 78]. Currently, the most accurate atomic clocks, with uncertainties less than 1×10^{-17} , use trapped-ions [Chou 10]. Other applications include ultra-sensitive mass spectroscopy [Van Dyck Jr 01, Rainville 04], quantum-logic aided spectroscopy [Hume 07], measurements of fundamental constants such as the electron magnetic moment [Hanneke 08], and the production of anti-hydrogen [Gabrielse 05]. The properties of trapped ions also make them excellent candidates for studying quantum information processing.

Earnshaw's theorem states that charged particle cannot be confined through the use of static electric fields alone because the field is divergenceless. This restriction can be overcome through the combination of static electric fields with static magnetic fields or oscillating electric fields. Such configurations are known as Penning traps [Ghosh 95] and rf Paul traps [Paul 90], respectively.

This dissertation will focus on variants of the rf Paul trap. In fact, one goal of this dissertation is to extend the linear rf Paul configuration to include more complicated features, such as junctions. Therefore, some effort will be expended in describing how an ideal rf Paul trap creates a three-dimensional harmonic confining potential for charged particles, including the simple adiabatic approximation and the more detailed Mathieu equation analysis. This will be followed by an analysis of the inhomogeneous Mathieu equation that arises in non-ideal Paul traps, such as surface-electrode traps and junctions. The chapter will conclude with a description of actual trap geometries and basic design considerations.

3.1 Static confining potential

Though Earnshaw's theorem implies a static electric field cannot independently confine an ion, the linear Paul trap includes a static-field component. Therefore, it is important to understand the characteristics of a static confining potential prior to considering the oscillating portion. In fact, Earnshaw's theorem is only true when all three dimensions are considered. A static field can confine an ion in one or two dimensions. We begin the discussion by restricting focus to a single direction.

Assume that an ion is placed in a 1-D quadratic electric potential of the form

$$\Phi_{\text{st}}(x) = U_0 + \frac{U_1}{R}x + \frac{U_2}{2R^2}x^2, \quad (3.1)$$

where U_i is in units of volts and R is the distance from the ion to the nearest electrode. Both a uniform offset term and a linear have also been included. However, neither term have a significant impact in the ion. The offset term, U_0 , can be removed by simply shifting the zero point of the potential. The linear term provides a shift, along x , of the minimum of the quadratic potential. This term can be incorporated into the quadratic term by making the change of variables $\tilde{x} = \left(x + \frac{U_1}{U_2}R\right)$, which will be accompanied by an offset shift in the potential that can be discarded as U_0 was. Therefore, without loss of generality, we can assume a purely quadratic potential,

$$\Phi_{\text{st}}(x) = \frac{U_2}{2R^2}x^2, \quad (3.2)$$

where $x = 0$ is chosen to coincide with the minimum of Eq. (3.1).

Terms that are higher order in x (cubic, etc.) may also present. However, it is usually possible to restrict focus to a region small enough that the higher order terms are negligible. For laser-cooled ions in a rf Paul trap, the ion is typically well-enough localized that this is the case.

The Hamiltonian for a particle with mass m and charge q in this potential is given by

$$H(x) = \frac{1}{2}m\dot{x}^2 + \frac{qU_2}{2R^2}x^2. \quad (3.3)$$

This is the Hamiltonian for a harmonic oscillator with frequency $\omega = \sqrt{\frac{qU_2}{mR^2}}$. Thus, a quadratic electric potential provides harmonic confinement in one dimension. However, this potential is only confining for $U_2 > 0$.

One might be tempted to construct a static electric potential that is quadratic in three dimensions to provide three-dimensional confinement. Take, for example, the potential

$$\Phi_{\text{st}}(\mathbf{z}) = \frac{U_x}{2R^2}x^2 + \frac{U_y}{2R^2}y^2 + \frac{U_z}{2R^2}z^2. \quad (3.4)$$

However, Laplace's theorem stipulates that $\nabla^2\Phi_{\text{st}}(\mathbf{z}) = 0$ for an electric potential in a charge-free region. The result,

$$\frac{U_x}{R^2} + \frac{U_y}{R^2} + \frac{U_z}{R^2} = 0, \quad (3.5)$$

implies that at least one of the U 's must be negative and anti-confining. This is the basis of Earnshaw's theorem. In the case of negative U_i , an oscillation frequency can still be defined as $\omega_i = \sqrt{\frac{qU_i}{mR^2}}$, though it is imaginary.

In practice, it is useful to calculate the Hessian matrix of the energy potential. The Hessian is defined as the matrix of second derivatives,

$$\mathcal{H} \equiv q \begin{bmatrix} \frac{\partial^2}{\partial x^2} & \frac{\partial^2}{\partial x \partial y} & \frac{\partial^2}{\partial x \partial z} \\ \frac{\partial^2}{\partial y \partial x} & \frac{\partial^2}{\partial y^2} & \frac{\partial^2}{\partial y \partial z} \\ \frac{\partial^2}{\partial z \partial x} & \frac{\partial^2}{\partial z \partial y} & \frac{\partial^2}{\partial z^2} \end{bmatrix} \Phi_{\text{st}}(\mathbf{z}). \quad (3.6)$$

The second derivative is often referred to as the 'curvature'. If the potential is predominantly quadratic, as intended, then the \mathcal{H} will not depend on position. The eigenvalues of the \mathcal{H} , λ_n (for

$n = 1, 2, 3$), can be used to find three harmonic frequencies

$$\omega_n = \sqrt{\frac{\lambda_n}{m}}. \quad (3.7)$$

These three motions occur along three perpendicular directions that coincide with the eigenvectors of \mathcal{H} . These three directions are referred to as the principal axes of the harmonic motion. They represent a basis along which the motion of the ion can be separated and is nominally independent of motion in the other directions. In cases where the ion moves far from the origin, higher order terms in $\Phi_{\text{st}}(\mathbf{z})$ can couple these independent motions, but for laser-cooled (relatively stationary) ions in nearly ideal quadratic traps, such coupling is seldom observed (though it was explored in [Roos 08]).

As we have seen, static electric potentials can be used to harmonically confine a particle in at least one direction, but lead to anti-confinement along at least one other direction. The inclusion of oscillating electric potentials can also provide confinement. If this rf confinement is stronger than the anti-confinement, the oscillating and static potentials can be used in conjunction to produce three-dimensional confinement.

3.2 rf Paul trap

The Paul trap is defined by the use of an oscillating quadrupole field to confine ions. It is instructive to first consider the Paul trap in two dimensions, where there are four electrodes each situated on the corner of a square (see Fig. 3.1(a)). Two of these electrodes have an oscillating radio frequency (rf) potential applied to them, while the other two are held at rf-ground (typically this is done with a capacitive connection to ground). The two grounded electrodes will be referred to as the control electrodes, for reasons that will become apparent later in the chapter. This electrode configuration results in a two-dimensional quadrupole potential. At the center of the trap is a region referred to as the rf null, as it has no field. At an instant in time, this potential will attract an ion (situated at the center) to both of the control electrodes, while repelling it from the rf electrodes. However, during the next half-cycle of the rf, the roles of the rf and control electrodes will switch,

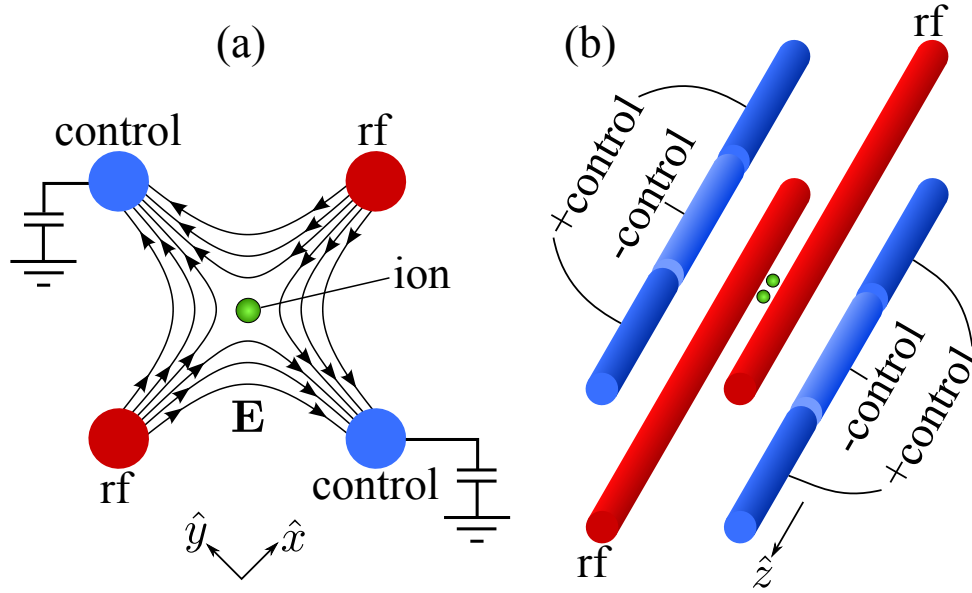


Figure 3.1: (a) An rf Paul trap in two dimensions. A common rf potential is applied to two electrodes while the other two, the control electrodes, are held at rf-ground using a capacitor. This results in a quadrupole \mathbf{E} -field, with the ion situated at the rf null in the center. The ion is confined in both directions. Note that here, \hat{x} and \hat{y} are defined as shown but will later be defined by the modes of motion of the ion. (b) In a linear rf Paul trap, this geometry is extruded along the \hat{z} axis to form four parallel rods. The control rods are each segmented into three sections, with the center segment receiving a lower static potential relative to the outer segments. This yields a static trapping potential along the \hat{z} axis, which, in conjunction with the rf-trapping, results in harmonic confinement in all directions. Two ions are depicted held in the center of the trap.

with the ion now being attracted to the rf electrodes and repelled from the control electrodes. To see how this situation might give rise to a trapping potential, it is helpful to consider the analogy of a marble on saddle [Paul 92].

Like the ion in the quadrupole potential, a marble placed at the center of a saddle will want to roll down the saddle in one direction (say either the $\pm\hat{y}$ directions) due to gravity, while being confined along $\pm\hat{x}$. However, if the saddle is quickly rotated by 90° (around \hat{z}) so that \hat{y} is now a direction moving uphill, the marble will no longer travel along \hat{y} and feel a force returning it to the center. If the saddle is spun fast enough, there is no direction the marble can move that it will not see a net uphill climb. Thus, the marble will stay right in the center of the spinning saddle. The speed above which the saddle must be rotated to achieve this depends on the mass of the marble,

the curvature of the saddle, and the force of gravity. But as long as the saddle rotates faster than the marble can make a big excursion from the center, marble will be trapped.

In the case of the Paul trap the situation is quite similar, with the electric force replacing the mechanical force in the marble's case. One small difference is that instead of using a rotating quadrupole potential, the Paul trap typically uses an inverting quadrupole potential. However, this does not significantly alter the situation. In fact, Paul traps can be made with rotating potentials [Hasegawa 05], though this introduces the technical challenge of realizing such a potential.

3.2.1 Adiabatic approximation

There are two commonly used methods for treating the confinement in an rf Paul trap. A complete treatment requires the Mathieu equation and this will be covered in a later section. In this section, we will use a static 'pseudopotential' to approximate the rf confining field, which is a helpful simplification that is reasonably valid for the operating parameters used in most quantum computing experiments.

The oscillating field must be quadrupolar to produce harmonic confinement. But to understand how an oscillating field can provide any net force on an ion, we will first examine a homogeneous oscillating electric field of the form $\mathbf{E}_{\text{rf}}(t) = -\nabla\Phi_{\text{rf}}(t) = -\nabla\Phi_0 \cos \Omega_{\text{rf}}t$. Here, Φ_0 is linear in position so that $\mathbf{E}_0 \equiv -\nabla\Phi_0$ does not have spatial dependence. A particle with charge q , mass m is located at position \mathbf{r}_μ . The force on the particle is given by

$$F(t) = m \frac{\partial^2 \mathbf{r}_\mu}{\partial t^2} = q \mathbf{E}_0 \cos \Omega_{\text{rf}}t. \quad (3.8)$$

Solving for both the position and the kinetic energy of the motion is straight forward and yields

$$\mathbf{r}_\mu(t) = -\frac{q}{m\Omega_{\text{rf}}^2} \mathbf{E}_0 \cos \Omega_{\text{rf}}t \quad \text{KE}_\mu(t) = \frac{q^2}{2m\Omega_{\text{rf}}^2} \mathbf{E}_0^2 \sin^2 \Omega_{\text{rf}}t. \quad (3.9)$$

Now assume $\mathbf{E}_{\text{rf}}(\mathbf{r}, t)$ has a small spatial dependence. At the same time, allow the ion to have a slower motion $\mathbf{r}_s(t)$, such that $|\Delta\mathbf{r}_s| \gg |\Delta\mathbf{r}_\mu|$ and the total motion is $\mathbf{r}_{\text{tot}}(t) = \mathbf{r}_s(t) + \mathbf{r}_\mu(t)$. The spatial dependence of \mathbf{E}_0 should be small enough that it is approximately constant over the displacement

\mathbf{r}_μ during a single rf period. Still, over many periods, \mathbf{r}_s can drift to a location where \mathbf{E}_0 is different. This spatial dependence will minimally affect the Eq. (3.9) dynamics over a single rf period. The Hamiltonian for this system is

$$H(t) = \frac{m}{2} \dot{\mathbf{r}}_{\text{tot}}^2(t) + q\Phi_{\text{rf}}(\mathbf{r}_{\text{tot}}(t), t). \quad (3.10)$$

As long as the $\mathbf{r}_s(t)$ motion is significantly slower and larger in amplitude than $\mathbf{r}_\mu(t)$, we can separate the two motions and perform a time average over a single period of $\mathbf{r}_\mu(t)$ to give

$$H(\mathbf{r}_s, t) = \left\langle \frac{m}{2} \dot{\mathbf{r}}_s^2 + \frac{m}{2} \dot{\mathbf{r}}_\mu^2 + m\dot{\mathbf{r}}_s \dot{\mathbf{r}}_\mu(t) + q\Phi_{\text{rf}}(\mathbf{r}_s + \mathbf{r}_\mu(t), t) \right\rangle_{\Omega_{\text{rf}}} \quad (3.11)$$

$$= \frac{m}{2} \dot{\mathbf{r}}_s^2 + \langle \text{KE}_\mu(\mathbf{r}_s, t) \rangle_{\Omega_{\text{rf}}} + 0 + 0, \quad (3.12)$$

where $\langle f(t) \rangle_{\Omega_{\text{rf}}}$ is the time average of $f(t)$ over a single period of Ω_{rf} . The third term vanishes because $\dot{\mathbf{r}}_s$ is nearly constant over a cycle and $\dot{\mathbf{r}}_\mu$ averages to zero. Similarly the last term vanishes because Φ_{rf} is oscillatory but (nearly) spatially constant over a single cycle. The second term is the \mathbf{r}_μ kinetic energy and is given in Eq. (3.9). Since there is a large discrepancy in time scales of the two motions, this second term behaves much like a potential energy for the slower \mathbf{r}_s motion. This term can be replaced by a time-independent quantity which is often referred to as the ‘pseudopotential’. This gives the simplified Hamiltonian,

$$H(\mathbf{r}_s, t) = \frac{m}{2} \dot{\mathbf{r}}_s^2 + \frac{q^2}{4m\Omega_{\text{rf}}^2} \mathbf{E}_0^2(\mathbf{r}_s). \quad (3.13)$$

It is important to note that the pseudopotential is actually a ponderomotive potential due to the kinetic energy associated with the fast, small-scale motion $\mathbf{r}_\mu(t)$. This fast motion is often referred to as the ‘micromotion’, while the slower $\mathbf{r}_s(t)$ motion is designated the ‘secular’ motion.

The force due to this ponderomotive potential at location \mathbf{r} is given by

$$\mathbf{F} = -\frac{q^2}{4m\Omega_{\text{rf}}^2} \nabla \mathbf{E}_0^2(\mathbf{r}). \quad (3.14)$$

In the case of the Paul trap, a quadrupole rf potential is used. For two-dimensions, as in Fig. 3.1(a), the potential is of the form,

$$\Phi_{\text{rf}}(\mathbf{r}, t) = \frac{\kappa V_0}{2} \left(1 + \frac{x^2 - y^2}{R^2} \right) \cos \Omega_{\text{rf}} t, \quad (3.15)$$

where V_0 is the rf potential applied to the rf electrodes, R is the distance from the center of the quadrupole to the electrodes, and κ is a geometry factor. For nominally quadrupole electrodes, as in Fig. 3.1(a), κ is near one, but may need to be numerically calculated. The two-dimensional Hamiltonian is,

$$H(t) = \frac{m}{2}(\dot{x}^2 + \dot{y}^2) + \frac{(\kappa q V_0)^2}{4mR^4\Omega_{\text{rf}}^2} (x^2(t) + y^2(t)). \quad (3.16)$$

This is the equation for a two-dimensional harmonic confining potential. The motional frequency is

$$\omega_{\text{rf}} = \frac{\kappa q V_0}{\sqrt{2}mR^2\Omega_{\text{rf}}}, \quad (3.17)$$

which is commonly referred to as the secular frequency and can be adjusted to a desired value by tuning V_0 . An interesting observation is that the secular frequency is inversely proportional to Ω_{rf} . This comes about because a lower value of Ω_{rf} allows the ion to travel farther during an rf cycle, resulting in more kinetic energy and a steeper pseudopotential. However, if Ω_{rf} is decreased too much, \mathbf{r}_μ amplitude can approach \mathbf{r}_s and the assumptions above break down. Generally, Ω_{rf} is chosen to be several times larger than both ω_{rf} and the linewidths of any laser interactions, so that any sidebands that may occur will be well resolved.

At this point, one should recall the assumptions made to reach Eq. (3.13). The first was that $\mathbf{r}_s(t)$ changed slowly on the time scale $2\pi/\Omega_{\text{rf}}$. The second was that $\nabla\Phi_0(\mathbf{r}_s)$ was approximately constant over one cycle of $\mathbf{r}_\mu(t)$. Together, these approximations are known as the ‘adiabatic approximation’. These assumptions are sufficient but not necessary to provide confinement in a quadrupolar rf field. Ion traps can be successfully operated outside of these assumptions, but to accurately describe the dynamics in that case requires the use of the Mathieu equation, which will be covered in Sec. 3.3 .

3.2.2 Linear Paul trap

The two-dimensional Paul trap can be extended to three dimensions using several different geometries. The first is to take the electrodes of Fig. 3.1(a) (shown in cross-section) and rotate them about either the \hat{x} or \hat{y} axis of the figure to form a ring with two electrodes above and below

the ring. This geometry provides harmonic confinement using only rf fields. However, the more relevant geometry for this dissertation is the linear rf Paul trap. In such a trap, the electrodes of Fig. 3.1(a) are extended along the \hat{z} direction to form four parallel rods, as in Fig. 3.1(b). By convention, the \hat{z} direction is referred to as the axial direction, while \hat{x} or \hat{y} are the radial directions. Along a radial cross-section, the rf fields will be identical to that of the two-dimensional trap and radial confinement is provided by the axially symmetric pseudopotential produced by the four rods. At the center of the trap, along the z axis, is a line with zero rf field (in all directions) which is referred to as the rf null or nodal line. If the rods are extended infinitely far along \hat{z} , this geometry will produce no axial rf field and the ion can freely travel along \hat{z} .

To achieve axial confinement, the control electrodes can be segmented into three sections each, for a total of six electrodes. The two outer sections on each rod are referred to as the endcaps and have a high static potential applied to them. The center sections are held at a lower static voltage, creating a potential minimum along the axis in which the ion sits. An expansion of this potential around the ion's equilibrium position will be quadratic to leading order. As long as the ion is localized close enough to this equilibrium position, this potential will create harmonic confinement in the axial direction.

Taking just the leading order of the three-dimensional static potential when expanded in multipoles, results in three orthogonal quadrupoles centered about the origin:

$$\Phi_{\text{st}}(\mathbf{z}) = \frac{U_{xy}}{2} \left(1 + \frac{x^2 - y^2}{R^2} \right) + \frac{U_{zx}}{2} \left(1 + \frac{z^2 - x^2}{R^2} \right) + \frac{U_{zy}}{2} \left(1 + \frac{z^2 - y^2}{R^2} \right). \quad (3.18)$$

In general, there are five orthogonal quadrupole terms in a spherical-harmonic expansion of a potential, but symmetry precludes two of these terms in the linear Paul trap. The magnitude of the U_{ij} 's are determined by the trap geometry and potentials applied to the segments. Here the \hat{x} , \hat{y} , and \hat{z} directions have been chosen to correspond with the eigenvectors (and principal axes) of the Hessian matrix, for convenience, with \hat{z} also being the axial direction of the linear trap. The

Hessian for the potential energy due to static electric potential is

$$\mathcal{H} = \frac{q}{R^2} \begin{bmatrix} U_{xy} - U_{zx} & 0 & 0 \\ 0 & -U_{xy} - U_{zy} & 0 \\ 0 & 0 & U_{zx} + U_{zy} \end{bmatrix}. \quad (3.19)$$

The eigenvalues of \mathcal{H} are related to the frequencies: $\tilde{\omega}_x = \sqrt{q(U_{xy} - U_{zx})/mR^2}$, $\tilde{\omega}_y = \sqrt{-q(U_{xy} + U_{zy})/mR^2}$, and $\tilde{\omega}_z = \sqrt{q(U_{zx} + U_{zy})/mR^2}$. From our discussion in Sec. 3.1, we know at least one of these frequencies must be imaginary and anti-confining.

Including the pseudopotential in the Hessian for the linear Paul trap gives,

$$\mathcal{H}_{\text{tot}} = m \begin{bmatrix} \omega_{\text{rf}}^2 + \tilde{\omega}_x^2 & 0 & 0 \\ 0 & \omega_{\text{rf}}^2 + \tilde{\omega}_y^2 & 0 \\ 0 & 0 & \tilde{\omega}_z^2 \end{bmatrix}. \quad (3.20)$$

Thus, in the radial directions, the overall secular frequencies are given by $\omega_x = \sqrt{\omega_{\text{rf}}^2 + \tilde{\omega}_x^2}$ and $\omega_y = \sqrt{\omega_{\text{rf}}^2 + \tilde{\omega}_y^2}$, while the axial frequency is entirely due to the static potential, $\omega_z = \tilde{\omega}_z$. Though the static frequencies can be imaginary, ω_{rf}^2 is always positive. The static and rf potentials can be chosen so that $\omega_i > 0$ and the trap will provide the desired three-dimensional harmonic confinement. Adjusting the voltages on the electrodes allows tuning of the secular frequencies. Typically, the parameters are set so that the linear Paul trap is highly anisotropic with $\omega_x \approx \omega_y \gg \omega_z$.

Above, the Cartesian coordinates were chosen to align with the principal axes of the potential. If the coordinates are not chosen in this manner, the secular frequencies and principal axes can be determined by finding the eigenvalues and the eigenvectors, respectively, of \mathcal{H}_{tot} .

For quantum operations that involve a mode of motion, such as Raman cooling or two-ion entangling operations, the axial mode is most often chosen because it is easy to adjust and more readily addressed. The axial mode is also important because the transport and rearranging operations in the multiplexed architecture take place along the trap axis. Thus in most quantum computing experiments, focus is placed on the axial mode, while the radial modes are simply spectators.

3.3 Mathieu equation

The adiabatic approximation discussed above is helpful for understanding the operation of the Paul trap and is reasonably accurate for most quantum-computing applications. However, it only applies when the secular frequencies are much smaller than the rf frequency ($\omega_i \ll \Omega_{\text{rf}}$). A more thorough analysis requires the Mathieu equation, which is discussed in a number of places throughout ion-trap literature. Here, we will focus on the most relevant details. For a more comprehensive discussion see, for example, [Ghosh 95].

In an ideal linear Paul trap, the rf potential is given by Eq. 3.15:

$$\Phi_{\text{rf}}(\mathbf{z}, t) = \frac{\kappa V_0}{2} \left(1 + \frac{x^2 - y^2}{R^2} \right) \cos \Omega_{\text{rf}} t. \quad (3.21)$$

In addition, potentials applied to the control segments introduce the static quadrupole potential given by Eq. (3.18). The equations of motion for a charged particle in the combined potential is given by

$$\begin{aligned} \ddot{x} + \frac{q}{mR^2} [(U_{xy} - U_{zx}) + \kappa V_0 \cos \Omega_{\text{rf}} t] x &= 0 \\ \ddot{y} + \frac{q}{mR^2} [-(U_{xy} + U_{zy}) - \kappa V_0 \cos \Omega_{\text{rf}} t] y &= 0 \\ \ddot{z} + \frac{q}{mR^2} [U_{zx} + U_{zy}] z &= 0. \end{aligned} \quad (3.22)$$

The z motion is easily recognized as harmonic motion with secular frequency $\omega_z = \sqrt{\frac{q}{mR^2}(U_{zx} + U_{zy})}$.

The x and y motion may not be as clear. Using the following substitutions,

$$\begin{aligned} q_x &= \frac{2\kappa q V_0}{mR^2 \Omega_{\text{rf}}^2} & q_y &= -q_x & q_z &= 0 \\ a_x &= \frac{4(U_{xy} - U_{zx})q}{mR^2 \Omega_{\text{rf}}^2} & a_y &= -\frac{4(U_{xy} + U_{zy})q}{mR^2 \Omega_{\text{rf}}^2} & a_z &= \frac{4(U_{zx} + U_{zy})q}{mR^2 \Omega_{\text{rf}}^2} \\ \xi &= \frac{\Omega_{\text{rf}}}{2} t \end{aligned} \quad (3.23)$$

all three equations can be transformed into the Mathieu equation

$$\frac{d^2 x}{d\xi^2} + [a_x + 2q_x \cos(2\xi)] x = 0. \quad (3.24)$$

(Only the x equation is written, but the y and z equations can be expressed in the same manner.)

The unitless parameters a_x and q_x are the Mathieu stability parameters. They should not be

confused with the ion-charge q nor the harmonic annihilation operator \hat{a} . The parameter a_x is related to the strength of the static confinement, while q_x is related to the rf confinement.

The Mathieu equation is well studied and shows up in a number of physical systems. The equation is satisfied by Mathieu functions [Abramowitz 64, Mathews 70] with the form

$$x(\xi) = e^{\mu\xi}\phi(\xi), \quad (3.25)$$

where $\phi(\xi)$ is a periodic function with period 2ξ . These solutions are stable and bounded only if $\mu \equiv \alpha_x + i\beta_x$ is purely imaginary, in which case $e^{\mu\xi}$ is sinusoidal instead of exponential in ξ (normalized time). The parameters a_x and q_x dictate the value of μ and, thus, whether the solution is stable or unstable. Equation (3.25) can be expanded using a Fourier cosine expansion to give

$$x(\xi) = A \sum_{n=-\infty}^{\infty} C_{2n} \cos [(2n + \beta_x)\xi + B]. \quad (3.26)$$

The values of A and B are determined by the initial conditions. To find β_x and the coefficients C_{2n} , Eq. (3.26) can be plugged back into Eq. (3.24) to obtain the recursion relation

$$G_{2n}C_{2n} + C_{2n+2} + C_{2n-2} = 0, \quad (3.27)$$

where $G_{2n} \equiv \frac{a_x - (2n + \beta_x)^2}{q_x}$.

From Eq. (3.27), two other equations can be produced. First, setting $n = 0$, gives the equation

$$\frac{\beta_x^2 - a_x}{q_x} = \frac{C_{2n+2}}{C_{2n}} + \frac{C_{2n-2}}{C_{2n}}. \quad (3.28)$$

Next, by rearranging terms in while Eq. (3.27), two continued fractions are found:

$$\frac{C_{2n\pm 2}}{C_{2n}} = -\frac{1}{G_{2n\pm 2} - \frac{1}{G_{2n\pm 4} - \frac{1}{G_{2n\pm 6} - \dots}}}. \quad (3.29)$$

The C_{2n} coefficients can be determined if initial conditions are specified. The more important parameter is β_x , which defines both the frequencies of the various motions and the stability of the trap. By combining Eq. (3.29) and Eq. (3.28), an equation for β_x that can be constructed and continued out to arbitrary accuracy.

The typical parameters used in ion trapping, and those used in this dissertation, fulfill the criteria $|a_x| < q_x^2 \ll 1$. In this case, finding β_x^2 to lowest order in a_x and q_x^2 is sufficient and yields

$$\beta_x = \sqrt{a_x + \frac{q_x^2}{2}}. \quad (3.30)$$

The lowest-frequency motion in Eq. (3.26) occurs when $n = 0$. This motion is the ‘secular’ motion described in the previous section and has frequency $\omega_x = \beta_x \Omega_{\text{rf}}/2$. If there are no static potentials ($\Phi_{\text{st}}(\mathbf{z}) = 0$), this frequency in the radial directions reduces to

$$\omega_x = \omega_y = \frac{\kappa q V_0}{\sqrt{2} m R^2 \Omega_{\text{rf}}}, \quad (3.31)$$

which agrees with the result obtained using the adiabatic approximation, Eq. (3.17). Including a non-zero $\Phi_{\text{st}}(\mathbf{z})$ introduces the corrections to ω_x , ω_y , and ω_z that are due to static trapping and also agree with results of the previous section.

However, this agreement between the adiabatic approximation and the Mathieu treatment only holds if $|a_x|, q_x^2 \ll 1$, where Eq. (3.30) is a good approximation. In fact, this requirement on a_x and q_x^2 is another way of stating the adiabatic approximation. In this dissertation, $a_x \approx 0.008$ and $q_x^2 \approx 0.2$ are typical Mathieu parameters. With these values, using Eq. (3.30) will result in $\sim 5\%$ error in β_x and ω_x relative to the actual value. The exact secular frequencies can be easily measured experimentally if needed for sideband or two-qubit gate operation.

Making the transformation back to the physical variables leads to the equation

$$x(t) = A \sum_{n=-\infty}^{\infty} C_{2n} \cos[(n\Omega_{\text{rf}} + \omega_x)t + B]. \quad (3.32)$$

Figure 3.2 shows solutions for different values of q_x (with $a_x = 0$), including $q_x = 0.45$ which is used in the X-junction trap. As q_x approaches its maximum stable value of 0.908, beating between the secular motion and micromotion becomes evident.

One issue the Mathieu treatment highlights, which the adiabatic approximation missed, is the existence of higher-frequency motion at $\omega_x = (2n \pm \beta_x)\Omega_{\text{rf}}/2$. For $n \neq 0$, these terms correspond to the micromotion that gives rise to the ponderomotive potential in the adiabatic approximation.

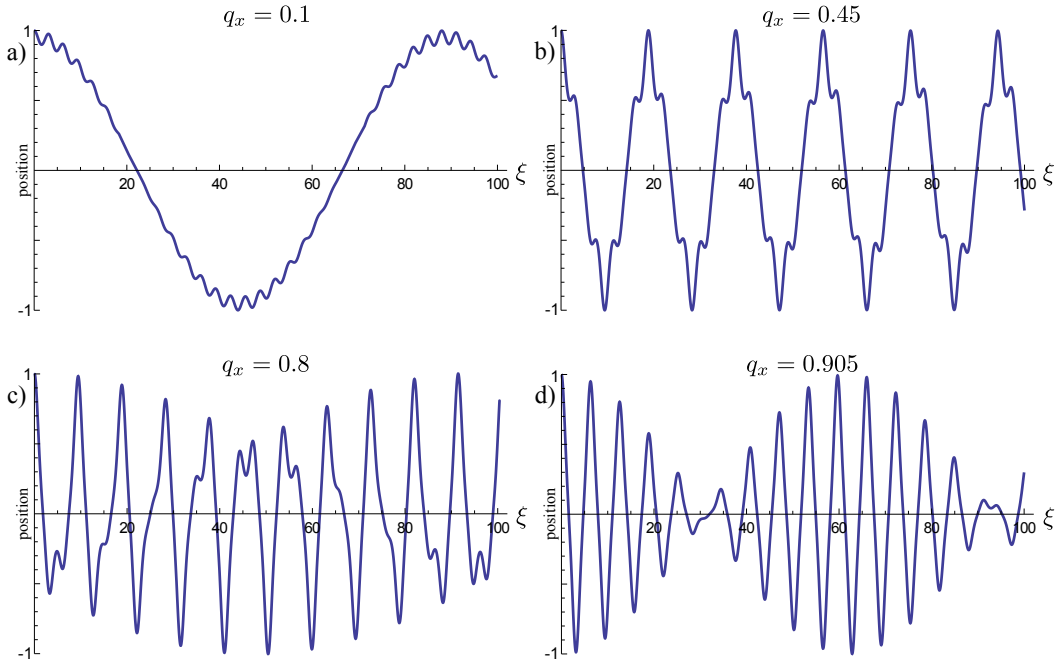


Figure 3.2: Solutions to the Mathieu equation for various q_x parameters when $a_x = 0$. The vertical axis is position, normalized to the maximum excursion, and the horizontal axis is $\xi = \Omega_{\text{rf}}t/2$. a) The solution for $q_x = 0.1$. The large sinusoidal motion is the secular motion while the smaller ripples are the micromotion. b) The solution for $q_x = 0.45$. This corresponds to the typical parameter used in this dissertation. c) The solution for $q_x = 0.8$. d) The solution for $q_x = 0.905$. Beating between the secular motion and micromotion is clearly visible. For $q_x \approx 0.908$, the secular motion and micromotion reach a resonance condition and the overall motion becomes unbounded.

When using the adiabatic assumption $|a_x|, q_x^2 \ll 1$, the C_{2n} coefficients fall off very quickly as n increases and only the terms with $n = 0, \pm 1$ contribute significantly to the motion. In that case, the motion, with the first-order correction, is given by

$$x(t) = x_0 \cos(\omega_x t) \left[1 - \frac{q_x}{2} \cos(\Omega_{\text{rf}} t) \right], \quad (3.33)$$

where x_0 is the magnitude of the secular motion. The ion's motion is predominantly sinusoidal, though there is a small modulation of that motion at Ω_{rf} . Nonetheless, the ion can, to a high degree, be thought of as a harmonic oscillator.

The Mathieu treatment is useful for determining the stability criteria and necessary parameters to achieve trapping. It turns out that the solutions are stable for $0 < \beta_x < 1$ (or $0 < \omega_x < \Omega_{\text{rf}}/2$), which places constraints on both a_x and q_x . As β_x approaches an integer, some of terms in Eq. (3.26) will become resonant with the rf drive at 2ξ . These motions can then be resonantly driven by the trap rf and become unbounded. There are higher bounded regions with $\beta_x > 1$, but most ion traps, including the trap in this dissertation, are operated with $0 < \beta_x < 1$. The β_x parameter can be adjusted through changing a_x and q_x , which depend on the voltages applied to the electrodes and Ω_{rf} . The adiabatic approximation holds when $|a_x|, q_x^2 \ll 1$, meaning β_x is close to 0. However, stable solutions also exist for $\beta_x \approx 1$ where the adiabatic criteria are not satisfied. Thus, the adiabatic approximation is a more restrictive requirement than simply having a bounded solution. This idea will return when discussing the potential near an ion-array junction.

It is also important to remember that there are three directions of motion, each having separate values of β . For complete confinement, all the β 's must satisfy $0 < \beta_i < 1$, simultaneously. However, these parameters are not independent, which further limits the parameter space where every motion is bounded. The stable parameter space can also be augmented by other factors, such as drag from Doppler laser cooling, which increases this range [Dawson 95].

So far, the treatment of the ion's motion has been classical and shows that the ion behaves much like a harmonic oscillator. This is important for using ions in quantum computation, as operations between different ion qubits is mediated by this harmonic motion. Quantum treatments

of the ion motion can be performed [Baseia 93, Ghosh 95, Bardroff 96], using a procedure similar to the Mathieu treatment above. Such analysis shows that ions in a Paul trap behave very much like harmonic oscillators on the quantum level, as well. The main difference between the Mathieu oscillator and a standard quantum harmonic oscillator is that the Mathieu wavepacket ‘breathes’ at the rf frequency. However, this breathing does not significantly alter the quantum behavior of the system and approximating the motion as a simple quantized harmonic oscillator is accurate to a high degree.

We can distill the main features of the ion motion in an ideal Paul trap, where both the rf and control potentials can be described as pure quadrupole fields, with the following points:

- (1) The equation of motion is described by the Mathieu equation, which is well known and has well understood solutions.
- (2) This motion can be accurately approximated as a harmonic oscillator.
- (3) Motion in the three orthogonal directions is decoupled and separable.
- (4) Two parameters, a_i and q_i , completely define whether the motion of the ion in one direction is stable or if the motion will increase in amplitude exponentially.
- (5) Stability of the ion’s motion is not dependent on the initial conditions of the ion.
- (6) The secular frequency, ω_i , is well defined (for small β_i) and independent of the motional amplitude.
- (7) No motion occurs at Ω_{rf} , so the rf drive does not pump energy into the motion.

So far, we have considered single ions, but one can also consider the motion of multiple ions in the same trapping potential. For multiple ions in the adiabatic limit, the rf trapping potential can be replaced by the pseudopotential equivalent and the problem becomes one of interacting particles in a harmonic potential. Ions cooled to the Doppler temperature will form a linear crystal along the axis of the trap (anisotropic trapping potential). This assumes the number of ions N , is small.

As N increases, other types of crystals can form, but in this dissertation N is small, being either 1 or 2.

Instead of having three harmonic motions (one in each direction), N ions in the same trapping well will have $3N$ harmonic modes of motion. Once the ions crystallize, the ions will remain well localized and the nonlinear Coulomb force can be approximated by a linear force. In this case, the $3N$ modes of motion become separable and individually addressable (as long as N is not too large).

If the ions are not well localized, perhaps due to a collision with some background gas molecule, the Coulomb force can lead to mixing of the motions and give rise to other nonlinear or chaotic phenomena. This can enable new avenues through which the rf potential can pump energy into the system [Walther 93, Hoffnagle 95] and lead to instability. Such mechanisms must be studied using the full Mathieu treatment instead of the adiabatic approximation. These complicated dynamics are one reason why the lifetimes of multiple ions in a Paul trap tend to be much shorter than those of single ions. Nonetheless, trapping multiple ions does not change the harmonic picture under standard operating parameters and at low ion kinetic energies.

3.4 Inhomogenous Mathieu equation

An ideal Paul trap consists of only quadrupole rf and static fields. However, non-quadrupolar fields are often present in real traps. These deviations from the ideal come about in many ways, such as trap misalignment. In quantum computing experiments, the ions are confined close enough to the rf null that these terms can usually be neglected. However, much larger deviations from the ideal quadrupole can arise in non-conventional geometries, such as those necessitated by junctions. Since junctions are a main topic of this dissertation, it is important to consider the effects of non-quadrupolar additions to the potentials.

To treat the complete electric potential, it is advantageous to expand around the ion's equilibrium position in terms of multipoles or spherical harmonics,

$$\Phi_{\text{rf}}(\mathbf{r}, t) + \Phi_{\text{st}}(\mathbf{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^l [\phi_{\text{rf},l,m} \cos(\Omega_{\text{rf}}t) + \phi_{\text{st},l,m}] \left(\frac{r}{R}\right)^l Y_l^m(\theta, \phi), \quad (3.34)$$

where R is a scaling factor, often taken to be the distance from the trapping minimum to the nearest electrode. The different rf and static multipole moment coefficients are $\phi_{\text{rf},l,m}$ and $\phi_{\text{st},l,m}$, respectively, while l defines the order of the multipole. The quadrupole terms ($l = 2$) are the only terms present in the ideal Paul trap. The rest of these terms can be divided into two categories: lower-order terms (monopole and dipole) and higher-order terms (hexapole, octopole, etc.). In either case, the addition of such terms can alter the equation of motion, giving an inhomogeneous Mathieu equation of the form,

$$\frac{d^2x}{d\xi^2} + [a_x + 2q_x \cos(2\xi)]x = A(x, y, z, t). \quad (3.35)$$

The effects of lower-order terms will be examined first with the conclusion that they play a relatively minimal role the ion behavior. Then the effects of the higher-order terms will be described, which can be more problematic for ion stability, as they lead to chaotic behavior.

3.4.1 Mathieu equation with monopole and dipole terms

There are only four lower-order terms, three dipole ($l = 1$) and one monopole ($l = 0$). The axes of the three dipole terms can be chosen to lie along the three principal axes of the trap for convenience, without loss of generality. This allows the potential to be separated into three independent directions. Taking one of the directions and expanding in terms of polynomials gives

$$\Phi_{\text{rf}}(x, t) + \Phi_{\text{st}}(x) = \sum_{l=0}^2 [\phi_{\text{rf},l} \cos(\Omega_{\text{rf}}t) + \phi_{\text{st},l}] \left(\frac{x}{R}\right)^l. \quad (3.36)$$

As before, each term of the sum corresponds to a different multipole term with $l = 2$ being the quadrupole term. The equation of motion is

$$\frac{d^2x}{d\xi^2} + [a_x + 2q_x \cos(2\xi)]x = - \left[\frac{4q\phi_{\text{rf},1}}{mR\Omega_{\text{rf}}^2} \cos(2\xi) \right] - \left[\frac{4q\phi_{\text{st},1}}{mR\Omega_{\text{rf}}^2} \right], \quad (3.37)$$

where $a_x = \frac{8q\phi_{\text{st},2}}{mR^2\Omega_{\text{rf}}^2}$ and $q_x = \frac{4q\phi_{\text{rf},2}}{mR^2\Omega_{\text{rf}}^2}$ are the same Mathieu stability parameters, but written a bit differently than before. The monopole term ($l = 0$) disappears when the gradient is taken, so has no effect on the motion, as expected. The left side of Eq. (3.37) is the homogeneous Mathieu equation with typical Mathieu solutions. To analyze the effect of the dipole term on the stability

of the motion, it is helpful to break the problem into three distinct cases, two special cases and one general case.

(i) When $\mathbf{q}_x = \phi_{\text{st},1} = 0$

In this case, the Mathieu equation simplifies to

$$\frac{d^2x}{d\xi^2} + [a_x] x = -\frac{4q\phi_{\text{rf},1}}{mR\Omega_{\text{rf}}^2} \cos(2\xi). \quad (3.38)$$

Here there is no rf confinement, which is often the situation along the axial direction. This is the equation for driven harmonic motion, with the rf dipole term acting as an off-resonant drive. From basic theory, we know that an off-resonant drive will not lead to a net energy gain. The motion is solved using the method of undetermined coefficients:

$$x(\xi) = A \cos(\sqrt{a_x}\xi + B) - \frac{4q\phi_{\text{rf},1}}{(a_x - 4)mR\Omega_{\text{rf}}^2} \cos(2\xi), \quad (3.39)$$

where A and B are arbitrary coefficients that allow for a specific solution. This is a bounded solution as expected, though the dipole term does introduce some residual micromotion.

(ii) When the ratio $\mathbf{a}_x : 2\mathbf{q}_x$ equals $\phi_{\text{st},1} : \phi_{\text{rf},1}$

This corresponds to a special case where the equation of motion can be rewritten in the form

$$\frac{d^2x}{d\xi^2} + [a_x + 2q_x \cos(2\xi)] x = -\chi [a_x + 2q_x \cos(2\xi)]. \quad (3.40)$$

Here the proportionality constant $\chi \equiv \frac{4q\phi_{\text{rf},1}}{mR\Omega_{\text{rf}}^2} \frac{1}{2q_x} = \frac{4q\phi_{\text{st},1}}{mR\Omega_{\text{rf}}^2} \frac{1}{a_x}$. Equation (3.40) can be transformed back into the homogeneous Mathieu equation using a simple change of variables, $\tilde{x} = x + \chi$. The solutions are Mathieu functions shifted by χ , and obey all of the stability criteria of the standard Mathieu function. This result can be thought of as simply using the wrong origin for solving the problem. This solution also holds when either $a_x = 0$ or $q_x = 0$, as long as $\phi_{\text{st},1} = 0$ or $\phi_{\text{rf},1} = 0$, respectively.

In the specific case that $q_x = 0$, but $a_x, \phi_{\text{st},1}, \phi_{\text{rf},1} \neq 0$, an exact solution can be constructed by first performing a change of variables to $\tilde{x} = x + \frac{4q\phi_{\text{st},1}}{mR\Omega_{\text{rf}}^2} \frac{1}{a_x}$ to remove the static dipole term. Then, following case (i), the rf dipole term can be handled as an off-resonant drive. The position

in this case is

$$x(\xi) = A \cos(\sqrt{a_x} \xi + B) - \frac{4q\phi_{\text{rf},1}}{(a_x - 4)mR\Omega_{\text{rf}}^2} \cos(2\xi) - \frac{4q\phi_{\text{st},1}}{mR\Omega_{\text{rf}}^2} \frac{1}{a_x}, \quad (3.41)$$

which is just a shifted version of case (i).

(iii) The general case of static and rf dipole terms

When a linear (dipole) term is added to a quadratic term, it only acts to shift the minimum of the quadratic term. A simple change of variables that moves the origin to the shifted minimum can be used to remove the linear term. However, Eq. (3.37) has a both an rf and a static dipole term, meaning that it is likely not possible to shift to the minima of both dipole terms simultaneously. It is possible to shift to either the rf or static minimum by making the change of variables $\tilde{x} = x + \chi$, with χ properly chosen, to remove one dipole term from the equation. Depending on which term is removed, this results in either of the following equations,

$$\frac{d^2 \tilde{x}}{d\xi^2} + [a_x + 2q_x \cos(2\xi)] \tilde{x} = -\frac{4q}{mR\Omega_{\text{rf}}^2} \left(\phi_{\text{st},1} - \frac{a_x}{2q_x} \phi_{\text{rf},1} \right) \quad (3.42)$$

$$\frac{d^2 \tilde{x}'}{d\xi^2} + [a_x + 2q_x \cos(2\xi)] \tilde{x}' = -\frac{4q}{mR\Omega_{\text{rf}}^2} \left(\phi_{\text{rf},1} - \frac{2q_x}{a_x} \phi_{\text{st},1} \right) \cos(2\xi). \quad (3.43)$$

In fact, one of these equations can be transformed into the other by making a different choice of χ where $\tilde{x} \leftrightarrow \tilde{x}'$. This implies the solution of the inhomogeneous Mathieu equation with dipole terms can be solved either as a static dipole term centered about one location or an rf dipole term centered around another location.

For now, we will work with Eq. (3.42). Make the ansatz [Rosenband 10]:

$$\tilde{x}(\xi) = A \sum_{n=-\infty}^{\infty} C_{2n} \cos[(2n + \beta)\xi + B] + \sum_{n=-\infty}^{\infty} D_{2n} \cos[2n\xi], \quad (3.44)$$

where the first sum is a Mathieu function and the second sum will solve the inhomogeneous equation. To determine the values of D_{2n} , the same procedure for finding the Mathieu functions in the homogeneous problem can be followed here. Plugging the ansatz into Eq. (3.42) leads to another recursion relation,

$$C'_{2n} + \frac{D_{2n+2}}{D_{2n}} + \frac{D_{2n-2}}{D_{2n}} = \frac{P\delta_n}{qD_{2n}}, \quad (3.45)$$

where δ_n is the Kronecker delta function, while we define $G'_{2n} \equiv \frac{a_x - (2n)^2}{q_x}$ and $P \equiv -\frac{4q}{mR\Omega_{\text{rf}}^2} \times \left(\phi_{\text{st},1} - \frac{a_x}{2q_x} \phi_{\text{rf},1}\right)$. For $n \neq 0$, two continued fractions for the D 's can be found:

$$\frac{D_{2n\pm 2}}{D_{2n}} = -\frac{1}{G'_{2n\pm 2} - \frac{1}{G'_{2n\pm 4} - \frac{1}{G'_{2n\pm 6} - \dots}}}. \quad (3.46)$$

Combining Eq. (3.46) with Eq. (3.45) for $n = 0$ finds a solution for D_0 that satisfies the inhomogeneous offset P :

$$D_0 = \frac{P}{q_x} \frac{1}{G'_0 - 2 \frac{1}{G'_2 - \frac{1}{G'_4 - \dots}}}. \quad (3.47)$$

This equation takes advantage of the fact that $D_{2n} = D_{-2n}$. From these equations, the coefficients of Eq. (3.44) can be determined to any desired accuracy.

A similar procedure can be followed to find solutions for the rf dipole in Eq. (3.43). Alternately, it may be simpler to transform Eq. (3.43) into Eq. (3.42) with the appropriate change of variables.

The additional motion, introduced by the by static and rf dipole term, is bounded and stable. Thus, it does not affect the stability of the overall motion. However, depending on the size of the dipole force, the extra induced micromotion can be substantial.

3.4.2 Mathieu equation with nonlinear terms

If higher-order multipole terms are present in the trapping potential, they can be of more concern than the monopole and dipole terms just discussed. Such terms result in Eq. (3.35) being nonlinear, which significantly complicates the analysis. In fact, nonlinear terms can easily lead to chaotic behavior. Some forms of the nonlinear Mathieu equation have been treated analytically in an attempt to determine stability criteria (see for example [El-Dib 01]). However, historically, solving the nonlinear Mathieu equation involves computer simulation.

Ion-trap research has generally tried to avoid higher-order multipole terms. Though the ideal Paul trap is purely quadrupolar, unwanted multipole terms result from trap misalignments or the use of approximations to the ideal hyperbolic electrode geometry with some simpler shape.

Higher-order terms due to such inaccuracies have typically been quite small. An extensive body of literature on this topic exists in the field of mass spectroscopy [Dawson 69, Alheit 96, Gudjons 97], where instabilities in the quadrupole mass filter due to nonlinear terms were described as early as [von Busch 61]. Mass spectroscopy often uses large collections of ions that extend over a significant portion of the trap. Some of these ions can lie far enough from the rf null that the influence of even small higher-order terms is noticeable. In fact, such terms can even be used to increase mass selectivity and are occasionally engineered into the apparatus [Michaud 05].

However, ion traps used in quantum computing typically employ a small number of ions confined very near the axis of the trap where the higher-order terms should be much smaller than the quadrupole terms. Thus, this topic has been largely neglected in the field of ion-trap quantum computing. However, with the advent of new trap designs, the issues of the nonlinear Mathieu equation may soon permeate this field as well. For example, these nonlinear rf terms are significantly larger for surface-electrode traps [Chiaverini 05], as will be discussed later in the chapter. And these terms may be unavoidable near junctions. Here, I will cover some of the issues that may arise with nonlinear potentials and why they should be kept in mind when designing ion traps for quantum computing.

The introduction of multipole terms with $l > 2$ into Eq. (3.34) results in four main differences from the features of pure quadrupole potentials listed at the end of Sec. 3.3:

- (1) The motion along different principal axes can be coupled and therefore not separable.
- (2) The secular frequency depends on the amplitude of the motion.
- (3) The ion motion may be chaotic and depend on the initial conditions, which may reduce the range of stable operating parameters.
- (4) Nonlinear resonances occur that lead to excitation of the ion's motion or even loss of the ion.

Examining the spherical harmonics that describe multipole potentials reveals why the motion

in different directions is coupled. Spherical harmonics of order l are composed of polynomials of the form $x^{n_x} \times y^{n_y} \times z^{n_z}$ where $n_x + n_y + n_z = l$. For example, the seven real (not complex) hexapole ($l = 3$) harmonics are

$$\begin{aligned}
H_1 &= \frac{1}{4} \sqrt{\frac{7}{\pi}} \frac{(2z^2 - 3x^2 - 3y^2)z}{R^3} \\
H_2 &= \frac{1}{4} \sqrt{\frac{35}{2\pi}} \frac{(3x^2 - y^2)y}{R^3} \\
H_3 &= \frac{1}{4} \sqrt{\frac{35}{2\pi}} \frac{(x^2 - 3y^2)x}{R^3} \\
H_4 &= \frac{1}{4} \sqrt{\frac{105}{\pi}} \frac{(x^2 - y^2)z}{R^3} \\
H_5 &= \frac{1}{2} \sqrt{\frac{105}{\pi}} \frac{xyz}{R^3} \\
H_6 &= \frac{1}{4} \sqrt{\frac{21}{2\pi}} \frac{(4z^2 - x^2 - y^2)y}{R^3} \\
H_7 &= \frac{1}{4} \sqrt{\frac{21}{2\pi}} \frac{(4z^2 - x^2 - y^2)x}{R^3}
\end{aligned} \tag{3.48}$$

All the terms include mixtures of x , y and z . Similarly, the gradient of each term and the equation of motion that depends on that gradient will be inseparable functions of multiple directions. This inseparability is a general feature of all multipoles with $l > 2$ and complicates the motion.

The second consequence of nonlinear terms is that the secular frequencies can shift as the amplitude of the motion increases. The secular frequency ω_i is related to the curvature of the potential. In hexapole and higher terms, this curvature is a function of the distance from the origin. Though a changing ω_i is not ideal for quantum computing purposes, it can actually increase the stability of the trap [Franzen 91]. For example, if an unwanted constant oscillatory force is present and resonant with ω_i , it will drive energy into the motion. In a quadrupole trap, this will continue indefinitely. If higher-order terms are present, ω_i shifts as the amplitude increases, tuning the motion out of resonance. Eventually, the relative phase between the motion and drive will change and energy will be pulled out of the secular motion.

However, in other cases, the combination of amplitude dependence on ω_i and the mixing of motion in different directions can lead to chaotic behavior and, potentially, ion loss. This raises the question, under what circumstances will nonlinear potentials lead to unstable trapping? Given the

complexity of this motion, that is difficult to determine and is generally treated using computer simulation. A different way to tackle this question is using the adiabatic assumption of Sec. 3.2.1. If the amplitude of the micromotion is small enough so that the rf field can be approximated as linear over a single rf cycle, the ion is expected to remain stable. To determine the applicability of this assumption, an adiabatic parameter can be defined [Gerlich 92] that measures the fractional change in the rf field over a single micromotion cycle:

$$\eta_x = \frac{|2(\mathbf{x}_\mu \cdot \nabla)(\nabla\Phi_{\text{rf}})|}{|(\nabla\Phi_{\text{rf}})|}. \quad (3.49)$$

The numerator is the micromotion amplitude \mathbf{x}_μ (due to a displacement along \hat{x}) multiplied by the derivative of the electric field along the direction of the micromotion, while the denominator is the electric field magnitude. This can be simplified to

$$\eta_x = \frac{2q|(\hat{x}_\mu \cdot \nabla)\nabla\Phi_{\text{rf}}|}{m\Omega_{\text{rf}}^2}, \quad (3.50)$$

where \hat{x}_μ is the unit vector in the direction of the micromotion at position x . This measures adiabaticity for displacements along the x axis. Adiabaticity in the other directions is measured by η_y and η_z . In a purely quadrupole rf potential, these η_i 's reduce to the Mathieu q_i parameters. In a higher-order potential, the η_i 's have a position dependence while the q_i 's do not.

When $\eta_i \ll 1$, the adiabatic criterion holds and the trap will be stable. Unfortunately, this does not necessarily tell us at what value of η_i the trap crosses over to unstable. From the Mathieu equation discussion, we know that the ideal quadrupole trap can be stable when q_i is near 1. Thus, in a strongly quadratic trap, $\eta_i \sim 1$ can result in stable behavior. However, for potentials with larger nonlinear terms, η_i will need to be smaller. In [Gerlich 92], it was suggested that $\eta_i < 0.3$ is sufficiently small to assure stability in arbitrary geometries, but that this value might be unnecessarily small for many traps. Calculating η_i for new trap designs can help identify if the trap may become unstable in certain locations. If η_i is small enough in all trapping locations, this can provide an alternative to exhaustive computer simulations of the ion dynamics in the potential.

The addition of higher-order terms to a predominantly-quadratic rf potential can lead to nonlinear resonances that excite the secular motion. These resonances arise from the mixing of the

different secular motions. Recall that the basic solution to the Mathieu equation can be written as Eq. (3.32)

$$x(t) = A \sum_{n=-\infty}^{\infty} C_{2n} \cos [(n\Omega_{\text{rf}} + \omega_x)t + B]. \quad (3.51)$$

If a small l^{th} -order multipole is added as a perturbation to a primarily-quadrupole potential, we would expect to see terms in the motion that go as

$$x^{n_x}(t)y^{n_y}(t)z^{n_z}(t) \propto \prod_{j=x,y,z} \left(\sum_{k=-\infty}^{\infty} C_{j,2k} \cos [(k\Omega_{\text{rf}} + \omega_j)t + B_j] \right)^{n_j}, \quad (3.52)$$

where $n_x + n_y + n_z = l$. Expanding this and finding all the cross terms is a tedious exercise infinite sums. Fortunately, we can focus on just a handful of those terms that give rise to the nonlinear resonances. The terms with $k = 0$ have the largest $C_{j,2k}$ coefficients, and the rest of the coefficients fall off quickly as $|k|$ increases (assuming $\beta_j \ll 1$). If we take the $k = 0$ terms and expand, the highest-frequency term is of the form

$$\cos [(n_x\omega_x + n_y\omega_y + n_z\omega_z)t + B]. \quad (3.53)$$

For certain secular frequencies, this term will oscillates at Ω_{rf} and the motion can be resonantly excited by the trap drive. If other terms from Eq. 3.52 are included, resonant excitations can also occur at any set of secular frequencies that fulfill

$$n_x\omega_x + n_y\omega_y + n_z\omega_z = \nu\Omega_{\text{rf}}, \quad (3.54)$$

where ν is an integer. This defines an extensive series of specific trap frequencies and parameters that should be avoided, especially when all the possible integer values of ν are considered. Fortunately, when ν is large, the resonances become increasingly difficult to drive and, in practice, are rarely seen.

These nonlinear resonances have been observed in a number of ion-trap mass-spectroscopy experiments [Alheit 96, Alheit 97, Drakoudis 06] and are described in greater analytical detail in many papers such as [Dawson 69, Wang 93, Konenkov 02].

Three conditions must be met for a nonlinear resonance to occur:

- (1) The secular frequencies must fulfill Eq. (3.54).
- (2) If Eq. (3.54) holds for a given set of n_i 's, then the l^{th} -order multipole term that satisfies $n_x + n_y + n_z = l$ must be present in the potential.
- (3) The ion must sample a region of the potential where that multipole term is large enough to influence the motion.

Condition (3) is rarely met for single ions confined close to the rf null, unless an occasional collision with a neutral background-gas molecule excites the ion far from the null.

In trap designs that are not as strongly quadrupolar as the traditional Paul traps were, nonlinear resonances may lead to heating of the ion motion if the ion is being pushed off axis by stray fields. This may explain some instabilities seen in surface-electrode traps. Similar instability could occur in the complicated potentials near junctions. In fact, as an ion transports through a trap array, the secular frequencies can sweep through a large range of values which makes it more likely that the resonance criteria are temporarily fulfilled at some point during the transport. Thus, paying attention to this and other issues of nonlinear potentials can be important when dealing with novel trap designs.

Some related issues where nonlinear terms affect ion dynamics have been studied. For example, the impact of higher-order terms on the axial secular frequency has been examined [Home 10]. Also, the influence of the Coulomb repulsion between multiple ions confined in the same trapping well has been considered [Marquet 03, Roos 08, Nie 09].

3.5 Real trap geometries

Now that the theory has been explained, we can discuss actual trap geometries and design requirements. This section will begin by covering a number of general design criteria and then focus on some of the common geometries for linear Paul traps that are currently being used. This topic of junction design will be covered in a subsequent chapter.

When beginning a trap design, perhaps the most obvious consideration is the size of the trap which we characterize with R . This affects the secular frequencies, the optical access, and the heating rate, among other things.

In current quantum computing research, the secular frequencies typically range from 10 kHz for heavier ions to 20 MHz for light ions such as ${}^9\text{Be}^+$. Higher secular frequencies can reduce both ion heating rates and the influence of stray fields, while increasing the speed of quantum operations. So increasing these frequencies can be beneficial. Equation (3.17) shows that the rf component to the secular frequency scales as R^{-2} . Thus decreasing R can dramatically increase ω_{rf} . The axial frequency ω_z also benefits from smaller traps, though it only scales as R^{-1} . Over the last few years, there has been a big push towards miniaturizing ion traps to take advantage of this scaling and traps as small as $R \approx 30 \mu\text{m}$ are becoming common.

Unfortunately, ions traps are plagued by an anomalous heating mechanism that increases for small R . The exact origin of this heating is unknown, but a general model of fluctuating patch potentials on the surface of the electrodes seems to fit [Turchette 00] and predicts a scaling of R^{-4} . This heating mechanism excites the harmonic motion of the laser-cooled ions and can reduce the fidelity of quantum operations if not kept in check. The dependence of this heating on R was measured using ions trapped between two translatable needles [Deslauriers 06], where the heating rate, \dot{n} , was shown to scale as $R^{-3.5 \pm 0.1}$. The anomalous heating mechanism is thermally activated and can be reduced significantly through cryogenic cooling of the trap electrodes [Labaziewicz 08]. The heating mechanism's undesirable scaling with R makes it difficult to work with very small traps, at least at room temperature.

When designing traps, attention must also be paid to electrical characteristics. For example, if the electrodes become too thin or narrow, their resistance can lead to electrode heating or unwanted voltage drops. Similarly, a large capacitance between rf electrodes and ground will increase the rf power consumption. The traps are typically driven by a rf resonator and an excessive trap resistance can spoil the Q of the resonator. This reduces the filtering factor of the resonator and can lead to increased motional heating of the ions due to noise entering along the rf line. This is of particular

concern in junctions, as discussed in Sec. 7.6.1. As experiments move to incorporate larger numbers of qubits, more trapping zones will be required, which could lead to larger capacitances and smaller Q -factors. Therefore, attention should be paid to keep the capacitance and other impedances in the trap as low as possible.

Another concern in trap design is the direction of the principal axes of secular motion. In a linear Paul trap, one principal axis will almost always align with the trap axis. The other two axes are defined by the static component of the trapping field. The direction of these axes are crucial for laser interactions that involve the ion motion, such as Doppler cooling, Raman ground-state cooling, and Raman gate interactions.

In the case of Doppler cooling, the laser beam must have a component of the beam's \mathbf{k} -vector along any mode of motion that is intended to be cooled. The Doppler cooling process involves radiating photons in random directions, which leads to heating in all directions. In any direction that is not simultaneously being cooled, this heating can continue unchecked. As long as the component of \mathbf{k} along all three principal axes is large enough, Doppler cooling can be done with a single beam. The angle between the beam and each axis should be $> 15^\circ$ for this to work [Itano 82]. If two axes are degenerate, the laser will define one axis along which it has maximal \mathbf{k} overlap and one where \mathbf{k} is perpendicular. This leads to heating in that second direction. Thus it is important that the principal axes are degenerate and that they point in convenient directions for laser cooling. In some trap geometries, this is not always the case and must be taken into consideration [Chiaverini 05].

There are a number of other considerations when designing with a trap geometry. First there needs to be sufficient optical access for both the lasers and the detection optics for collecting resonantly-scattered photons. This tends to favor more open geometries. On the other hand, more closed geometries provide better shielding of the ion from stray charges that might build up on insulating components near the trap, such as vacuum windows. The confinement depth of the trap is also highly dependent on the geometry. A larger depth makes loading ions more easy while increasing trap lifetimes.

Of course, the ease with which a trap geometry can be fabricated is an important practical consideration in trap design. The use of microfabrication techniques for ion-trap quantum computing has increased significantly over the past decade. Microfabricated traps can be made more precisely and smaller than other traps. Microfabrication techniques have been developed well enough that the time required to build a trap can drop by an order of magnitude compared to conventional construction. The surface-electrode trap architecture [Chiaverini 05], which will be discussed shortly, lends itself to microfabrication techniques.

3.5.1 Common linear trap architectures

A number of general trap architectures have been used in recent years, each with its own benefits and drawbacks. These traps can be divided into three general categories: macroscopic architectures, wafer architectures, and microfabricated architectures.

Macroscopic architectures

These traps generally resemble the four-rod design depicted in Fig. 3.1 and have been commonly used for decades. They are typically fabricated out of macroscopic rods positioned ~ 1 mm apart, making relatively large traps. Their size requires large rf voltages of ~ 1 kV to get high secular frequencies. However, the size also results in low electrical impedances, low anomalous heating, good optical access, and deep trapping depths. These traps are ideal for precision applications, such as atomic clocks, or for trapping large collections of ions. For large-scale quantum computing, these traps are less promising. We would like to have arrays of traps, each with a small number of ions. The size of macroscopic traps would make such arrays prohibitively large.

Wafer architectures

This architecture is of special importance to this dissertation as it is what the X-junction array is based on. In the wafer architecture, the four rod design of Fig. 3.1 is approximated using two wafers [Rowe 02]. Each wafer has a single channel cut completely through it. The two sides of the channels are electrically conducting (typically with a metal coating) and will approximate two of the electrodes in Fig. 3.1. A second wafer is held just beneath the first, as in Fig. 3.3. To provide




Figure 3.3: a) Two wafers can be used to approximate the ideal Paul trap [Rowe 02]. Each wafer has a channel cut through it and then is metalized to define the electrodes. Small slits on the control electrode side create the control segments. The two wafers are then placed very close to each other (with a spacer not shown), creating a linear Paul potential in the channel, between wafers. b) The cross section of the two wafer trap showing that the two wafers approximate the four-rod design with an ion trapped along the axis.

axial trapping, the control electrodes are segmented using narrow slits. A three wafer version is also possible (Fig. 3.5(c)). The advantage of these designs is that all the electrodes can be produced in two or three wafers, reducing the alignment requirements. These traps have been made using laser machining which can produce $\sim 15 \mu\text{m}$ features, allowing the overall size of the trap to be reduced to $R \sim 100 \mu\text{m}$. This allows for more traps in a smaller area, and higher secular frequencies. The wafer architecture has been used commonly over the past decade in experiments at NIST and elsewhere.

The two-wafer designs have a few advantages over three-wafer traps. First, the smaller number of wafer makes assembly and alignment of the wafers easier. Though seemingly trivial, precision alignment is one of the most difficult requirements to satisfy with wafer traps. A second benefit it that the principal axes in the two-layer design point towards the electrodes while the trap has easy beam access in a plane perpendicular to the wafers. This allows the principal axis determination for Doppler cooling to be easily satisfied. However, in a three-wafer trap, the principal axes are rotated by $\sim 45^\circ$, causing the Doppler beam to be perpendicular to one axis. This must be circumvented using different laser angles, alternate trap designs, or more complicated static potentials to realize Doppler cooling in the three-layer architecture.

The two-wafer design is often constructed so that the width of the channel is approximately

equal to the distance between wafers, as shown in Fig. 3.3. In [Stick 06], this aspect ratio was increased from 1 to 15 due to details necessitated by the microfabrication of that trap. The effects of such high aspect ratios were discussed in [Madsen 04]. However, an aspect slightly below one is advantageous because it results in a more ideal quadrupole. The electrodes in Fig. 3.3(b) extend farther in the horizontal direction than in the vertical direction, giving rise to a larger field between the wafers than vertically. By pulling the wafers farther apart, this effect is balanced to create a more ideal quadrupole. This balanced two-wafer architecture was used for the X-junction trap. The proper wafer spacing depends on the shape of the electrodes in cross section, but for the wafers used in the X-junction trap, the desired ratio is 0.91.

The disadvantage of wafer architectures is their relatively large size and difficulty of fabrication. Though both of these aspects are improved over macroscopic architectures, further improvement is desired for large-scale quantum computing. The main issue with fabrication is that laser machining of ceramic wafers is a specialized technique with limited resolution. For example cut widths are typically not smaller than $15 \mu\text{m}$. Also, since the trap is assembled from wafers, it is prone to misalignments when the wafers are stacked together.

Microfabricated architectures

Microfabrication techniques are widely used in many fields and are becoming highly refined. Leveraging this technology for ion traps allows a number of exciting possibilities. The tolerances and alignment of microfabricated structures are typically extremely good ($\sim \mu\text{m}$). This allows much smaller traps to be produced, with some current traps having $R \sim 30 \mu\text{m}$. Associated with this is an increase in secular frequencies and the potential for faster quantum operations. Bringing the ion in close proximity to the trap structure can also provide for stronger coupling between the ion and objects on the surface of the electrode. For example, coupling to the magnetic, electric, vibrational, or even optical properties of the trap structure could be realized. Microfabrication also promises the ability to incorporate other technologies into the ion trap, such as on-chip optics or CMOS electronics.

The standard four-rod trap architecture is difficult to approximate with most microfabrication

technologies. Some realizations of such traps have been proposed [Folman 05] and tested [Stick 06, Britton 08]. However, other architectures are better suited for microfabrication. These generally fall under the category of surface-electrode architectures, or surface traps for short [Chiaverini 05, Seidelin 06].

In a surface trap, the four rods of a linear Paul trap are deformed and pulled into a single plane, as depicted in Fig. 3.4. This can either create what can be called a four-wire or a five-wire geometry. In both cases, the ion is trapped in an approximately quadrupole field located above the surface by a distance approximately equal to the electrode width. These two geometries are quite similar, with the only significant difference being that the principal axes tend to point parallel and perpendicular to the surface for a five-wire design, while they point $\sim \pm 45^\circ$ in the four-wire design. This allows an ion in a four-wire architecture to be Doppler cooled by a single laser beam parallel to the surface. Cooling with a single beam in a five-wire design implies the beam must come in at an angle to the surface, and might need to be reflected off the surface, which could lead to excessive scattered light. Alternatively, by applying more complicated potentials to the control electrodes, both transverse principal axes can have a component parallel to the surface. Also, by adjusting the widths of the five wires, the principal axes can be rotated enough that a single beam parallel to the surface can cool the ion [Chiaverini 05].

Having all the electrodes in a single plane makes surface traps significantly easier to fabricate. Alignment of trap features can be completely taken care of using fabrication techniques, such as photolithography.

There are penalties incurred with surface traps. The surface architectures include large multipole contributions to their potential. A primary problem with this is that it reduces the trap depth, which can be two orders of magnitude smaller than the other architectures [Chiaverini 05]. These terms can also lead to unwanted nonlinear effects discussed above. Surface traps also lack the symmetry of multi-dimensional designs. Though this is not necessarily a problem, the rotational symmetry (about the trap axis) of a traditional linear Paul trap simplifies the trap characteristics. For example, in a traditional linear trap, the symmetry suppresses rf fields that might occur along

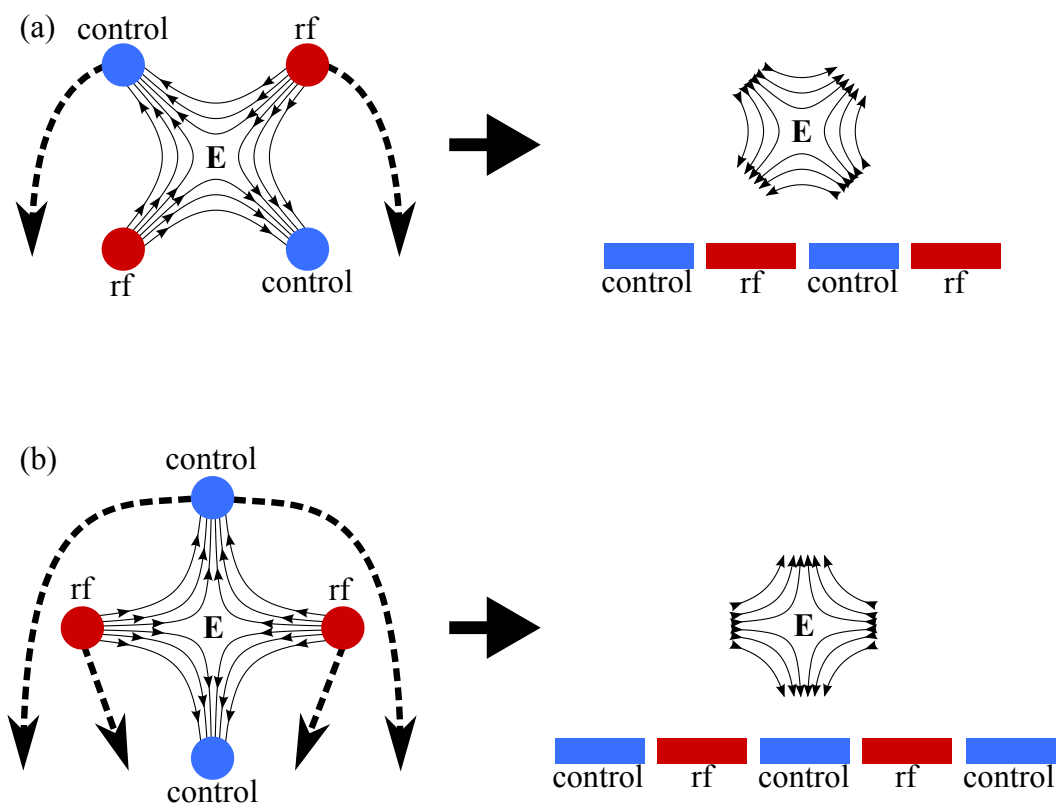


Figure 3.4: a) Deformation of a four-rod Paul trap into a four-wire surface trap. The top two rods are pulled down and outward until they lie in the same plane as the lower rods. This creates a quadrupole field located above (and another below) the surface. The four-wire design has principal axes that are aligned at $\sim 45^\circ$ to the surface. b) Showing the deformation of a four-rod Paul trap into a five-wire surface trap. The top rod is split in two, and then all the rods are pulled down into a plane. One of the principal axes in this case is parallel to the surface.

the axis. However, in a surface geometry, axial rf fields and corresponding axial micromotion can arise from notches in the electrodes or if the trap electrodes taper along the axis. Furthermore, ions are more susceptible to unwanted external fields due to charges on distant parts of the apparatus, such as the vacuum windows. Recent experience with surface traps have generally shown that ions are better behaved in the multi-dimensional designs [Amini 09]. Hopefully, as the technology matures surface architectures will become as reliable as their traditional counterparts.

Other microfabricated architectures are possible. For example, another surface design, the in-plane four-wire design, can be made by deforming the four-rod design into four wires in a plane, but with both rf electrodes on the inside or outside (see Fig. 3.5(e)). In this case the ion is trapped in the plane of the electrodes with the principal axes also parallel or perpendicular to that plane. This geometry shares many characteristics with the other surface architectures.

3.5.2 Comparison of multipole terms in different trap geometries

The desire to miniaturize traps in recent years has led to a number of novel configurations that have certain disadvantages compared to the older-style traps. Higher-order potentials and nonlinear dynamics may be present in such geometries. This can be mitigated by using ions accurately positioned along the rf null, where higher-order corrections to the quadrupole field are vanishingly small, even for these novel designs. In practice, it is not always simple to move an ion onto the rf null. This may contribute to some of the performance issues seen in these newer configurations [Amini 09]. Regardless, understanding the design or operation requirements of these new trap configurations is important. Here we will compare the magnitude of the higher-order corrections to the ideal quadrupole potential for eight trap configurations that have been demonstrated. This should help to quantify the extent to which nonlinear dynamics may play a role in these designs.

This comparison will focus only on the radial directions of ion motion. Axial considerations, such as junctions and electrode segmentation along the axis, will not be covered. Such issues would introduce too many variables to draw any clear conclusions. It seems reasonable to assume that if a design is based on a radial geometry with fewer problems, the overall design will have fewer

multipole term	N	function
monopole	0	$P_{0,a}(x, y) = 1$
dipole	1	$P_{1,a}(x, y) = x$
	1	$P_{1,b}(x, y) = -y$
quadrupole	2	$P_{2,a}(x, y) = x^2 - y^2$
	2	$P_{2,b}(x, y) = -2xy$
hexapole	3	$P_{3,a}(x, y) = x^3 - 3xy^2$
	3	$P_{3,b}(x, y) = y^3 - 3x^2y$
octopole	4	$P_{4,a}(x, y) = x^4 - 6x^2y^2 + y^4$
	4	$P_{4,b}(x, y) = 4xy^3 - 4x^3y$

Table 3.1: The first few multipole functions found using Eq. (3.55)

problems.

Restricting the model to the transverse directions reduces the number of applicable multipole terms to those that have translational symmetry along the axis. Only two distinct multipole moments per order, N , fulfill this. An analytical representation for the two multipole functions, $P_{N,i}(x, y)$, can be found by calculating the real and imaginary parts of

$$(x + iy)^N. \quad (3.55)$$

The first few multipole functions found in this manner are given in Table 3.1. The two separate multipole functions for each order are identical except that one is rotated from the other in the xy -plane by an angle $\frac{\pi}{2N}$. The transverse potential can then be expanded about the location of the ion as

$$\Phi(x, y) = p_0 + \sum_{N=0}^{\infty} \frac{1}{R^N} [p_{N,a}P_{N,a}(x, y) + p_{N,b}P_{N,b}(x, y)], \quad (3.56)$$

where $p_{N,i}$ are the multipole moments, and R is the distance between the trapping location and the nearest electrode. By using rotated coordinates, it is possible to pick an orientation where only one multipole term is needed for a given order. These new coordinates are $x'_N = x \cos \theta_N - y \sin \theta_N$ and $y'_N = x \sin \theta_N + y \cos \theta_N$, where the rotation angle is given by $\theta_N = \frac{\arctan(p_{N,b}/p_{N,a})}{N}$. In general, a different set of rotated coordinates (x'_N, y'_N) is needed to reduce each multipole order. However,

using different coordinates allows a simplification of Eq. (3.56) to

$$\Phi(x, y) = p_0 + \sum_{N=0}^{\infty} \frac{1}{R^N} p'_N P_N, a(x'_N, y'_N), \quad (3.57)$$

where $p'_N = \sqrt{p_{N,a}^2 + p_{N,b}^2}$. This allows more direct comparison of the magnitude of the multipole moments at the expense of keeping track of the relative orientations.

The eight geometries to be considered are depicted in Fig. 3.5 and include:

- (1) A standard 2-wafer trap. The dimension between wafers and between opposing electrodes is equal, which is typical for traps used at NIST.
- (2) The balanced 2-wafer trap. This is nearly identical to the previous geometry except that the spacing between the wafer is expanded by $\sim 10\%$ to more closely approximates a perfect quadrupole. This is the geometry the X-junction trap uses.
- (3) A 3-wafer trap. As will be discussed later, three layer designs can be used for T-junction arrays; this design was taken from [Hensinger 06].
- (4) A 2-layer trap with a high aspect-ratio. The high aspect ratio of this design makes it more conducive to GaAs/AlGaAs microfabrication [Stick 06].
- (5) In-plane 4-wire trap [Chiaverini 05, Blain 08]
- (6) The 4-wire surface trap [Chiaverini 05]. The dimensions are typical of NIST surface traps.
- (7) The 5-wire symmetric surface trap [Chiaverini 05]. Also using typical dimensions for NIST traps.
- (8) The 5-wire asymmetric surface trap [Chiaverini 05]. This is identical to the previous geometry, but widths of two electrodes are adjusted to rotate the principal axes to a more convenient angle for laser cooling. In this example, the width of one electrode is reduced by 10% while the other is increased by 10%.

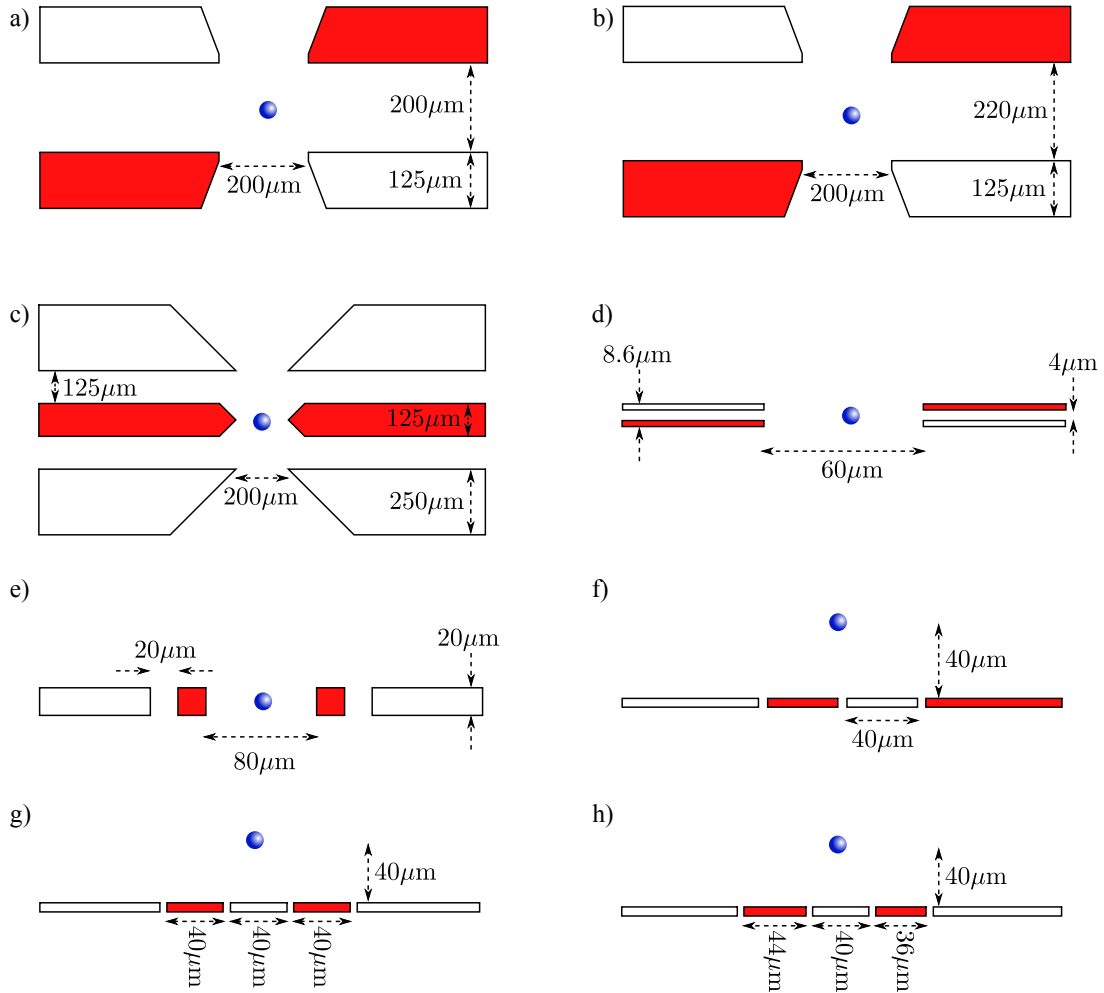


Figure 3.5: Cross-section of various trap geometries. The rf electrodes are red, the ground electrodes are white, and the trapping location of the ion is depicted by a blue dot. a) A typical 2-wafer geometry. b) A balanced 2-wafer geometry with a larger inter-wafer spacing that creates a more-quadrupole field. c) The 3-wafer design used in [Hensinger 06]. d) The 2-layer design used in the GaAs/AlGaAs trap in [Stick 06]. e) An in-plane 4-wire trap. f) A 4-wire surface trap. g) A 5-wire symmetric surface trap. h) A 5-wire asymmetric surface trap.

Table 3.2: The magnitudes of the higher-order multipole moments for the different trap geometries described in Fig. 3.5. The moments are reported as the ratio of the n^{th} -order moment to the quadrupole moment. This normalizes the moments relative to the secular frequency and removes the R dependence for each geometry.

Trap geometry	p'_3/p'_2	p'_4/p'_2	p'_5/p'_2	p'_6/p'_2	p'_7/p'_2	p'_8/p'_2
2-layer	0.0015	0.0284	0.0005	0.1171	0.0036	0.0039
balanced 2-layer	0.0011	0.0158	0.0009	0.1156	0.0005	0.0242
3-layer	0.0006	0.1904	0.0006	0.0554	0.0005	0.0203
2-layer AlGaAs	0.0026	0.6005	0.0054	0.4340	0.0143	0.2642
in-pane 4-wire	0.0008	0.3443	0.0009	0.1191	0.0010	0.0330
4-wire surface	0.9731	0.6931	0.4255	0.2442	0.1481	0.1074
5-wire symm. surface	0.9963	0.6202	0.2640	0.0408	0.1504	0.1426
5-wire asymm. surface	1.0100	0.6401	0.2737	0.0569	0.1195	0.1455

To compare different geometries, the rf potentials were modeled (see Sec. 5.2) and then fit to extract the p'_N moments. The ratio of the higher-order moments to the quadrupole moments are reported in Table 3.2. This ratio normalizes out the secular frequency and is a better measure of the extent to which the design deviates from the ideal quadrupole. The higher-order moments are much larger in some geometries implying they will be significantly more prone to nonlinear dynamical effects.

In Table 3.2, some of the entries that are non-zero (such as p'_3/p'_2 for (a)-(e)) should be zero due to symmetry considerations. However, the data in the table requires both an electrostatic model and a multipole fitting routine, which both have some inherent numerical error. This gives an idea of the precision of these calculations.

The dramatically larger hexapole term in the surface geometries can be intuitively understood by considering the electric field. Figure 3.6(a) shows the field lines above a 4-wire surface trap. These lines must terminate on the electrodes, forcing them to bend downward farther from the rf null. For comparison, Fig. 3.6(b) and (c) show the field of an ideal quadrupole and hexapole, respectively. An appropriately weighted sum of these two components will very closely match the field in (a), as demonstrated in Fig. 3.6(d). It is the addition of the hexapole term that, to first order, allows the field lines to bend and conform to a geometry where the electrodes are in a single plane. Any 4-wire surface geometry will have the same requirement, making it difficult to




Figure 3.6: a) The electric field above a 4-wire surface trap. The inner-most $50 \mu\text{m}$ is modeled by computer, while the dotted line shows the general behavior at greater distances. The quadrupole and hexapole components of this field are shown in b) and c), respectively, with a colored ring representing the orientation of the given component. The appropriately weighted sum of the quadrupole and hexapole terms is shown in d). Together, these lowest two terms closely match the total field in a), showing that a hexapole term is paramount in providing the bending of the field lines that allows the field to conform to a surface geometry.

significantly reduce the size of the hexapole (and higher) terms.

A similar argument can be made for 5-wire surface geometries, though the quadrupole and hexapole terms will be oriented differently. Much like the surface geometries may mandate a large hexapole term, the in-plane 4-wire and the 3-layer designs require a large octopole term to conform to their electrodes.

An interesting observation is that the magnitude of the multipole terms is comparable between all the surface geometries examined. Thus there seems to be little advantage to a four- or five- wire geometry, in terms of reducing non-quadrupole behavior.

Chapter 4

Junction architectures and the X-junction array

An important task in computing is the rearranging of information to prepare for subsequent gate operations. In the multiplexed architecture for trapped-ion quantum computing [Wineland 98, Kielpinski 02], this rearranging takes the form of physically moving ion qubits and placing selected ions in the same trapping well so that multi-qubit gate operations can be performed between the ions. The standard linear Paul trap can easily be extended to produce a linear array of trapping regions with ions distributed amongst different trapping wells along that array. Using transport, separation, and combination waveforms, these ions can be placed in the same trapping well as their neighbors, or distributed to different wells. However, switching the order of ions requires moving into a second dimension. Using multiple linear arrays that intersect at an angle, forming a junction, will allow ions to be moved into the second dimension, though other methods exist [Splatt 09]. Such junctions are an efficient way to perform rearranging of information in a trapped-ion architecture and should lead to a scalable quantum computing system.

Junctions can come in many possible configurations. Three linear arrays can connect at right angles to form a ‘T’ junction or at 120° to form a ‘Y’ junction. Another possibility is to use four linear arrays connecting at right angles to form an ‘X’ junction. Of course, there are limitless number of potential designs, incorporating non-equal angles or a larger number of intersecting arrays. One could even imagine making a three-dimensional junction that can move ions into the third direction, though this may be challenging to fabricate and difficult to deliver the required potentials to the control electrodes. Since the main purpose of a junction, the ability to switch ion

order, can be achieved with any of these designs, there is not a strong computational advantage for having more than three incoming arrays, or legs, to the junction. Perhaps having more legs in a junction may allow for a more efficient or compact architecture, but for now, the question is whether any type of useful junction can be realized. With that in mind, the discussion here will be limited to the ‘X’ and ‘T’ designs, which are the most straightforward configurations.

The potentials in and around a junction are more complicated than in the standard linear Paul trap and the different junction designs provide different relative advantages. This chapter will compare a couple of junction designs based on the multi-wafer traps discussed in the previous chapter. After considering these different junction configurations, I decided on a specific two-wafer X-junction configuration. The second part of this chapter will describe the design of the trap I built, including both the junction and the rest of the trap geometry.

4.1 Junction architectures

The state of the art trap technology when I embarked on this dissertation work involved gold-coated laser-machined alumina wafers stacked to form a wafer trap like those described in Sec. 3.5 [Rowe 02]. The plan was to extend such a structure to incorporate a junction based on either the two- or three-wafer geometry (Fig. 3.5(a)-(d)). The first junction demonstrated used the three layer geometry [Hensinger 06] while here I have opted for two-layer design.

4.1.1 Two-layer junctions

The two-wafer trap (Fig. 3.5(a)-(b)) most closely resembles the ideal linear Paul trap. The electrodes are arranged so that in cross-section opposite electrodes (e.g., top-left and bottom right) are held at the same rf potential. Figure 4.1(a) depicts a possible X-junction using a two-wafer geometry. The control and rf electrodes are swapped in the top layer relative to the bottom layer so that in all four legs leading up to the junction, the diagonally opposite electrodes are at the same rf potential. This produces the desired radial quadrupole field in each leg. The control electrodes would be further segmented (not shown) to produce the an axial trapping well along the axis of




Figure 4.1: a) X-junction using two layers of electrodes. This junction cannot harmonically confine ions at the center. The x, y axes have been defined rotated by 45° from what was used in the previous chapter to better align with the junction and this new definition will be used for the remainder of the dissertation. b) Placing two rf bridges across the junction, as shown, will allow for harmonic confinement at any location along the four legs of the X-junction (see text).

each leg and to allow transport of ions through the array.

The design in Fig. 4.1(a) has a key flaw: it is not harmonically confining at the center of the junction. This can be seen by considering the trap symmetry.

Three perpendicular planes, xy , yz , and xz , cross at the center of the junction. Across each of these planes the geometry is mirrored but the potentials are antisymmetric.¹ On the surface of each of these planes, the rf electric field *must* be perpendicular to that plane. Where two perpendicular planes cross, the field *must* be zero, as it cannot be perpendicular to both planes simultaneously. Thus along this line, there is no pseudopotential confinement (which is what occurs on the axis of a linear Paul trap). At the center of the junction, all three planes cross, which results in no pseudopotential confinement along \hat{x} , \hat{y} , or \hat{z} . From Sec. 3.1 and Laplace's equation, the static fields from the control electrodes cannot simultaneously harmonically confine in three perpendicular directions and must anti-confine in at least one direction. Thus, the addition of a static control potential on top of the pseudopotential cannot lead to harmonic confinement at the center of the junction.

Although it might be possible to steer the ion off of the trap axis to avoid the problem at the center [Hucul 08], another approach is to introduce two bridges that cross the junction, as in Fig. 4.1(b). These bridges break the symmetry and effectively ‘plug’ the $\pm\hat{y}$ escape routes at

¹ The control electrodes are effectively the inverse of the rf electrodes. Though the control electrodes are held at rf ground in practice, identical trapping could be achieved by applying $\frac{1}{2}V_{\text{rf}}$ to the rf electrodes and $-\frac{1}{2}V_{\text{rf}}$ to the control electrodes, since the potential difference is what matters.

the center. This allows harmonic confinement at the center. The bridges shown connect the rf electrodes, one on the top wafer and one on the bottom wafer. It is also possible to produce the same effect using bridges that connect the control electrodes instead, though this is not optimal because it would force the electrodes on opposite sides of the junction to have the same voltage. When moving ions through the junction, these control voltages are adjusted and it is helpful to be able to apply different voltages to different sides of the junction.

Though these bridges remove the escape routes, they do it at the expense of introducing an axial rf field along the four legs of the junction ($\pm\hat{x}$ and $\pm\hat{z}$). This axial field produces an axial pseudopotential barrier along all four entrances to the junction. Figure 4.2 shows an example of two pseudopotential barriers as a function of the axial position along z . The height of these barriers is large and is a significant fraction of the transverse pseudopotential trapping depth (the amount of energy needed to escape the pseudopotential in the radial direction). This fraction depends on the specifics of the geometry used, but is generally $\sim \frac{1}{3}$ to $\sim \frac{1}{2}$. Varying the width of the rf bridges was not seen to have a strong effect on the confining potentials or the height of these barriers. Two additional barriers exist in the two perpendicular legs of the junction.

The pseudopotential barriers depicted in Fig. 4.2 are for the specific X-junction array used in this dissertation. The height of these barriers depends on the exact trap geometry and dimensions, as well as the rf voltage, rf frequency, ion mass, and ion charge. The standard values for these last four parameters in this dissertation are 204 V (peak), 83.2 MHz, 9 amu, and e^+ , respectively, resulting in a barrier height of $\sim \frac{1}{3}$ eV. The asymmetry is due to misalignment of the two trap wafers. Barriers in traps using different dimensions or operating conditions would be qualitatively similar but have a different magnitude.

When transporting through this junction, the ions must be transported over these large barriers. An ion at the apex of one of these barriers will feel a strong anti-confinement in the axial direction, due to the pseudopotential. (This anti-confinement in the case of the X-junction array, with standard operating parameters, was $\omega_{\text{rf},z} = 2\pi \times 5i$ MHz, where the square of the frequency is proportional to the curvature of the potential.) In the radial directions, the pseudopotential is




Figure 4.2: Two of the axial pseudopotential barriers along z produced by inserting rf bridges in the two-wafer X-junction. This is a plot of the pseudopotential for the actual junction used in the X-junction array, where the channel width of the linear array is $200 \mu\text{m}$ and the center of the junction is at 0 in the plot. The asymmetry between the two barriers is due to a slight misalignment of the trap wafers.

still confining at the apex (with $\omega_{\text{rf},x} \approx \omega_{\text{rf},y} \approx 2\pi \times 13 \text{ MHz}$). If the control electrodes create a larger axial confinement than this anti-confinement, the total potential will be a harmonic confining well. However, if the control potential is strongly confining along the axis, it must be strongly anti-confining along one or both of the radial directions. To achieve three-dimensional trapping, we require a set of control voltages such that $\tilde{\omega}_i^2 + \omega_{\text{rf},i}^2 > 0$ for all three directions. This will not be possible with every geometry and one should check the junction design for compliance with this, especially at the apex of the barrier. The X-junction array does satisfy this criteria.

In the center of the junction, the pseudopotential produces three-dimensional harmonic confinement. Ions can be trapped in this location without any aid from the control electrodes.

In summary, the introduction of rf bridges has produced a trap geometry that allows for harmonic confinement at any location during a transport protocol that moves ions along the axis of the linear arrays and through the junction.




Figure 4.3: a) A T-junction design using the three-wafer geometry. b) An X-junction design. In both of these junctions the rf and control layers can be switched.

4.1.2 Three-layer junctions

An alternate approach to a junction is to use the three-wafer geometry (Fig. 3.5(c)). The middle wafer can be used as the rf electrode, while the top and bottom wafers would be segmented to form control electrodes. Conversely, this geometry can be inverted, where one segmented control-electrode wafer is in the middle and two rf-electrode wafers on top and bottom. There is little difference between these two configurations.

Figure 4.3 shows two possible junction designs using a three-wafer geometry, a ‘T’ and an ‘X’. A ‘Y’ junction would also be possible by distorting the ‘T’ into a ‘Y’.

Unlike the two-wafer X-junction (without bridges), the three-wafer X-junction does not have the problem with confinement at the center. Here the xy and yz planes have reflective rf symmetry, where any field in the plane must be parallel to that plane. This does not force the same requirement that was seen in the two-wafer case where the field must be zero along the intersection line of two of these planes. In fact, the three-wafer junction looks very similar to the bridged version of the two-wafer X: the ion is harmonically confined in all three directions at the center and along each entrance to the junction ($\pm\hat{x}$ and $\pm\hat{z}$) there are pseudopotential barriers. The height and appearance of the barriers is quite similar to the bridged two-wafer X.

It has been shown that these pseudopotential barriers will be present in any possible X-junction, as long as there are four perpendicular legs and all the legs lie in the same plane (and the third direction, \hat{y} , is required to be confining) [Wesenberg 09]. This same paper suggests that three-dimensional X-junctions, where the four legs cross at a shallow angle, can be created without

barriers, though fabrication of such a trap would be challenging. Small barriers can be achieved for X-junctions where the linear arrays lie in a plane but cross at a very shallow angle. Such designs are intriguing for future work.

The three-layer T-junction design is qualitatively very similar to the X. It also has pseudopotential barriers, one for each leg. This was the general design used in [Hensinger 06] and the geometry is described there and in [Hucul 08]. A Y-junction would have similar characteristics to a T.

Given the similarities between the two- and three-layer X-junctions, there is little theoretical advantage of one over another. However, there are practical considerations that can differentiate the two. When the top and bottom layers are chosen to be control electrodes, the three-layer design requires twice the control electrodes to achieve the same level of complexity and number of trapping regions as a two-layer trap. This doubles the required number of electrical connections. In addition, when fabricating the ion trap, there will be three layers to assemble instead of two. Since all the layers need to be aligned to each other, adding a third layer makes alignment more challenging. For a laser-machined alumina trap, like the one in this dissertation, where the wafers are hand assembled, alignment of even two wafers will lead to errors and adding a third wafer significantly compounds the problem. Primarily for these reasons, I chose to pursue the two-wafer design with rf bridges over a three-wafer design.

In the time since the X-junction array was designed, simple surface-electrode geometries (Fig. 3.5(f)-(h)) have been demonstrated [Seidelin 06]. These surface-electrode traps are more amenable to microfabrication techniques which will allow easy scaling to large number of electrodes. A surface-electrode ‘Y’ junction has been designed and tested [Amini 10], but motional excitation and reliability were not measured.

One key goal of the design in [Amini 10] was to reduce the height of the pseudopotential barriers in the legs of the ‘Y’, due to concern over the magnitude of these barriers. The rf and control electrodes were distorted to create higher-order terms in the pseudopotential that canceled each other to produce a lower overall height for the pseudopotential barriers. At the same time, this




Figure 4.4: Trap wafers. There are three alumina wafers: the top and bottom trap wafers have electrodes on them, and a wafer between those two produces the correct spacing between the electrode wafers. The white area are exposed alumina, the yellow are gold traces and electrodes. Two screw holes (for holding the stack together) and two gauge-pin holes (for alignment) can be seen.

produced a spatial ripple in the pseudopotential. It is possible to use similar techniques to reduce the height of the barriers in the two- and three-wafer junctions. However, as we will see in Ch. 7, the height of barrier seems less important to the ultimate performance of the junction than the slope of the barrier. Therefore, any attempt to reduce the barrier height that introduces a greater slope (e.g., this ripple) may not be advisable. As long as the control potentials are capable of producing a sufficiently-deep harmonic trapping minimum at any point along the pseudopotential barrier, the height of that barrier seems less important.

4.2 X-junction array

After considering multiple possibilities for junction configurations, the two-wafer X-junction with rf bridges was determined to be the most reasonable given the fabrication constraints of a hand-assembled, laser-machined, gold-coated alumina wafer trap [Rowe 02]. The next step was to design a trap array incorporating such a structure.

Figure 4.4 shows the trap wafers that incorporate such a junction. The trap is made out of alumina (an Al_2O_3 ceramic) that has been laser machined to produce through slots or trenches in the wafer. In addition, narrow cuts ($20 \mu\text{m}$) are made to separate the different control electrodes. The wafers are then coated with gold to produce electrodes. There are three wafers: the top and bottom trap wafers contain the electrodes and a spacer wafer the separates the trap wafers from each other. The spacer wafer is a square ring in which the center has been removed so that trap wafers are supported at the edge but the electrodes do not contact the spacer. A set of screw holes allow screws to clamp the wafers together and a pair of gauge-pin holes are meant for temporary gauge-pins to be inserted to help align the three wafers. The bottom wafer is made slightly larger in one direction so that the gold traces are visible at the edges and wires can be more easily attached to the all the traces.

This section will introduce the trap geometry, focusing on the electrode structure. The wafers, how they were fabricated, and how they were assembled will not be covered here and will, instead, be explained in Ch. 6.

The trap design for the X-junction array can be seen in Fig. 4.5. The trap contains 46 control electrodes that produce 18 possible trapping zones, including one zone at the center of the junction labeled \mathcal{C} . The array includes two separate experiment zones, \mathcal{E} and \mathcal{F} , where the lasers will interact with the ions. Next to \mathcal{E} and \mathcal{F} are two separation zones where trapping wells can be separated and combined over a narrow $100 \mu\text{m}$ wide electrode pair (electrodes 19/20 and 33/34).

In the vertical ($+\hat{x}$) leg of the junction is another zone labeled \mathcal{V} . Though the junction has four legs, only three of them contain enough control electrodes to independently produce a trapping well that can hold ions. This was done to make fabrication and wiring more simple, and there is no fundamental reason why the fourth leg of an X-junction could not be used.

Zones \mathcal{E} , \mathcal{C} , \mathcal{F} , and \mathcal{V} , comprise the four destinations that will be used in the transport protocols later in this dissertation. There is also a load/reservoir zone labeled \mathcal{L} which is where Be atoms are ionized and first trapped.

The trap dimensions are similar to those in [Barrett 04, Jost 10]. The width of the channel

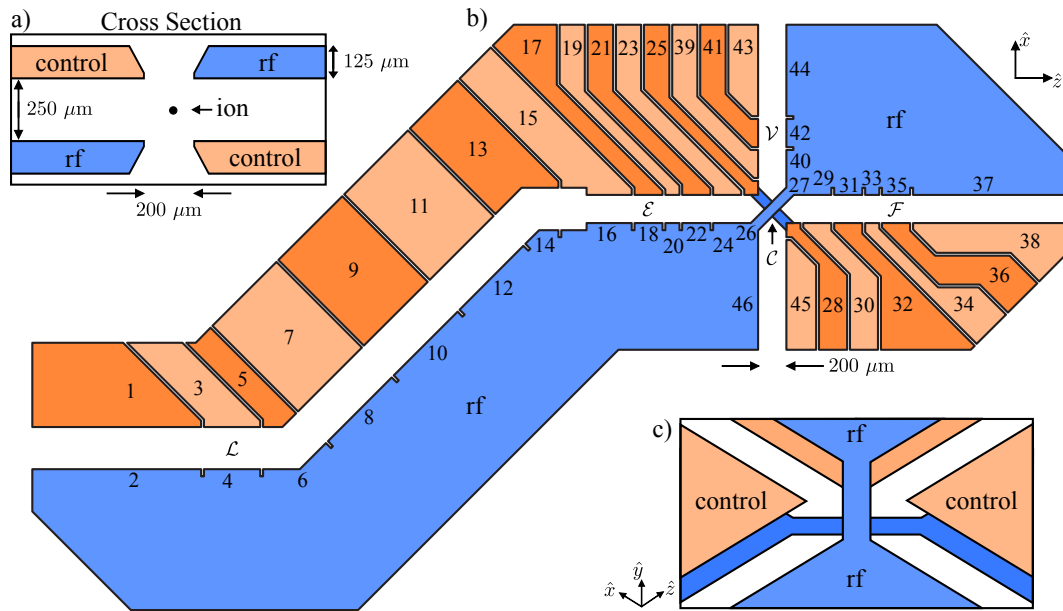


Figure 4.5: a) Cross-sectional view of the two layers of electrodes in the X-junction array. b) Top view of the electrode layout, with the rf electrodes indicated, and all other (control) electrodes held at rf ground. A nearly identical set of electrodes sits below these electrodes, with rf and control electrodes interchanged. Each control electrode is numbered, including the obscured electrodes on the second layer. Forty-six control electrodes support 18 different trapping zones. The load zone (\mathcal{L}), the main experiment zone (\mathcal{E}), the vertical zone (\mathcal{V}), the horizontal zone (\mathcal{F}) and the center of the junction (\mathcal{C}) are labeled. c) Schematic of the rf bridges from an oblique angle (not to scale).

between the rf and control electrodes is $200\ \mu\text{m}$. The trap in [Barrett 04] had a $200\ \mu\text{m}$ space between wafers to match the channel width, but in the X-junction array this distance was designed to be $220\ \mu\text{m}$. This large spacing provides a more quadrupolar field in the radial directions as discussed in Sec. 3.5.2. For unknown reasons (perhaps the wafers warped when screwed together), this distance in the assembled trap was larger than intended and was measured to be $250\ \mu\text{m}$.

Each of the following subsections will be dedicated to a specific region of the trap and will discuss the design considerations at that location.

4.2.1 X-junction

When designing the junction, there are two main aspects that must be considered: the rf bridges and the control electrodes. The bridges were designed to be $70\ \mu\text{m}$ wide. This dimension was chosen primarily because the laser-machining company advised against a narrower width for reasons of strength. However, as stated earlier, computer models suggest that the trap potentials do not strongly depend on the bridge width and $70\ \mu\text{m}$ was a reasonable value.

The width of the control electrodes is more important to the functioning of the junction. As shown in [Reichle 06], the width should be chosen to produce maximum curvature of the control potential on axis of the trap. If the electrodes are too wide, this curvature is small and it becomes difficult to counter the anti-confinement of the pseudopotential at the apex of the barrier. The control electrodes on the four corners of the junction, electrodes 25 to 28, extend $100\ \mu\text{m}$ along the trap axis (both the \hat{z} axis for the horizontal legs of the junction and the \hat{x} for the vertical legs). This width is close to optimal in terms of producing a good curvature, but this width was primarily chosen because there was concern that narrower electrodes might not survive the laser-machining process.

4.2.2 Experiment zone

The X-junction array contains two experiment zones, \mathcal{E} and \mathcal{F} , one on each side of the junction. These are the regions where the lasers will interact with the ion. The lasers can pass

through the channel if they are incident from $\pm\hat{y} \pm \hat{z}$. The two zones are roughly identical, so only \mathcal{E} will be described. Zone \mathcal{E} is centered between electrodes 17 and 18, which are both $200 \mu\text{m}$ wide. A harmonic axial trapping well is produced by setting 17 and 18 to a low voltage and placing electrodes 15, 16, 19, and 20 (and usually 21 and 22) at higher voltages.

The experiment zones were placed far enough from junction to ensure the pseudopotential barrier did not extend into these zones because interacting with the lasers, one ideally wants minimal micromotion. Thus one should try to minimize the axial pseudopotential at the experiment location so the trap looks like a standard linear Paul trap there. At $880 \mu\text{m}$ from the center of the junction, the axial pseudopotential at \mathcal{E} is 2.9^{-5} eV with a 8.7×10^{-8} eV/ μm axial gradient. This gives rise to $z_\mu = 47$ nm micromotion.

There are notches in the rf electrode across from each gap between control electrodes. These notches, and the gaps, are $20 \mu\text{m}$ wide, which is the minimum focus of the laser used for machining the wafers. The notches are about $50 \mu\text{m}$ long. The purpose of the notches was primarily to increase the radial symmetry of the rf potential in order to reduce axial micromotion. These also served as a valuable tool for examining the alignment of the two trap wafers, since they should overlap with the gaps in the wafer below them if the trap is viewed normal to the trap surface.

Next to each experiment zone are two pairs of narrow electrodes, 19/20 and 33/34. These electrodes are $100 \mu\text{m}$ wide (nearest to the ion) and can be used for separating or combining trapping wells. A significant reduction in the motional excitation while performing the separation protocol was seen in [Barrett 04] over experiments in a previous trap [Rowe 02]. This improvement is largely credited to the use of a narrower pair of separation electrodes that were $100 \mu\text{m}$ wide. Thus this geometry was repeated here.

The \mathcal{V} zone may also appear as a reasonable location to interact ions with lasers. Unfortunately, due to the aspect ratio of the trap opening, laser beams cannot pass through the channel if they are incident from $\pm\hat{y} \pm \hat{z}$, as they can in \mathcal{E} and \mathcal{F} . Instead, the laser beam would have to come along $\pm\hat{x} \pm \hat{y}$. A magnetic field is applied along $-\hat{y} - \hat{z}$ to define a quantization axis and allow the $\pm\hat{y} \pm \hat{z}$ beams to be σ^\pm polarized. Since σ^\pm polarization is necessary for many of these

beams, it is not possible to simultaneously obtain the correct laser polarization in both \mathcal{E} and \mathcal{V} if the B -field is uniform between the zones (though it is possible for \mathcal{E} and \mathcal{F}). Therefore, \mathcal{V} is intended only as a place to store ions. In spite of this limitation, an incorrectly polarized beam can be used in \mathcal{V} to perform crude detection or Doppler cooling and this was done when verifying that ions were successfully arriving in this zone during some of the transport experiments.

4.2.3 Load zone

The X-junction array also includes a load region located far away from the experiment zones. The process of loading ions involves heating a sample of metal (Be or Mg) to evaporate neutral atoms. These travel ballistically from the source and a fraction of the atoms pass through the center of the trapping potential. A laser beam is used to photo-ionize the atoms in this region. The newly created ions will have significant kinetic energy but, as soon as they are ionized, they begin to interact with the Doppler cooling beam that intersects the load zone. A high power, far detuned beam, such as the BDD beam, is used for initial Doppler cooling to quickly remove the ion's kinetic energy and stabilize it in the trapping potential.

There is circumstantial evidence that buildup of the neutral Be on the gold trap electrodes may lead to an increase in anomalous heating of the ion [Turchette 00]. For this reason, a separate loading zone was included in the X-junction array so that the loading process and the experiment can be performed in different locations. By placing the loading zone far from \mathcal{E} , the neutral Be built up on the electrodes at \mathcal{L} will be too far from \mathcal{E} to strongly affect the ions during the transport experiments. In addition, a physical barrier is included to shield the experiment region from neutral Be flux from the oven. This barrier is a sheet of stainless steel bent at a right angle to form an 'L' shape in profile. This is positioned just above the surface of the top trap wafer (the ions are trapped between the two wafers). The positioning of this shield can be seen in Fig. 4.6. This shield creates a wall with the oven and \mathcal{L} on one side and the experiment region on the other. This shield is electrically grounded.

The axis of the load zone is displaced laterally from that of \mathcal{E} through the use of two 135°




Figure 4.6: The oven is positioned above the plane of trap (positive y), in roughly the location depicted. Neutral Be sprays from the oven down onto the trap and some of the oven flux will pass through the trap's channel, where it is photo-ionized by a laser beam. This beam passes through the trap parallel to $-\hat{y} + \hat{z}$, making a 45° angle to the xz plane of the page. Laser beams also pass through the experiment zone at 45° to the xz plane. To block oven flux from accumulating on the surfaces of the electrodes near the junction, an oven barrier obscures the line-of-sight between the oven and the experiment zone. This shield is made of a small sheet of stainless steel bent at a right angle. It is positioned just above the trap electrodes as shown in the figure, but extends 1.6 cm along \hat{y} out of the plan of the page.

bends. The purpose of this is to allow laser beams to pass into both the \mathcal{L} and \mathcal{E} without impacting the L-barrier. These laser beams enter parallel to $-\hat{y} + \hat{z}$. Since the shield extends a significant distance along \hat{y} in Fig. 4.6, the shield would block one of the beams passing through \mathcal{E} if these two zones had the same x coordinate. By translating the load zone and the shield to lower values of x , the beams can pass by the shield into both zones. The shield does block access to the load zone from the $-\hat{y} - \hat{z}$ direction. However, the load zone only needs one beam path for the Doppler and photo-ionization beam (as opposed to the experiment zone which needs perpendicular Raman beams). Blocking the other path is not a problem in \mathcal{L} .

The two 135° bends do not pose a problem for transport. The potential in such a bend, though distorted by the curving of the trap axis, is very close to the potential that exists in a standard linear array.

The channel in the load zone is $300 \mu\text{m}$ wide which is slightly wider than in the experiment zone. This provides a larger volume through which the oven flux can pass and is intended to increase the numbers of atoms being ionized to improve loading rates. There is a step transition

from this wider channel to the narrower channel in electrodes 15 and 16. If the trap wafers are well aligned, this step will not produce any rf field along the axis of the trap and transporting ions along this region will not be a problem.

Electrodes 3 and 4, at the load zone, are $400\ \mu\text{m}$ wide. Electrodes 7 to 14 are used primarily for transporting ions between the load and experiment zones and were made relatively wide ($641\ \mu\text{m}$).

Note that the load zone of the trap can also be used as a reservoir zone. During the loading process, a large number of ions could be loaded and only the desired number would be transported to the experiment zones. The rest of the ions could stay in the load zone as a reserve. As ions in the experiment zone are lost due to collisions, they could be replaced by ions out of the reservoir, without repeating the entire load process. This would help increase ion lifetime because the load process requires heating a filament that degrades the vacuum for up to an hour after loading. By avoiding subsequent loading cycles, the frequency of collisions with background gas molecules will be reduced and the ion lifetime will increase.

Chapter 5

Ion transport

The ultimate goal of ion transport procedures is to reliably move ions between two separate locations in the minimum amount of time while leaving them in the same harmonic motional state that they started in (ideally the ground state). If the transport is not highly reliable, ions will be lost and the quantum computation will suffer. Also, an excessive transport duration will slow the quantum processor and allow more time for qubit decoherence. If the secular motion is excited, the fidelity of any subsequent gate operation will be reduced. Motional excitation can be corrected by implementing sympathetic cooling [Barrett 03], but such cooling is slow and would, again, allow additional time for the qubit to decohere. Thus it is important to carefully consider the transport procedures in an ion-trap array.

The act of moving an ion is rather simple, as ions respond readily to electric potentials. Application of a positive voltage to one electrode will cause the ion to move away from that electrode. However, if the intent is to move ions quickly, over long distances, while maintaining low excitation of the ion's secular motion, and perhaps traversing non-trivial potential landscapes such as those near junctions, such a simple approach to transport is unlikely to yield the desired results. Instead, this requires carefully constructed transport protocols where multiple control potentials are adjusted simultaneously in a choreographed waveform that transports the ion.

The first demonstration of ion transport occurred in 2002 [Rowe 02]. In that experiment, ions were transported 1.2 mm down a linear array in $\sim 50 \mu\text{s}$ with minimal excitation of the secular motion, no qubit decoherence, and essentially no ion loss. In addition to this linear trans-

port, that report also described a protocol where two ions were separated from a single trapping well into two wells, albeit with reduced success probability and significant motional excitation. This separation protocol was later perfected and used in a quantum teleportation experiment [Barrett 04]. Since then, transport protocols have been extended to transporting through a T-junction [Hensinger 06] and switching ion order [Hensinger 06, Splatt 09]. To realize the multiplexed architecture in [Kielinski 02] will require further refinement of these transport protocols, possibly in new trap geometries.

In this section, I will explain the process and considerations used for constructing the transport protocols that transport ions through the X-junction trap array. The chapter will begin with a simple description of ion transport in a linear array. This is followed by a brief discussion of the procedure used to create a computer model of the trap potentials. Section 5.3 will cover the singular value decomposition (SVD) method for determining a set of voltages that produces a harmonic trapping potential and the constraints used to define that potential. Section 5.4 will discuss how these voltages are then combined into a waveform that produces a moving potential well. When doing this, some key weaknesses in the SVD method are uncovered. The linear least squares (LLS) method addresses many of these shortcomings while still using the same framework as the SVD method, and it is the LLS method that is actually used to find the waveforms in this dissertation. The chapter will conclude by reviewing some of the waveforms used to transport through the X-junction array.

5.1 Transport introduction

Before describing the details of generating transport waveforms, it will be helpful to consider the transport in general terms. Moving ions in a linear array, as in [Rowe 02], is the simplest protocol to consider and a good place to start. In the linear Paul trap, a nodal line of (ideally) zero pseudopotential exists along the axis of the trap. The goal is to adjust the control electrodes in a way that the ions move along a trajectory that follows this nodal line, from one region of the trap to another.

There are two different approaches to ion transport: non-adiabatic and adiabatic transport. Non-adiabatic transport involves accelerating the ion quickly relative to the ion's harmonic motion. At the beginning of the transport, the potential is abruptly changed to impart a kick to the ion. When the ion reaches the final destination, the trap potentials are switched in such a way as to catch the ion and, at the same time, remove all of the motional excitation that was imparted during the transport. Adiabatic transport, on the other hand, attempts to move the ion slowly to minimize the amount of motional excitation in the first place. The potential is gradually shifted along the trajectory so that the ion only sees adiabatic changes in its local environment.

An analogy I find useful is to consider soup in a bowl. If the goal is to transport the soup without starting it sloshing in the bowl, the adiabatic approach would be to move the bowl slowly so that the sloshing motion is never excited. The non-adiabatic approach would be to jerk the bowl quickly over the transport distance, but, at the end, rock the bowl in just the correct way to remove the sloshing motion.

The advantage of non-adiabatic transport is clear: the ions can be moved much faster, potentially in a time shorter than the gate duration ($\sim 10 \mu\text{s}$). This would significantly reduce the time cost of moving ions during algorithms.

The issue with non-adiabatic transport is that it is challenging to realize. Some theoretical work on such ion transport has been performed [Schulz 06, Hucul 08] and experimental results have been reported [Rowe 02, Hensinger 06, Huber 08]. In these demonstrations, the secular motion of the ion was excited to an undesirable level. Non-adiabatic transport was demonstrated using a collection of neutral atoms that were moved in a single harmonic oscillation of the center-of-mass motion without any excitation of that motion (at 8.1 MHz) [Couvert 08]. However, the relatively higher motional frequencies ($\sim 1 \text{ MHz}$) and smaller wave packet size ($\sim 20 \text{ nm}$) for trapped ions makes it more difficult to achieve the necessary control to perform non-adiabatic transport.

Realizing control over the distance scale of the $\sim 20 \text{ nm}$ wave packet is not as important in a simple linear trap array, where the harmonic state is unlikely to evolve significantly during the transport. However, transporting through the complicated pseudopotential landscape near a

junction promises to be much more complicated. For example, in regions near the X-junction, the ion’s micromotion has up to 5 μm amplitude. Such micromotion could easily introduce secular energy, in a non-trivial manner, as the ion passes non-adiabatically through this region. If there is any nonlinear evolution of the ion’s motional state, this unwanted energy could be very difficult to remove using the potentials alone. Thus, it is in these types of regions where very precise control over the spatial form of the potential would be needed, and would be difficult to achieve.

In addition, switching the potentials quickly compared to 1 MHz can also be challenging. The digital-to-analog converter (DAC) system used with the X-junction array had a maximum update rate of 500 kHz and was not capable of such tasks. Faster DAC technology does exist but is quite expensive, especially when considering the number channels needed for a large-scale array (the X-junction array used 40 DAC channels). The problem of moving ions quickly is exacerbated by RC low-pass filters (see Ch. 6) which are needed to suppress electronic noise at the secular frequency and to short the control electrodes to ground at the rf frequency (83 MHz) to ensure proper trapping. The RC filters used in the X-junction array have a low-pass corner frequency of 160 kHz. Any attempt to drive control potentials non-adiabatically must occur at, or well above, this corner frequency and would be significantly attenuated by the filters. Given the technological limitations of today’s DACs and filtering solutions, non-adiabatic transport seems intimidating, especially in complicated trap arrays. Therefore, this dissertation will limit consideration to adiabatic transport.

In adiabatic transport, the control electrodes are used to produce a harmonic trapping potential along the trap axis. The electrode voltages are adjusted to shift the minimum of this potential along the trap axis until the harmonic well arrives at the final destination. Defining ‘adiabatic’ in this context is tricky. Perhaps the best practical definition is simply, “slow enough so that the harmonic motion is not excited.” A more quantitative requirement for adiabaticity in a harmonic well with constant frequency ω whose minimum is accelerating at $a(t)$ can be found using time-dependent perturbation theory [Hucul 06],

$$\omega \gg \frac{1}{a(t)} \frac{d}{dt} a(t). \quad (5.1)$$

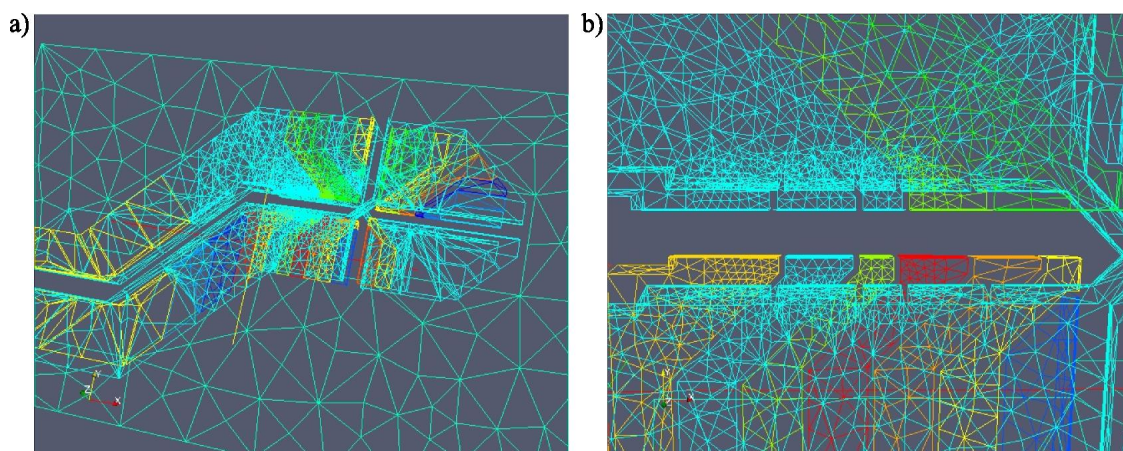


Figure 5.1: This is a sample mesh of the electrode surfaces for calculating the trap potentials using the BEM method. This mesh was used when modeling the region near the experimental zone of the X-junction array. Each electrode was designated a certain color (though some colors were repeated). a) The entire trap, including the bias electrode, is meshed. b) A zoom in on the experimental zone. See Appendix B for a description of meshing and BEM.

The series of control potentials that create the moving harmonic well is referred to as the ‘waveform’. Determining these waveforms can be broken down into four steps: modeling the trap, determining the constraints, solving for the appropriate voltages, and assigning the timing. Much of the remainder of this chapter will be devoted to exactly how the waveforms are determined for experiments in the X-junction array.

5.2 Modeling of the trap potentials

Before waveforms are constructed and ions manipulated, an accurate model of the trap potentials is needed. There are a number of possible methods for generating models of the trap potentials and several are briefly described in Appendix B. The boundary element method (BEM) [Sadiku 09] was the preferred method for modeling the X-junction array.

BEM modeling software, in this case [CPO], was used to calculate the potentials created by the various electrodes in the vicinity of the trap axis. Given the size and complexity of the X-junction array, the trap was broken into several models for different spatial regions to allow for reasonable computations. Each model incorporated a section of the trap axis that ranged from 400

to 600 μm . Every model included the entire trap geometry, but the parts of the trap that were closer to the given section of axis were ‘meshed’ (see Appendix B) with a finer mesh to improve the accuracy. A sample mesh is shown in Fig. 5.1. A three-dimensional rectangular grid of points was defined, centered on the axis and extending the length of the particular section of axis. The grid points were spaced by 5 μm and generally extended eight points (40 μm) to each side of the axis. The BEM code calculated the potential at each grid point.

Instead of calculating the potential due to voltages being applied to all electrodes simultaneously, the BEM model was used to find the potential due to a single electrode held at 1 V while all other electrodes were grounded. This was then repeated for all N electrodes, producing a series of N arrays $\check{\phi}_N(i, j, k)$, where (i, j, k) is the (x, y, z) identifier for each grid point. To find the potential due to multiple electrodes, these different arrays can be weighted by the voltage, V_N , being applied to the given electrode N . Then, due to the linear nature of electric potentials, the weighted arrays can simply be summed to find the total potential. A mathematical manipulation program, in this case [MATLAB 07], can be used to easily calculate these sums.

Though the potentials are determined throughout some grid, interpolation can be used to find the potential at any location. It will be easier to consider a continuous function rather than an array, so, at this point, I will switch to the notation $\phi_N(\mathbf{r})$.

The rf potential ϕ_{rf} , which is calculated as a *static* potential at 1 V in the same way the other potentials are calculated, must be treated with an additional step. Instead of directly using this potential in the waveform calculations, it is converted into a time-independent pseudopotential using

$$\phi_{\text{ps}}(\mathbf{r}) = \frac{q}{4m\Omega_{\text{rf}}^2} (V_{\text{rf}}\nabla\phi_{\text{rf}}(\mathbf{r}))^2, \quad (5.2)$$

where V_{rf} is the peak voltage applied to the rf electrode. This transforms the rf contribution into a pseudopotential that can be directly added to the control potentials ϕ_N to determine the total potential seen by the ion. In this dissertation, the ϕ potentials (including ϕ_{ps}) are reported as electric potentials (in units of V) and *not* energy potentials (units of eV), though converting from

one to another is as simple as multiplying by q .

The total potential when each electrode is held at voltage V_N and the pseudopotential is included is

$$\Phi(\mathbf{r}) = \phi_{\text{ps}}(\mathbf{r}) + \sum_N V_N \phi_N(\mathbf{r}). \quad (5.3)$$

Using $\Phi(\mathbf{r})$, the behavior of an ion in a potential defined by V_N and V_{rf} can be determined, including such things as secular frequencies, principal axes, and stability parameters.

These computer models seem reasonably accurate and predict experimentally measured secular frequencies within a few percent (in one example, the model anticipated an axial frequency of 3.6 MHz and the measured value was 3.62 MHz). The errors were likely due to incomplete knowledge of the trap electrode misalignment (estimated to be 5%) and not due to numerical errors in the model, which were determined to be small by trying a smaller mesh and observing little change in the solution.

5.3 Solving for a single harmonic potential

Using the modeled potentials $\phi_N(\mathbf{r})$, it is possible to determine a set of voltages, $\{V_N\}$, that will produce a desired harmonic potential. The characteristics of the desired harmonic well define a set of constraints. These constraints are then used to construct a matrix equation whose solution yields a set of $\{V_N\}$ that satisfies the constraints.

The trick to finding good solutions and waveforms is determining the best set of constraints to use. These constraints can be broken down into three categories:

- (1) **Physical constraints** - defined by the laws of physics.
- (2) **Geometry constraints** - defined by the trap geometry.
- (3) **User constraints** - defined by the user to produce the desired harmonic potential. This category includes practical limits on available voltages V_N .

In practice, there can be complicated interplay between these three different types of constraints which gives rise to unintended problems. Therefore, it is important to understand all the constraints

and how they complement or contradict each other. However, first we will discuss the mathematical process of solving for $\{V_N\}$. For this purpose, it will be helpful to ignore the first two types of constraints and focus on a specific subset of user constraints. By doing so, we are stripping out the details that are specific to trapped ions or a certain trap geometry, and instead, focusing on the more general idea of a constrained harmonic potential. Once the calculation process has been explained, we will return to the topic of constraints to examine it more carefully. At that point, the physical and geometric constraints specific to trapped ions will be introduced, as will a more appropriate set of user constraints.

5.3.1 Singular value decomposition and the pseudoinverse

A three-dimensional harmonic potential can be completely defined by nine attributes: three position attributes, three harmonic frequencies, and three angles that define the orientation of principal axes. Assume our goal is to produce a harmonic potential located at $\mathbf{r}_0 = (x_0, y_0, z_0)$ with frequencies ω_x , ω_y and ω_z and with principal axes that lie along \hat{x} , \hat{y} and \hat{z} . Since this is an example only, for the time being the directions \hat{x} , \hat{y} and \hat{z} are defined relative to an arbitrary coordinate system and not the coordinate system of a specific trap. Likewise, \mathbf{r}_0 can be any location and we are not currently concerned with placing it in a location ideal for ion-trapping, such as the pseudopotential minimum. (Later, when we return to the topic of geometry constraints, the direction of the coordinate system will be defined relative to the trap and \mathbf{r}_0 will be chosen to be in a reasonable location.) To solve for this desired harmonic potential, nine constraints, corresponding to the nine attributes, must be represented mathematically in a manner that is useful for calculating $\{V_N\}$.¹

The three position constraints imply \mathbf{r}_0 is a local minimum of the potential. At this point, the slope of Φ must be zero, which allows the position constraint to be written as

$$\nabla\Phi(\mathbf{r}_0) \doteq 0. \tag{5.4}$$

¹ These nine constraints will help explain the mathematical procedure used to solve for waveforms. In practice, there are better sets of constraints and those will be covered later.

The symbol \doteq is used to distinguish this as being constrained to be equal rather than being implicitly equal. If the three frequencies are greater than zero, then this constraint forces \mathbf{r}_0 to be a local minimum. This single equation actually defines three separate constraints: one for the slope in each of three directions.

For the frequency constraints, it is useful to first consider a one-dimensional harmonic oscillator taken to be along the z axis. The frequency is constrained by defining the following relationship between that frequency ω_z and the second derivative (curvature) of the potential at the local minimum:

$$\frac{d^2}{dz^2}\Phi(z_0) \doteq \frac{m}{q}\omega_z^2. \quad (5.5)$$

For the three-dimensional harmonic oscillator the frequencies and the directions of oscillation (the principal axes) cannot necessarily be considered separately. However, they can be considered collectively by using the Hessian matrix (Eq. (3.6)),

$$\mathcal{H}(\mathbf{r}_0) \equiv q \begin{bmatrix} \frac{\partial^2}{\partial x^2} & \frac{\partial^2}{\partial x \partial y} & \frac{\partial^2}{\partial x \partial z} \\ \frac{\partial^2}{\partial y \partial x} & \frac{\partial^2}{\partial y^2} & \frac{\partial^2}{\partial y \partial z} \\ \frac{\partial^2}{\partial z \partial x} & \frac{\partial^2}{\partial z \partial y} & \frac{\partial^2}{\partial z^2} \end{bmatrix} \Phi(\mathbf{r}_0). \quad (5.6)$$

As stated in Sec. 3.1, the harmonic frequencies of the potential will be given by the eigenvalues λ_i of the Hessian,

$$\lambda_i = m\omega_i^2. \quad (5.7)$$

Furthermore, the principal axes of the three-dimensional harmonic potential lie along the eigenvectors of the Hessian. Thus the Hessian can be used to constrain the frequencies and principal axes. However, it may not be simple to separate the constraints on the frequencies from the constraints on the axes when using the Hessian.

For now, consider a special case where the principal axes coincide with the directions in which the derivatives of the Hessian are taken. In this case, the Hessian will be diagonal and the

frequency and direction constraints can be written as

$$\mathcal{H}(\mathbf{r}_0) \doteq m \begin{bmatrix} \omega_x^2 & 0 & 0 \\ 0 & \omega_y^2 & 0 \\ 0 & 0 & \omega_z^2 \end{bmatrix}. \quad (5.8)$$

This is actually nine simultaneous constraint equations on all of the second derivatives of the potential written as a single matrix equation. Since the derivatives of the potential commute, the Hessian will be symmetric ($\mathcal{H} = \mathcal{H}^T$) and three of the nine constraints equations are duplicates. This leaves six linearly independent constraints. In this special case of a diagonal Hessian, the frequency and direction constraints are separable. The secular frequencies are constrained by the three equations corresponding to the diagonal of Eq. (5.8). The off-diagonal constraints force the cross-derivatives to be zero and ensure the principal axes align to the $(\hat{x}, \hat{y}, \hat{z})$ directions in which the Hessian is calculated.

Together Eqs. (5.4) and (5.8) constitute the set of nine constraints that completely define the harmonic potential. All of these constraints have been written in a form that depends on the first or second derivatives of $\Phi(\mathbf{r}_0)$. And, as we soon be seen, it is possible to combine these into a formalism that can be used to solve for $\{V_N\}$.

In the interest of compact nomenclature, it is convenient to define several vectors and matrices. The set of $\{V_N\}$ can be written as a vector

$$\mathbf{V}^T \equiv \begin{bmatrix} V_1 & V_2 & \dots & V_N \end{bmatrix}, \quad (5.9)$$

where \mathbf{V}^T is \mathbf{V} transposed. The potentials at location \mathbf{r}_0 , contributed by each control electrode, will also be written as a vector:

$$\Psi^T(\mathbf{r}_0) \equiv \begin{bmatrix} \phi_1(\mathbf{r}_0) & \phi_2(\mathbf{r}_0) & \dots & \phi_N(\mathbf{r}_0) \end{bmatrix}. \quad (5.10)$$

In addition, we would like to include a potential contribution from a non-scalable or fixed potential, such as the pseudopotential ϕ_{ps} . There is no fundamental reason why ϕ_{ps} cannot also be actively adjusted during the course of a waveform. However, adjusting V_{rf} in the experimental

apparatus is inconvenient and, unless there is a compelling reason to do so, it is typically held constant. In this dissertation, ϕ_{ps} was fixed to a value that produced ~ 12 MHz radial secular frequencies.

To incorporate a fixed potential, another set of vectors is defined:

$$\mathbf{V}_F \equiv [1], \quad (5.11)$$

$$\Psi_F(\mathbf{r}_0) \equiv [\phi_{\text{ps}}(\mathbf{r}_0)], \quad (5.12)$$

where the subscript F implies that these vectors are fixed and will not be adjusted to fulfill constraints. Setting $\mathbf{V}_F = [1]$ results in the fixed potential always being applied at full magnitude.² In addition to pseudopotential, other fixed potentials could be included into the problem by placing them in Eq. (5.12) and increasing the size of Eq. (5.11) accordingly.

Finally, define the operator

$$\mathcal{P}^T \equiv \left[1 \quad \begin{array}{c} \vdots \\ \frac{\partial}{\partial x} \quad \frac{\partial}{\partial y} \quad \frac{\partial}{\partial z} \\ \vdots \end{array} \quad \begin{array}{c} \frac{\partial^2}{\partial x^2} \quad \frac{\partial^2}{\partial x \partial y} \quad \frac{\partial^2}{\partial x \partial z} \\ \frac{\partial^2}{\partial y \partial x} \quad \frac{\partial^2}{\partial y^2} \quad \frac{\partial^2}{\partial y \partial z} \\ \frac{\partial^2}{\partial z \partial x} \quad \frac{\partial^2}{\partial z \partial y} \quad \frac{\partial^2}{\partial z^2} \end{array} \right]. \quad (5.13)$$

The dotted lines are intended as a guide to the reader to visually separate the different components. When acting upon a single potential, such as $\phi_N(\mathbf{r}_0)$, this operator returns a column vector with 13 entries; the first entry is the potential, the next three entries are the slope of the potential, and the final nine entries are related to the various components of the Hessian.

Using the above definitions and the constraints, the following matrix equation can be constructed

$$\mathbf{C}_1(\mathcal{P} \otimes \Psi^T(\mathbf{r}_0))\mathbf{V} + \mathbf{C}_1(\mathcal{P} \otimes \Psi_F^T(\mathbf{r}_0))\mathbf{V}_F \doteq \mathbf{C}_2. \quad (5.14)$$

Here \mathbf{C}_1 is a $j \times 13$ matrix and \mathbf{C}_2 is a $j \times 1$ vector, and together they contain the information about the constraints. The parameter j is the number of constraints placed on the system, in this case $j = 9$. The $(\mathcal{P} \otimes \Psi^T(\mathbf{r}_0))\mathbf{V}$ portion of the equation can be adjusted with \mathbf{V} to satisfy the constraints.

² Since ϕ_{ps} is already scaled by the rf voltage V_{rf} , as defined in Eq. (5.2), ϕ_{ps} is not scaled by V_{rf} in the same way that V_N is used to scale ϕ_N .

To see how the nine constraints are encoded into \mathbf{C}_1 and \mathbf{C}_2 , first consider the position constraints. The second, third, and fourth entries in $\mathcal{P}\Phi(\mathbf{r}_0)$ are the three components of the gradient in Eq. (5.4). If the constraint matrices are chosen to be

$$\mathbf{C}_1 = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (5.15)$$

and

$$\mathbf{C}_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \quad (5.16)$$

\mathbf{C}_1 will pick out the spatial derivatives while \mathbf{C}_2 constrains them to be zero. In this case, Eq. (5.14)

becomes

$$\begin{bmatrix} V_1 \frac{\partial}{\partial x} \phi_1(\mathbf{r}_0) + \cdots + V_N \frac{\partial}{\partial x} \phi_N(\mathbf{r}_0) + \frac{\partial}{\partial x} \phi_{\text{ps}}(\mathbf{r}_0) \\ V_1 \frac{\partial}{\partial y} \phi_1(\mathbf{r}_0) + \cdots + V_N \frac{\partial}{\partial y} \phi_N(\mathbf{r}_0) + \frac{\partial}{\partial y} \phi_{\text{ps}}(\mathbf{r}_0) \\ V_1 \frac{\partial}{\partial z} \phi_1(\mathbf{r}_0) + \cdots + V_N \frac{\partial}{\partial z} \phi_N(\mathbf{r}_0) + \frac{\partial}{\partial z} \phi_{\text{ps}}(\mathbf{r}_0) \end{bmatrix} \doteq \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \quad (5.17)$$

which is exactly Eq. (5.4). (Through setting the appropriate entry in \mathbf{C}_2 to a non-zero value, one or more of these derivatives can be set to a different value, if the user had the desire to do so.)³

Appending additional rows to \mathbf{C}_1 and \mathbf{C}_2 will add more constraints to Eq. (5.14).

The frequency and principal axes constraints in Eq. (5.8) are easy to implement and involve the last nine entries of $\mathcal{P}\Phi(\mathbf{r}_0)$. The constraint matrices that incorporate these six constraints are

$$\mathbf{C}_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{C}_2 = \begin{bmatrix} (m/q)\omega_x^2 \\ (m/q)\omega_y^2 \\ (m/q)\omega_z^2 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (5.18)$$

³ Though possible to implement, such non-zero constraints were never used in this dissertation as they appear to be of no practical value for the problem addressed here.

The first three rows in Eq. (5.18) correspond to the diagonal constraints in Eq. (5.8) while the next three rows correspond to the upper-half off-diagonal constraints and will constrain the Hessian to be diagonal. Thus, this exactly reproduces all six independent frequency and principal axes constraints of Eq. (5.8) with $q\mathbf{C}_1$ replacing $\mathcal{H}(\mathbf{r}_0)$ and $q\mathbf{C}_2$ replacing the right-hand-side of Eq. (5.8).

By appending Eq. (5.18) to Eqs. (5.15) and (5.16), all nine constraints needed to define a three-dimensional harmonic oscillator will be contained within the constraint matrices and can be encoded into Eq. (5.14).

From Eq. (5.14), the solution for \mathbf{V} that satisfies all of these constraints is

$$\mathbf{V} = \mathbf{G}^{-1}(\mathbf{C}_2 - \mathbf{D}), \quad (5.19)$$

where $\mathbf{G} \equiv \mathbf{C}_1\mathcal{P} \otimes \Psi^T$ and $\mathbf{D} \equiv \mathbf{C}_1\mathcal{P} \otimes \Psi_F^T\mathbf{V}_F$. To find \mathbf{V} , \mathbf{G} must be inverted, which may not be strictly possible. The problem is that the dimension of \mathbf{G} depends both on the number of electrodes and the number of constraints used. There is no guarantee that \mathbf{G} is square; in fact, it typically is not. If the number of electrodes exceeds the number of constraints, then there will not be a unique solution. If, on the other hand, the number of constraints exceeds the number of electrodes, the problem will be over-constrained and there may be no solution at all.

A mathematical method for solving this problem, even when the problem is under- or over-constrained, exists and uses singular value decomposition (SVD) [Golub 65]. SVD is an important technique in linear algebra for factorizing matrices and has many applications. We will use it here for finding the ‘pseudoinverse’ of \mathbf{G} . The SVD method allows the matrix \mathbf{G} , which is an $m \times n$ matrix, to be decomposed into three matrices:

$$\mathbf{G} = \mathbf{J}\mathbf{S}\mathbf{K}^\dagger. \quad (5.20)$$

Here, \mathbf{J} and \mathbf{K} are unitary matrices with dimension $m \times m$ and $n \times n$, respectively. \mathbf{S} is an $m \times n$ rectangular diagonal matrix with only non-negative values along the diagonal. Those values are called the ‘singular values’ of \mathbf{G} . The SVD is similar to an eigenvalue decomposition. In fact, in cases where \mathbf{G} is invertible and has all positive eigenvalues, the two decompositions are identical.

However, SVD has the advantage that it can be performed on any matrix. The matrix need not be square nor invertible (both of which are prerequisites for the eigenvalue decomposition).

When solving Eq. (5.19), if \mathbf{G} is not invertible, obviously finding the true inverse \mathbf{G}^{-1} is not possible. However, the SVD technique allows one to calculate the Moore-Penrose pseudoinverse [Moore 20, Penrose 55], which is denoted as \mathbf{G}^+ . The pseudoinverse satisfies $\mathbf{G}\mathbf{G}^+\mathbf{G} = \mathbf{G}$. This does not require $\mathbf{G}\mathbf{G}^+$ to be the identity matrix, but it does require that $\mathbf{G}\mathbf{G}^+$ maps all column vectors of \mathbf{G} onto themselves. The pseudoinverse can be found from the SVD decomposition as follows:

$$\mathbf{G}^+ = \mathbf{K}\mathbf{S}^+\mathbf{J}^\dagger. \quad (5.21)$$

Here, \mathbf{S}^+ is also a pseudoinverse, but because \mathbf{S} is diagonal, the pseudoinverse can be trivially found by replacing every non-zero entry in \mathbf{S} with its reciprocal and then transposing the matrix. The solution for \mathbf{V} is then, more accurately,

$$\mathbf{V} = \mathbf{G}^+(\mathbf{C}_2 - \mathbf{D}). \quad (5.22)$$

In cases where there are more independent constraints than there are electrode variables, the pseudoinverse in Eq. (5.22) will automatically find a least-squares solution that minimizes the deviation between the solution and all the constraints collectively.

In other cases where there are more electrode variables than constraints, the pseudoinverse will pick a single \mathbf{V} , though multiple solutions exist. To access the other solutions, the SVD technique can also be used to find the null space $\tilde{\mathbf{V}}$ of the Eq. (5.14). $\tilde{\mathbf{V}}$ is the vector space where any vector, $\tilde{\mathbf{V}}^i$, in that space satisfies the homogeneous equation $\mathbf{C}_1(\mathcal{P} \otimes \Psi^\top)\tilde{\mathbf{V}}^i = 0$. Thus, any superposition of the vectors in $\tilde{\mathbf{V}}$, when added to \mathbf{V} , will still satisfy Eq. (5.14). The null space $\tilde{\mathbf{V}}$ can be written as a matrix where each column is a linearly independent vector in the null space. The rank of $\tilde{\mathbf{V}}$ is equal to the size of \mathbf{V} minus the number of independent constraints.

There is a practical issue that can arise when the problem is under-constrained. Since any vector of the form $\mathbf{V} + \alpha\tilde{\mathbf{V}}^i$ (where α is a scalar) satisfies Eq. (5.14), there is no limit to how large individual $|V_N|$'s can be. Furthermore, there is nothing in the SVD technique that prefers solutions

with small voltages. If the voltages become too large, they cannot be realized in the apparatus, so it important to use small voltages.

The null space $\tilde{\mathbf{V}}$ can be used to find the solution \mathbf{V}' which has the smallest Euclidean norm $\|\mathbf{V}'\|_2 \equiv \sqrt{\sum_n (V'_n)^2}$,

$$\mathbf{V}' = \mathbf{V} - \tilde{\mathbf{V}}\tilde{\mathbf{V}}^T\mathbf{V}. \quad (5.23)$$

Finding $\|\mathbf{V}'\|_2$ places a preference against solutions that involve excessively large voltages, but does not necessarily guarantee that every $|V'_N|$ will be small.

This section has described the process of solving for \mathbf{V}' by constructing a matrix equation that incorporates nine constraints. These constraints define the desired harmonic potential located at \mathbf{r}_0 with frequencies ω_x , ω_y and ω_z and with principal axes that lie along \hat{x} , \hat{y} and \hat{z} . It is instructive to ask, ‘what is the effect of deleting one of the constraints to leave just eight?’

Deleting one of the position constraints from Eqs. (5.15) and (5.16) can have a significant. In this case, the position of the local minimum is no longer constrained to be at \mathbf{r}_0 and could move somewhere else entirely (unless there is a physical or geometry constraint that prevents this). If the minimum moves, then the frequency and principal axes constraints may not be relevant at the actual location of the ion, as they are enforced at \mathbf{r}_0 . Thus, it is not advisable to remove a position constraint.

If a frequency constraint from the first three rows of Eq. (5.18) is removed, the only impact on the solution is that the corresponding frequency will be unconstrained. This would not affect the rest of the constraints (because \mathcal{H} is still constrained to be diagonal, making the frequency constraints independent).

Finally, assume the last row of Eq. (5.18) is removed. This deletes the constraint $\frac{\partial^2}{\partial y \partial z} \Phi(\mathbf{r}_0) \doteq 0$ and allows \mathcal{H} to be non-diagonal. The two principal axes that were constrained to lie along \hat{y} and \hat{z} are now free to align anywhere in the y/z plane (though they still must be orthogonal). However, the second and third rows of Eq. (5.18) will not rotate with the axes and will constrain curvatures that no longer correspond to independent harmonic frequencies. The frequency constraints are

dependent on the principal axes constraints. (Appendix C suggests some alternate methods for expressing the frequency and axes constraints so that they are more independent.)

Now that the mathematics has been explained, it is useful to briefly consider how practical these nine constraints are in the case of a trapped ion. Though the position \mathbf{r}_0 can be arbitrarily chosen, it is best to have the ion located at the micromotion minimum (the pseudopotential minimum), which is usually on the axis of the trap. If that axis is along z , then the user should select $\mathbf{r}_0 = (x = 0, y = 0, z = z_0)$. In addition, the user should choose secular frequencies that are useful. In the X-junction array, the typical parameters are $\omega_z = 2\pi \times 3.6$ MHz and $\omega_x \approx \omega_y \approx 2\pi \times 13$ MHz. Finally, the principal axes should be oriented in a direction that is amenable to laser cooling. So far the axes were chosen to lie along the direction of the derivatives in \mathcal{H} . It is possible to define \mathcal{H} in a direction that is good for laser cooling, but in general \mathcal{H} will be defined along the direction of the computer model. Appendix C includes a method for using the above formalism with radial principal axes that are rotated to an arbitrary angle relative to \mathcal{H} .

In the upcoming sections, we will see that when the physical and geometry constraints are included, the above set of nine constraints is not the optimal set of constraints and it is better to omit some of them.

5.3.2 Constraints

The formalism for solving for \mathbf{V} is actually quite flexible. There are a number of possible constraints that can be implemented through simple modifications to the above procedure. At the same time, there are a number of pitfalls and caveats that must be kept in mind. One must have a good understanding of the constraints being applied and how they will influence the solutions. This section will consider some other constraints, both implicit and explicit, that can be applied to the Eq. (5.14).

Physical Constraints

As alluded to earlier, the physical laws that govern the ion system place implicit constraints on Eq. (5.14). It is important to realize how these constraints influence the problem. Most of the

physical constraints that are encountered can be traced back to either Laplace's equation or the nature of a harmonic potential, and can be found in one form or another in Ch. 3.

One physical constraint that was already mentioned was the symmetry of the Hessian, $\mathcal{H} = \mathcal{H}^T$. There is nothing preventing the user from defining a constraint that would violate this requirement. However, if this is done, Eq. (5.14) will return a bad solution that does not fulfill the user's requirements.

The three principal axes of the harmonic potential are physically constrained to be orthogonal. Again, this comes from the Hessian being diagonalizable. This can also be viewed as arising from the fact that a quadratic potential can be expressed as uncoupled quadratic terms along orthogonal directions.

The physics also constrains the portion of the secular frequencies that is due to the control potentials (and not the pseudopotential) to fulfill the following relationship:

$$\tilde{\omega}_x^2 + \tilde{\omega}_y^2 + \tilde{\omega}_z^2 = 0. \quad (5.24)$$

This is seen in Eq. (3.19) and is a direct result of Laplace's equation. When the pseudopotential is included, as in Eq. (3.20), the total secular frequencies become $\omega_i = \sqrt{\tilde{\omega}_i^2 + \omega_{\text{rf}}^2}$, where ω_{rf} is the pseudopotential contribution defined in Eq. (3.17). For a *linear* Paul trap, this implies that the physics of the system forces the following constraint on the three secular frequencies

$$\omega_x^2 + \omega_y^2 + \omega_z^2 = 2\omega_{\text{rf}}^2. \quad (5.25)$$

Recall that ω_{rf} is considered to be fixed in this problem. If the user constrains two of the secular frequencies, the third secular frequency will also be constrained. In the previous section, three user constraints were placed on the secular frequencies. If those user constraints do not coincide with this physical constraint (and even, to some extent, if they do) the solution for \mathbf{V} may be compromised, returning a voltage set that does not do what was intended. Thus the three secular frequencies cannot be arbitrarily set.

Geometry Constraints

In addition to the physics of the system, the geometry of the trap will introduce a series of implicit constraints.

For example, the trap geometry completely defines the structure (though not necessarily the amplitude) of the pseudopotential. The location of the pseudopotential nodal line is fixed and cannot be adjusted without changing the trap geometry (or allowing different rf amplitudes on the rf electrodes, which is a topic that will not be covered here but is discussed in [Cetina 07, Herskind 09, VanDevender 10]).

The trap geometry also places constraints on the principal axes. For the Paul trap in Fig. 3.1, the perfect four-fold symmetry requires that the radial principal axes lie along the diagonal of the trap and point toward the electrodes, regardless of what static potentials are applied to the control electrodes (unless the two motions are degenerate, in which case the principal axes are poorly defined). Thus the user does not have complete freedom over the orientation of the principal axes. The two-layer wafer trap does not have this exact four-fold symmetry and the above argument does not strictly hold. Still, to get the radial principal axes to rotate far from the trap diagonals can require excessively large control potentials depending on how close the trap geometry is to having the four-fold symmetry.

Geometric constraints are something that must be considered during the trap design process. Once the trap is built, if one of these constraints conflicts with the goal of an experiment, there may be no recourse other than building a new trap.

User Constraints

The rest of this section will be devoted to the user constraints that can be introduced into Eq. (5.14). Using the formalism in the previous section, there are many possible constraints that can be applied by adding rows to \mathbf{C}_1 and \mathbf{C}_2 . Appendix C contains explanations of some of these constraints. Here the focus will be on what constraints were actually used rather than what constraints are possible.

In the preceding section, three types of user constraints were used. The position constraints

in Eqs. (5.15) and (5.16) are very important when designing waveforms. Thus, these three position constraints are almost always used in practice. There may be some circumstances where other considerations may be more important than the position of the trapping minimum, and in those cases, perhaps one or more of the position constraints could be omitted (by omitting the appropriate lines in Eqs. (5.15) and (5.16)). However, for the waveforms used in the X-junction array, these three position constraints were always present.

There were also three constraints placed on the secular frequencies and another three placed on the orientation of the principal axes using Eq. (5.18). In practice, this set of six constraints is too rigid, especially because physical and geometry constraints also implicitly define some of these attributes. This can lead to a poorly-conditioned matrix problem where the solutions become poor and dominated by numerical errors. It may be desirable, or even necessary, to reduce these six constraints to a more flexible set of constraints. Furthermore, the principal axes defined by Eq. (5.18) may not be convenient for practical purposes (e.g., laser cooling). Some methods for increasing the flexibility of these constraints are mentioned in Appendix C. However, the usual approach is to simply drop some of the constraints, as many of them are not needed for the success of the experiment.

For example, quantum computing operations are usually performed on the axial mode of motion, making axial constraints much more important than the radial constraints. In many circumstances, the radial mode frequencies and directions will not need to be explicitly constrained. In this situation, all that is required of the radial modes is that they not be degenerate with the axial mode and that the axes are not perpendicular to the Doppler cooling beam. Here is where the geometry and physical constraints can be used for our benefit. As stated earlier, geometric constraints prefer (or demand, in the case of the four-fold-symmetric trap in Fig. 3.1) radial principal axes that point towards the electrodes. Since the Doppler beams in such a trap usually propagate along a different direction, this geometric constraint makes it difficult to realize a principal axis orientation that would lead to problems. Thus the user constraint that explicitly fixed the orientation of the radial principal axes can be omitted and the implicit geometry constraint will be sufficient.

Similarly, the constraints on ω_x and ω_y may not be needed. Most experiments are designed to have a much lower axial frequency than the radial frequencies. As long as the rf is set so that $\omega_{\text{rf}}^2 \gg \omega_z^2$, the physical constraint in Eq. (5.25) will often force $\omega_x, \omega_y \gg \omega_z$ and $\omega_x \neq \omega_y$. This is especially true when the solution \mathbf{V}' , found using Eq. (5.23), is used because its small voltages are less likely to create large discrepancies between ω_x^2 and ω_y^2 . This physical constraint is usually sufficient and we rarely need to explicitly constrain the radial secular frequencies.

Thus, instead of using the combined constraint matrices in Eqs. (5.15), (5.16), and (5.18), it is often sufficient (and preferable) to use the reduced constraint matrices

$$\mathbf{C}_1 = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{C}_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ (m/q)\omega_z^2 \\ 0 \\ 0 \end{bmatrix}, \quad (5.26)$$

where the first three rows constrain the position, the fourth row constrains the axial secular frequency. The last two rows pick out two of the cross-derivatives and set them to zero: $\frac{\partial^2}{\partial x \partial z} \Phi \doteq 0$ and $\frac{\partial^2}{\partial y \partial z} \Phi \doteq 0$. This prevents the principal axis along \hat{z} from pointing in any other direction and ensures that the fourth row is constraining the frequency in just the \hat{z} direction.⁴

A second variety of user constraint that is helpful is to fix one or more of the electrodes to a predetermined value. For example, it can be convenient to set the control electrodes at the very ends of a linear array to some high voltage (say 10 V) to ensure no ion will ever escape out the ends of the trap (even when hot). This can also be used to reduce the number of free parameters in the problem. Such a constraint cannot be achieved through altering \mathbf{C}_1 and \mathbf{C}_2 . Instead $\Psi^T \mathbf{V}$ and $\Psi_F^T \mathbf{V}_F$ must be directly adjusted. If V_2 is to be held at 10 V, this is achieved by moving the

⁴ Without the last two rows, the principal axes would not necessarily coincide with the axes in which the derivatives were taken and the fourth row would no longer independently constrain a single secular frequency.

electrode 2 dependence from the first term of Eq. (5.14) to the second term:

$$\mathbf{C}_1 \mathcal{P} \otimes \begin{bmatrix} \phi_1(\mathbf{r}_0) \\ \phi_3(\mathbf{r}_0) \\ \dots \\ \phi_N(\mathbf{r}_0) \end{bmatrix}^T \begin{bmatrix} V_1 \\ V_3 \\ \dots \\ V_N \end{bmatrix} + \mathbf{C}_1 \mathcal{P} \otimes \begin{bmatrix} \phi_{ps}(\mathbf{r}_0) \\ \phi_2(\mathbf{r}_0) \end{bmatrix}^T \begin{bmatrix} 1 \\ 10 \end{bmatrix} \doteq \mathbf{C}_2. \quad (5.27)$$

Any number of electrodes can be set to fixed voltages in this manner.

It can also be helpful, in certain situations, to constrain the ratio of two control voltages, e.g., $V_2 \doteq \alpha V_3$. This also cannot be achieved by appending rows to \mathbf{C}_1 and \mathbf{C}_2 . In fact, this requires a slight modification to Eq. (5.14) of

$$\mathbf{C}_1 (\mathcal{P} \otimes \Psi^T(\mathbf{r}_0)) \mathbf{M} \mathbf{V} + \mathbf{C}_1 (\mathcal{P} \otimes \Psi_F^T(\mathbf{r}_0)) \mathbf{V}_F = \mathbf{C}_2. \quad (5.28)$$

If no voltage ratios are to be fixed, the matrix \mathbf{M} is the $N \times N$ identity matrix. To introduce a ratio constraint, one of the voltages must be removed from the equation. For the above example, \mathbf{M} becomes

$$\begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & \frac{1}{\alpha} & 0 & 0 & \dots \\ 0 & 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad (5.29)$$

where V_3 is removed from the equation by setting the third column to zeros, while the contribution from the third electrode is added to the second column, with the proper weighting. This technique can be used to set any linear relationship between voltage on any number of control electrodes. However, it is most often used to force control electrodes on one side of a linear Paul trap to be equal those on the other side, so that the voltages are symmetric about the axis. The trap symmetry suggests that is often the optimal solution.

These are the main constraints that can be applied. This formalism provides significant flexibility for constraining the monopole, dipole, and quadrupole terms of Φ . One could imagine

expanding \mathcal{P} to include third order derivatives or higher, which would allow constraints on hexapole or higher-order terms in the potential. However, doing this may create limits to numerical accuracy due to the large number of derivatives involved. Quadrupole terms are sufficient for defining a harmonic potential so \mathcal{P} was only taken out to the second derivatives.

5.3.3 Issues

When using the SVD technique, there are a number of potential problems that can result in bad solution to Eq. (5.14). It is important to examine these issues.

First, consider the case when Eq. (5.14) is over-constrained. The SVD technique is still able to find a solution and that solution is a least-squares type solution that minimizes the difference between the solution and the constraints. One problem is that this assumes all constraints are equally important. For example, assume the user constrains both the axial and radial secular frequencies but places significantly more importance on axial constraint. If the problem is over constrained, the SVD solution cannot fulfill all the constraints. The solution that is returned will not take into account the importance of the axial frequency to the user. In fact, having two radial constraints and only one axial constraint dilutes the importance of that axial constraint in the SVD algorithm. The result is the axial frequency will likely deviate farther than the user would want.

It is possible to weight each constraint encoded in \mathbf{C}_1 and \mathbf{C}_2 by its relative importance. This is done by simply multiplying the each row of both matrices by the weight factor for that row. (Constraints such as fixing voltages or voltage ratios cannot be weighted in this manner.) These weights can help emphasize certain constraints over others, but also begs the question ‘what is the exact relative importance of any one constraint over any other?’ The output solution will depend very much on precisely what weights are given to each row. Thus defining a good set of weight factors can be an ambiguous process.

In the X-junction array, there were 46 electrodes and 40 independent voltage variables. Distant electrodes were often fixed to 0 V to reduce the size of the problem, but still the solutions calculated in that trap were always under-constrained rather than over-constrained. Under-constrained

problems suffer from a series of similar issues.

The SVD technique, on an under-constrained problem, will always produce an exact solution which can actually lead to difficulties. Assume, for example, that the user desires to constrain all three secular frequencies and, being aware of the physical constraint in Eq. (5.25), the user has decided to constrain those frequencies such that they fulfill that equation. Thus there are three user constraints and one physical constraint placed on the same three quantities. (This is not over-constrained, because there are more than three voltage variables.) The user might not anticipate an issue because his constraints should not contradict the physical constraint. However, when numerical errors are present, the physical constraint, as expressed within the computational problem, can deviate slightly from its ideal form, which will cause it to disagree with the user-defined constraints. Such numerical errors may arise during the modeling or the differentiation, or can even be due to floating-point errors. In this case, the four constraints, when considered in an orthogonal basis, include three constraints on the frequencies and a fourth, unintended (and non-zero) constraint on the disagreement between the physical and user constraints within the computer model. This last constraint is entirely due to numerical errors and does not represent anything physical in the apparatus. The SVD method can still produce an exact solution that mathematically satisfies all four constraints, but will likely involve extreme control voltages used to fulfill the erroneous constraint. This is an example of a ‘poorly-conditioned’ matrix equation.

Poor conditioning occurs when significantly different scales are present within the problem. In a poorly-conditioned matrix, there will be a large discrepancy in the amplitudes of the singular values, which are calculated when finding the SVD. Tiny singular values then dominate the solution when the pseudoinverse is computed. There are many opportunities for Eq. (5.14) to become poorly-conditioned. A common cause is two ideally-contradicting constraints that can be simultaneously realized due to numerical errors. This can be a consequence of two poorly chosen user constraints. It is more often the result of hidden physical or geometry constraints contradicting a single or combination of user constraints, which can be especially frustrating depending on how well hidden the contradiction is. In the above example, the constraint on the disagreement has a very small

singular value associated with it because it is constraining an extremely small numerical error, much smaller than the scale of the other frequency constraints. Poorly-conditioned matrix equations often lead to divergence of the solution.

Fortunately, there are two convenient methods for fixing a poorly-constrained matrix. The first is to remove the extra user constraints. It is often not advisable to create a user constraint that is not linearly independent of all other constraints. Had the user above only used two frequency constraints, the third secular frequency would likely have been constrained well enough for his purposes and his problem would have been avoided. Allowing the system of equations more flexibility by removing the least important constraints will likely improve the solution.

On the other hand, there are occasions when a user constraint can be used to enforce a physical constraint to produce beneficial results. For example, in some instances, numerical errors led to a violation of the physical constraint $\mathcal{H} = \mathcal{H}^T$. When this occurred, it resulted in bad solutions. This was easily fixed by going against the advice above and introducing a user constraint that duplicated the (failing) physical constraint. In particular, the problem was avoided by introducing the user constraint $\frac{\partial^2}{\partial x \partial z} \Phi \doteq \frac{\partial^2}{\partial z \partial x} \Phi$. Similar constraints on the x/y and/or y/z cross derivatives are also possible but unadvisable: though one of these constraints would occasionally be used to help reinforce the symmetry of the Hessian, using a second was typically too constraining and would lead to poor-conditioning. This was a case when duplicating a physical constraint with a user constraint was permissible, but duplicating more than one caused problems.

Another method of dealing with a poorly-constrained matrix is to examine the singular values of the matrix that will be pseudo-inverted. The S matrix in the decomposition in Eq. (5.20) contains the singular values, which are similar to eigenvalues. If some of the singular values are significantly smaller than the rest of the singular values (say 10^{-6}), they can be set to zero. This will remove that constraint from the equation and will improve the matrix conditioning, along with the quality of the solution. Of course, one must be careful not to remove essential constraints simply because they have a small singular value associated with them. In practice, the singular values of useful constraints tend to be large enough that the user would not accidentally remove them.

There is one important cause of poor-conditioning that deserves extra attention. It is common to use both position and frequency constraints simultaneously to solve for the intended harmonic potential. But these are very different types of constraints and can have very different scales associated with them. The position constraint is a constraint on the $\frac{\partial}{\partial z}\Phi$ and has units of V/m. The frequency constraint is a constraint on $\frac{\partial^2}{\partial z^2}\Phi$ and has units of V/m². If the size scale of the trap is 10⁻⁴ m and the problem is being calculated in meters, the frequency constraints will be weighted by a factor of 10⁴ that comes purely from the choice of units. This makes the frequency constraint much more important and the position constraint much more susceptible to numerical errors. It is preferable to use a distance unit equal to some size scale relevant to the trap (such as the ion-to-electrode distance) rather than meters. In that case, the weighting factor disappears and both constraints are of similar importance. A related source of poor-conditioning occurs when very large (or small) physical constants (e.g., \hbar , m , or q) are included as part of a constraint. This too can be fixed by using units that do not bias one constraint over another.

The X-junction array contains many different types of regions with different geometries. No single set of constraints seemed work well for every region, and the constraints were adjusted depending on where the trapping well was being calculated. However, a certain basic set of constraints was useful in most locations and provided a starting point in the more troublesome areas:

- (1) The position of the potential minimum was constrained in three directions to be at \mathbf{r}_0 .
- (2) One of the principal axes was constrained to lie along the trap axis (which involves two constraints on axes orientation).
- (3) The axial frequency was constrained (usually 3.6 MHz).
- (4) The voltages applied to electrodes on opposite sides of the trap axis were constrained to be equal (unless something broke this symmetry, such as in the bends near the load zone).

Constraints (1), (2), and (3) are the same as given in Eq. (5.26). All of the derivatives were performed in units of $d = 150 \mu\text{m}$, which was the distance from the ion to the nearest electrode.

This relatively sparse set of constraints tended to give good solutions at most locations considered.

5.4 Solving for transport waveforms

Producing a transport waveform is not conceptually difficult once the process of solving for a harmonic potential has been perfected. The waveform is simply a series of these harmonic potentials, calculated so that the minimum moves along a defined path, and then combined sequentially in time to produce a moving harmonic well. However, there are a number of issues that arise when working with a series of solutions to Eq. (5.14) and attempting to fit the different solutions together in a smooth manner. This section will examine these issues.

The waveform technique used in the X-junction array differs slightly from what was used in previous NIST traps [Rowe 02, Barrett 04]. In these original transport waveforms, the SVD technique was used to solve for a harmonic potential at just two points, once at an initial location \mathbf{r}_i and once again at the final location \mathbf{r}_f . These two solutions were \mathbf{V}_i and \mathbf{V}_f , respectively. The waveform was constructed by using a sinusoidal ramp to smoothly transition between the solutions:

$$\mathbf{V}(t) = \mathbf{V}_f + \frac{1 + \cos(\pi t/T)}{2} (\mathbf{V}_i - \mathbf{V}_f), \quad (5.30)$$

where T was the duration of the transport. As long as \mathbf{r}_i and \mathbf{r}_f were close enough that \mathbf{V}_i and \mathbf{V}_f did not differ drastically, Eq. (5.30) would produce a moving well that would transport the ion. If \mathbf{r}_i and \mathbf{r}_f were too far apart, a third, intermediate, solution could be included. Other ramp functions, involving something other than a sinusoidal time dependence, could also be used.

This method does not constrain the potential during the ramp, which could result in undesired behavior at some point during the waveform. For example, the secular frequencies could vary substantially with position and possibly even become anti-trapping at some location. The farther apart \mathbf{r}_i and \mathbf{r}_f are, the more likely problems will arise in the waveform. In a large-scale array, such as the X-junction array, the ions move far enough that the above approach would require a significant number of intermediate steps to ensure optimal transport.

Therefore, the approach adopted for the X-junction array is to solve Eq. (5.14) for \mathbf{V}_s at every

step along the ion trajectory, where the step size is defined by the grid spacing in the computer model (5 μm). The set of different voltage steps, $\mathbf{V}_{\{s\}}$, are applied sequentially to the trap electrodes, in discrete steps, at an update rate near 500 kHz. A series of low-pass RC filters (160 kHz corner frequency) smooth the transition between the discrete voltage steps. With the harmonic well being constrained at these closely-spaced intervals, the properties of that well do not change noticeably between steps. Appendix D discusses the maximum permissible distance between these waveform steps and what errors are caused when this is exceeded, but the 5 μm spacing used here is well under that value ($\sim 45 \mu\text{m}$). When moving through the complicated potentials near a junction, this approach ensures the ions are always behaving as intended.

The first step to constructing a waveform is to describe the desired trajectory by defining an array of points $\mathbf{r}_{\{s\}}$ that follows that path. The trajectory is typically chosen lie along the trap axis, which we will assume is parallel to \hat{z} . In this case, a series of z values is chosen, starting with the initial location, and it is convenient for these values to coincide with the z component of the grid of trap model. In the ideal geometry, the pseudopotential minimum will lie along the trap axis, but this is not always the case and the location of the minimum may not be exactly known. Therefore, to determine the x and y radial components of $\mathbf{r}_{\{s\}}$, the pseudopotential is searched for the minimum in the radial plane at each z location in trajectory.

The set of $\mathbf{r}_{\{s\}}$ is then used to solve Eq. (5.14) for $\mathbf{V}_{\{s\}}$ at each point to produce the steps of the waveform.

Since the waveforms steps are calculated at equal spacing, if the waveform is stepped through at a constant rate, the ions will be transported at constant speed, with an instantaneous acceleration at the beginning and end. Adjusting this update rate will change the speed of the ion during the waveform.

By slowing the update rate at the beginning and end of the waveform, the ion can be smoothly accelerated to minimize any non-adiabatic behavior. An alternate approach is to use a constant update rate but have steps that are more closely spaced at the beginning and the end of the waveform. In practice, the waveform is calculated using equally spaced points but the computer

controlling the experiment has the ability to interpolate between the $\mathbf{V}_{\{s\}}$ to produce steps at variable distances to achieve the proper acceleration profile using a constant update rate.

For adiabatic transport, the optimal acceleration profile will minimize the acceleration of the moving trapping well. Some promising profiles that have been considered include an error function [Reichle 06] and a hyperbolic tangent [Hucul 08]. The acceleration profile becomes more important as the transport speed increases and approaches the non-adiabatic regime. However, this was not the case for transports in the X-junction array. These transports used either a linear ramp or a sinusoidal acceleration profile and no discernible difference in the motional excitation was observed. This suggests that both transports were well within the adiabatic regime at the speeds the ions were being moved. Attempts to move the ions faster resulted in ion loss, most likely due to the control electronics being limited to a 500 kHz update rate. Though the number of waveform steps could be reduced to increase the speed, eventually the number of steps became too small and between those steps, the potential likely deviated significantly from the intended potential, resulting in the ion being lost. Faster electronics (DACs) would likely allow the ion to be transported closer to the non-adiabatic regime.

The user may desire to change the electrode voltages faster than the RC filter constant, which could drastically distort the set of $\mathbf{V}_{\{s\}}$ applied to the electrodes. However, using the desired set of $\mathbf{V}_{\{s\}}$ along with the appropriate time profile, it is possible to calculate the voltages that need to be applied at the filters to produce the correct voltages on the trap electrodes. In this way, the speed of waveform can be increased beyond the limit set by the filters. This was not done for the X-junction waveforms but may be necessary as transport speeds increase.

As when solving for a single \mathbf{V} , the set of constraints used when finding a waveform is extremely important. The waveform constraints are the same constraints used for a single potential well, but are defined separately at each step of the waveform. The matrix $\mathcal{P} \otimes \Psi^T(\mathbf{r}_s)$ is recalculated at each \mathbf{r}_s . By adjusting the rest of the constraints from step to step, a waveform can be produced where the properties of the trapping well are held constant or follow a prescribed change as the ion moves. For example, the secular frequency of the ions could be made to increase linearly from one

value to another as the waveform proceeds.

The constraints most commonly used for the transport waveforms in the X-junction array were the same set of typical constraints enumerated at the end of Sec. 5.3.3. Given the small spacing between waveform steps, the $\mathbf{V}_{\{s\}}$ usually vary smoothly from step to step, producing well-behaved waveforms. Only the axial secular frequency and its corresponding principal axis is constrained. The other two frequencies and axes are free to drift but tended to do so slowly and smoothly as the waveform progressed, without discontinuities.

However, there were some trouble spots in the waveform where this set of constraints was not sufficient and led to a poorly-conditioned matrix problem or one of the unconstrained frequencies becoming too low. In these situations, additional constraints were added on a case-by-case basis. This often required trial and error to determine a slightly different set of constraints that would produce good results in the given region.

5.4.1 Problems encountered solving for waveforms

In addition to the above problem, there were two other common problems that deserve specific mention: ‘run-away’ voltages and ‘solution jumping’.

The Eq. (5.23) describes a method for finding a solution \mathbf{V}' that ensures the electrode voltages do not become needlessly large. Prior to using this method, it was common to solve for \mathbf{V}_s at location \mathbf{r}_s and get a good solution, but at another location \mathbf{r}_{s+u} several electrodes farther down the trajectory (say 200 μm) the solution \mathbf{V}_{s+u} would no longer be practical. Instead the SVD would return a solution that used distant electrodes at high voltages to create the desired potential rather than local electrodes at low voltages to achieve the same effect. Implementing Eq. (5.23) solved the problem of ‘run-away’ voltages, but at the expense of adding a new problem that I will call ‘solution jumping’.

For the most part, $\mathbf{V}'_{\{s\}}$ varies smoothly from step to step, but occasionally the voltages would jump to a completely different configuration from one step to the next. This can occur when an under-constrained problem has multiple solutions. In that case, any combination of

vectors out of the null space $\tilde{\mathbf{V}}$ can also be added to \mathbf{V}_s to form a solution. If $\tilde{\mathbf{V}}$ is rank 3, then there are four linearly independent solutions. Since the $\mathbf{r}_{\{s\}}$ positions are so close, neither \mathbf{V}_s nor $\tilde{\mathbf{V}}$ change significantly between steps. Thus, the four linearly independent solutions do not change significantly. Equation (5.23) will pick out the solution with the lowest norm. However, as the waveform progresses through the $\mathbf{r}_{\{s\}}$, the norms of the four independent solutions will vary smoothly as a function of position. When one solution drifts below the previous lowest solution, Eq. (5.23) will immediately jump to this new solution, causing a discontinuity in the voltages as function of position. Such jumps are common when transporting over significant distances, whether it be through a junction or in straight linear array. Jumps typically occur when an ion is transitioning from being closest to one pair of diagonally-opposite electrodes (across the trap channel) to being closest to another pair. There will be some threshold where it suddenly becomes easier to produce the desired potential using a new combination of electrodes now that the ions moved closer to those electrodes.

One might wonder if these jumps will cause problems, since both the old solution and the new solution produce identical potentials at the ion. The ion should not notice if the electrode voltages are discontinuous as long as the potential well is continuous. The concern is that the jump will never be perfectly instantaneous, especially since the voltage jump will be filtered by the RC network. While these voltages are transitioning, the potential at the ion could briefly become undesirable.

These jumps can be avoided by spending time to find constraints that do not result in jumps within the regions that are prone to jumping. However, the approach taken when solving for waveforms in the X-junction array (both near the junction and in straight sections) was more drastic: the SVD technique was scrapped in favor of a new method for finding the solutions to Eq. (5.14) which uses quadratic programming to solve a linear least-squares problem. It must be stressed that this new method changes only the algorithm used for solving the matrix equation. It does require changing the form of that equation or the constraints. All the formalism described so far in this chapter still applies with this new method. In addition to maintaining all the abilities of

the SVD method, the new method provides a few new types of constraints that will address both the problem of jumps in the solution and the problem of runaway voltages.

5.4.2 Linear least-squares solution method

By switching from SVD to a linear least-squares (LLS) method, a number of issues with the SVD can be avoided and useful new constraint types can be added.

The LLS solution uses a quadratic programming algorithm. Quadratic programming is a mathematical optimization procedure that minimizes a quadratic function of multiple variables subject to a set of linear constraints. The specific algorithm being used is an iterative active-set method similar to that described in [Gill 81] and is used by the `lsqlin` function in [MATLAB 07]. The general LLS method is used to solve the problem

$$\min_{\mathbf{V}} (\|\mathbf{A} \cdot \mathbf{V} - \mathbf{B}\|_2)^2, \quad (5.31)$$

subject to the constraints

$$\tilde{\mathbf{A}} \cdot \mathbf{V} \leq \tilde{\mathbf{B}} \quad (5.32)$$

$$\tilde{\mathbf{A}}' \cdot \mathbf{V} = \tilde{\mathbf{B}}'. \quad (5.33)$$

Here \mathbf{A} , $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{A}}'$ are constraint matrices while \mathbf{B} , $\tilde{\mathbf{B}}$ and $\tilde{\mathbf{B}}'$ are the associated constraint vectors, all of which are defined by the user. This method finds the \mathbf{V} that minimizes Eq. (5.31). The LLS method is commonly used to find least-squares solutions when fitting data sets. But here it will be used to solve Eq. (5.14), which can easily be written in the form $\mathbf{A} \cdot \mathbf{V} - \mathbf{B}$.

With the LLS method, there are two distinct approaches to entering constraints into the problem. The constraints can be included using Eq. (5.32) and/or Eq. (5.33). Alternately, constraints can be incorporated into Eq. (5.31) by constructing \mathbf{A} and \mathbf{B} in the same manner used for the SVD technique, in particular $\mathbf{A} = \mathbf{C}_1 \mathcal{P} \otimes \Psi^T$ and $\mathbf{B} = \mathbf{C}_2 - \mathbf{C}_1 \mathcal{P} \otimes \Psi_F^T \mathbf{V}_F$. If the user takes this second approach and leaves Eqs. (5.32) and (5.33) undefined, the LLS method will return the same solution that the SVD method would have found. (This is true for both under- and over-

constrained problems.) In this case, all issues, concerns, and tips that applied to the constraints for the SVD method will transfer to this new method.

There is a difference between placing the constraints within Eq. (5.31) versus using Eqs. (5.32) and (5.33). All constraints defined in Eqs. (5.32) and (5.33) will be strictly satisfied (assuming the problem is well formulated). On the other hand, constraints defined within Eq. (5.31) using \mathbf{C}_1 and \mathbf{C}_2 will be used to find a least-squares solution. That is, a solution that minimizes the squared difference between the solution and constraints will be found. If the problem is under-constrained and the Eq. (5.31) constraints do not contradict with the constraints defined in Eqs. (5.32) and (5.33), then all the constraints will be satisfied. Otherwise, the solution's deviation from the $\mathbf{C}_{1,2}$ constraints is minimized but not completely removed. In this case, weighting factors can be used to emphasize the more important constraints in Eq. (5.31).

In practice, constraints are incorporated into all three locations (Eqs. (5.31), (5.32) and (5.33)) when solving for waveforms. The problem is formulated in nearly the same way that it would have been using the SVD method and the constraints are primarily placed within Eq. (5.31). The main advantage of the LLS method comes from using Eqs. (5.32) and (5.33) judiciously to implement a relatively small number of constraints, that were not possible using SVD, to restrain the solution from approaching regimes where the SVD method was having problems.

In the SVD method, all constraints had to be formulated as equalities, while the new method allows for inequality constraints. This is useful, for example, to constrain all the voltages to be within a certain range. The voltage sources driving each electrode will have some maximum voltage output and this inequality constraint can be used to address this practical limit. (For the X-junction, that limit is ± 10 V.) By setting $\tilde{\mathbf{A}}$ equal to the identity matrix, \mathbf{I} , and $\tilde{\mathbf{B}}$ equal to a vector with every entry equal to 10, Eq. (5.32) constrains all the voltages to be less than 10 V. If, instead, $\tilde{\mathbf{A}} = -\mathbf{I}$, the voltages will be constrained to be above -10 V. Combining these two sets of constraints imposes both the upper and lower bound to produce the desired result. This will ensure the voltages are within the specified limits and removes the need to find the smallest set of voltages, \mathbf{V}' , using Eq. (5.23). By not using Eq. (5.23), the propensity for the waveform solutions

to jump will be reduced, but not completely avoided.

A second trick can be used to further alleviate the jumping problem. In this case, we constrain how much a voltage can change from step to step. This is done again using Eq. (5.32), but this time setting

$$\tilde{\mathbf{A}} = \pm \mathbf{I} \quad (5.34)$$

$$\tilde{\mathbf{B}} = \pm \mathbf{V}_{s-1} + \delta, \quad (5.35)$$

where \mathbf{V}_{s-1} is the solution to the previous step in the waveform and δ is a vector of the same size as \mathbf{V} where each entry is the maximum deviation the user will allow between steps for the given electrode voltage. By setting δ small enough, the solutions cannot jump to a completely new configuration and will instead take several waveform steps as it transitions smoothly from one configuration to another. In the waveforms for the X-junction array, $\delta_i = 0.2$ V (for each electrode voltage) and the steps are spaced by $5 \mu\text{m}$.

Another useful constraint is to set the maximum difference between two of the secular frequencies. Equation (C.2) gives the constraints needed to fix the secular frequencies so that $\omega_i^2 - \omega_j^2 \doteq w^2$. This constraint is useful, for example, to make sure the two radial frequencies do not drift too far apart, which may allow one to become degenerate with the axial frequency. However, forcing an equality may be too rigid. The same constraint can be used, but by including it in Eq. (5.32) where it will be constrained as an inequality. This provides more flexibility and makes this constraint much more useful in regions where the frequencies tend to drift excessively without using some type of constraint. A similar inequality constraint can be placed on the orientation of the principal axes to discourage the axes from aligning in an undesirable direction.

The LLS method can also be seeded with an ‘initial guess’, which tells the optimization where to begin. Seeding each step in the waveform with the solution to the previous step can provide more continuity between steps in the waveform.

Though the LLS method provides a means for bounding the individual control voltages to be within practical limits, this does not ensure that the algorithm will pick the solution that uses

the smallest voltages. When the same potential at the ion can be constructed using either small voltages on the nearest electrodes or larger voltages on more distant electrodes, the first solution is likely preferred by the user (but not necessarily by the algorithm). Thus it may be convenient for the user to place a weak constraint to force a preference for small voltages. This can be performed by placing a series of constraints within Eq. (5.31) that set every voltage to zero. At the same time, the weight of these constraints should be set to a very small value (say 10^{-8}). In doing this, the problem is made over-constrained and a least-squares solution will be found. Since these zero-voltage constraints are very weakly weighted, they will be largely ignored by the solving algorithm and priority will be given to the other constraints. Yet, when two solutions exist, these weak constraints will ensure that the solution involving the smallest voltages will be chosen.

The new LLS method is quite similar to the SVD pseudoinverse method and uses much of the same formalism. The identical constraints used in SVD are still included in Eq. (5.31). The advantage is that the new method provides more guidance by placing inequality constraints that are not forcing unnecessarily strict conditions on the solution. In addition, when all the constraints cannot be fulfilled, the LLS method provides a convenient means for ensuring the most critical constraints (such as $|V| \leq 10$ V) are met while allowing deviation from the rest of the constraints based on their relative importance. This resolves many of the pitfalls of the previous method while enabling a new set of handy constraints that were previously missing.

5.5 Other types of waveforms

Thus far, the discussion has been limited to transporting single ions, assuming only one ion is in the trapping well. However multiple ions can easily be placed in a single trapping well and moved together, without changing the waveform. Certain parameters may be more important when multiple ions are present, such as frequencies of the different modes. So specialized waveforms that use extra constraints to account for these extra parameters may be needed. However, in the X-junction array, the same waveforms were used to transport both a single ion and a pair of ions. These waveforms worked well in both cases despite not being specifically designed for pairs of ions.

As the number of ions in a large-scale array increases and rearranging of the ions becomes important, it will be necessary to have multiple trapping wells moving simultaneously and independently. The methods used in this chapter can be easily expanded to incorporate multiple wells. By expanding Eq. (5.14) to include not only $\Psi^T(\mathbf{r}_0)$ but also the potentials calculated at the minima of other wells, $\Psi^T(\mathbf{r}'_0)$, the SVD and LLS methods can find solutions that create two or more simultaneous trapping wells. However, as the number of wells grows large, the problem may become too large to solve quickly. It will likely become necessary to separate the problem into multiple regions that are spatially distant and do not interact. That way a trapping well in one region can be moved without considering how it might influence the potential in other regions.

Operations such as the crystal rotation in [Splatt 09] can be tackled with the techniques in this chapter as well. By creating a series of constraints on the principal axes that cause them to rotate, a linear crystal of ions within that potential will rotate with the potential. A similar operation can be performed in the center of an X-junction.

Another type of waveform, the separation and combination of trapping wells, will also be necessary. These waveforms will require extra consideration and are not provided by use of techniques in this chapter. To perform separation, two trapping wells must move towards each other, and at some point they become close enough that the total potential is better described by a quartic potential centered between the two wells. Once the wells are this close, the quartic term in the potential can be gradually changed to a quadratic term and the two wells will join [Home 06]. Separation will be performed by running this procedure in reverse. Since these process depend on the quartic potential, one could consider extending \mathcal{P} in Eq. (5.13) to fourth-order derivatives, which should allow control over the quartic term.

5.6 Waveforms for the X-junction array

The regions of the X-junction array can be broken into three different categories for transport: (1) straight sections, (2) from the load zone to the experiment zone, and (3) near the X-junction. These are listed in ascending order relative to the difficulty of producing good waveforms in the




Figure 5.2: Three different types of regions for transport. The straight sections (outlined in blue) are easiest to transport in and there are three of them on the three legs of the junction. The region from the load zone to the experiment zone is seen on the left (outlined in black). The region near junction (outlined in red) is most difficult location in which to find decent waveforms due to pseudopotential barriers.

given region and can be seen in Fig. 5.2.

Straight sections

An example of a ‘straight sections’ is the region near the experiment zone, where the channel width is $200 \mu\text{m}$, from $z = -1200$ to -250 . There are similar regions on the other sides of the junction, one from $z = 200$ to $z = 900$ and another from $x = 200$ to $x = 500$. These regions are far enough from the junction that they closely resemble the standard linear Paul trap.

Solving for waveforms in these regions is straight forward using the standard set of constraints listed at the end of Sec. 5.3.3. The axial secular frequency is constrained to be 3.6 MHz, while the radial frequencies are unconstrained and varied between 12 and 14 MHz. A sample waveform solution for this type of region can be seen on the left side of Fig. 5.3.

The load zone to the experiment zone

This region of the trap array has a channel width of $300 \mu\text{m}$, making the radial confinement slightly weaker. This region includes two 135° bends in the channel and a transition from the $300 \mu\text{m}$ width to the $200 \mu\text{m}$ channel width, near the experiment zone. This transition in channel width does not produce any axial pseudopotential, due to the axial symmetry of the trap. For this

reason, the process of solving for waveforms in the vicinity of this transition is not different than in the straight sections.

In the regions where the channel bends, the trap axis smoothly transitions between lying along \hat{z} to lying along $\hat{x} + \hat{z}$, rounding the corner gradually. Though intuition may suggest this bend may complicate the process for solving for waveforms, this did not prove to be the case. The only difference in the constraints used here versus those used in the straight sections was that the electrodes on opposite sides of the axis were not constrained to be equal. Since they will have different proximities to the ion, those electrodes will need different voltages.

The only real complication, as far as waveforms are concerned, in the region between the load zone and the experiment zone, was that the control electrodes were unusually wide in this region (641 μm wide). To produce a reasonable axial frequency at certain locations along the trap axis, larger voltages were needed. For example, when the ion is moving between the edges of electrode pairs 9/10 and 11/12, the axial trapping must be created by applying large positive voltages to 7/8 and 13/14 while applying large negative voltages to 9/10 and 11/12. The axial frequency, which was constrained to be 2.0 MHz at the load zone and 3.6 MHz at the experiment zone, was relaxed to 1.2 MHz during most of this transport and constrained to dip to 0.95 MHz at some points to ensure all the electrode voltages remained within ± 10 V.

As in the straight section, waveforms in this region of the array were found using the standard set of constraints listed at the end of Sec. 5.3.3, with the exception that electrodes on opposite sides of the bends were not constrained to be equal. The unconstrained radial frequencies fluctuated between 7.6 and 9.7 MHz while in the 300 μm -wide channel.

Near the X-junction

This was the most challenging location in which to find suitable waveform solutions, due to the axial pseudopotential barriers seen in Fig. 4.2.

The waveform procedure calls for a harmonic potential to be calculated at each step of the waveform. Along the axis of an ideal linear Paul trap, the axial harmonic potential is constructed entirely of the control potentials without a contribution from the pseudopotential. This is not the

case near the X-junction, where the pseudopotential contribution is large. The problem is that the pseudopotential provides a barrier and acts contrary to the user's typical constraints. In locations on the side of the barrier, the pseudopotential provides a strong force pushing the ion away from the barrier and the desired trapping location. In order to counter this and produce a well at correct location, the control potentials must provide an equally strong force in the other direction. At the top of the pseudopotential barrier, the pseudopotential is strongly anti-confining along \hat{z} with $\omega_{\text{rf},z} = 2\pi \times 5.03i$ MHz. Here, the control electrodes must provide enough positive curvature to overcome this anti-confinement and produce the proper axial frequency, while still maintaining trapping in the radial directions.

In both cases, to overcome the pseudopotential's unwanted effects along the barrier, much larger control voltages are needed than in the other regions of the trap array. In addition, it is advantageous to have narrow control electrodes in this region to provide the maximum possible control-potential gradients. This second issue was considered during the design phase of the junction and the control electrodes were made sufficiently narrow. This allows proper waveforms to be defined without exceeding the ± 10 V limit of the DACs.

A plot of the final waveform used between the experiment zone and the center of the junction is displayed in Fig. 5.3. As can be seen, when the trapping minimum is near a certain electrode, that electrode is held at a low voltage, while its neighbors are held at high voltages to produce a trapping well. Electrodes are paired with their counterpart on the opposite side of the trap axis (e.g., 15 with 16, 17 with 18, etc.), so only odd electrodes are depicted in the figure, with the exception of 46. Any electrode that is held constant during the waveform is excluded from Fig. 5.3(a) for clarity.

The left side of the waveform, up to $z = -250 \mu\text{m}$, is typical of the waveforms in the other straight sections of the trap.

The pseudopotential barrier is located between $z = -200 \mu\text{m}$ and $z = 0 \mu\text{m}$. In this region, some of the control voltages get very large. In particular, electrodes 23, 24, 25, and 26 reach the voltage limit of ± 10 V constrained using the LLS method. Some of the other voltages in this region

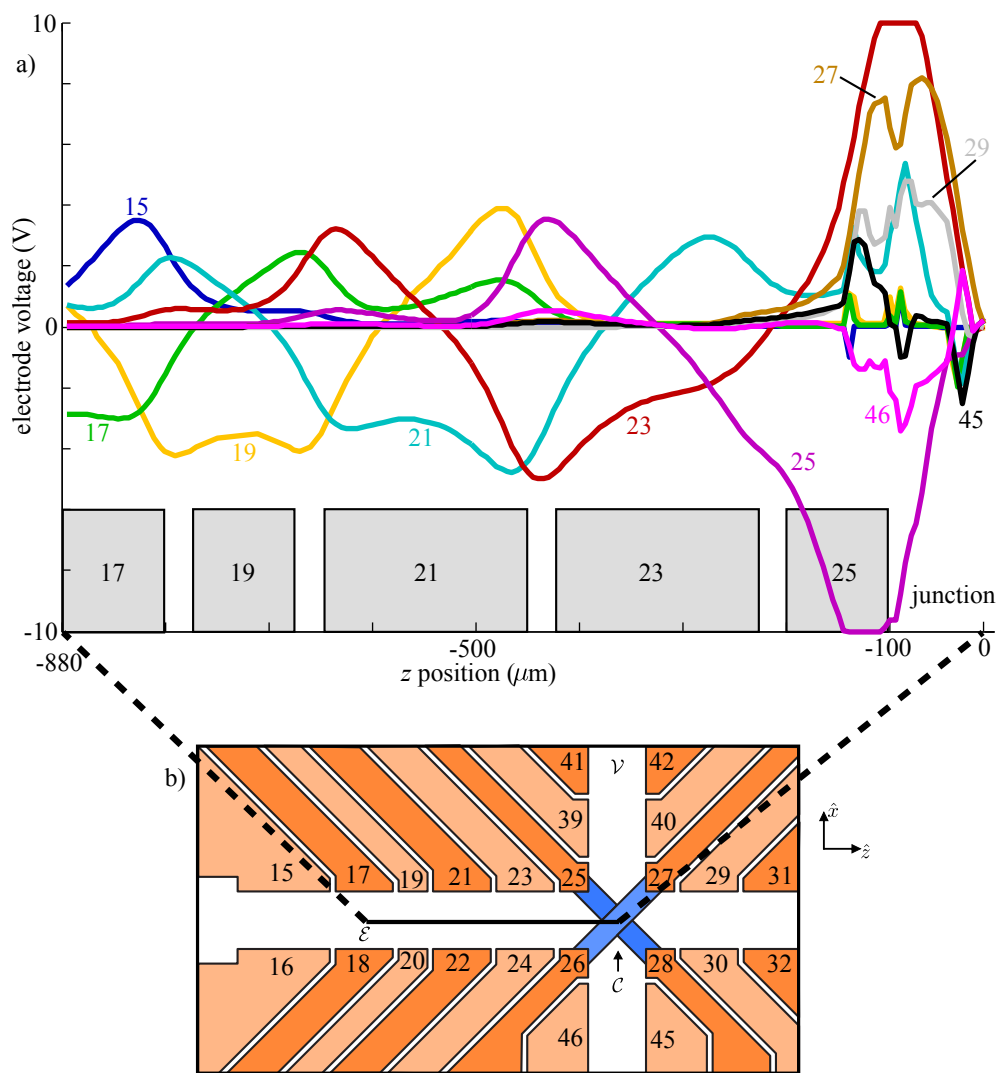


Figure 5.3: a) The waveform used to transport an ion from the experiment zone, located at $z = -880 \mu\text{m}$, to the center of the junction at $z = 0 \mu\text{m}$. The relative locations of some of the electrodes are depicted, along with their electrode number, by the rectangles in the bottom of the figure. The right-most $100 \mu\text{m}$ is inside of the junction. The voltage traces are numbered to show which electrode they correspond to. b) A schematic of the trap, showing the range over which this waveform transports. The control electrodes from both layers are shown without the rf electrodes (with the exception of the bridges).

have sharp and abrupt changes (e.g., electrode 21). These result from the LLS constraint described in Eqs. (5.34) and (5.35), where a limit is set to how much the solution at a given point can deviate from the previous step. At these locations, the solution wants to jump, but the constraint instead forces the voltages to ramp linearly over several steps. The spike occurs when the solution wants to jump again into a different configuration and the same constraint forces the voltage to ramp back down gradually. Though these voltages are not smooth, they are continuous, and that is sufficient to ensure the ion sees a smooth potential. Neither these sharp features nor the voltages that rail against the ± 10 V limit result in abnormalities or discrete changes in the trapping potential seen by the ion.

To move an ion into the center of the junction, the waveform in Fig. 5.3 is run from left to right, and to move the ion out of the junction, it is just run in reverse. A typical transport duration (to move $880 \mu\text{m}$) was approximately $140 \mu\text{s}$, with $40 \mu\text{s}$ to cross the pseudopotential barrier. The waveforms that move ions into the other two branches of the junction ($z > 0 \mu\text{m}$ or $x > 0 \mu\text{m}$) very closely resemble this waveform due to the symmetry of the trap, but the voltages are applied to a different set of electrodes along the other legs of the junction.

In addition to issues with the pseudopotential, a second complication occurs at the center of the junction. At this location, the pseudopotential will confine an ion in all three dimensions without the need for control electrodes. If no control electrodes are applied, one of the radial secular frequencies will be degenerate with the axial frequency at 5.7 MHz. (This is a result of the trap's symmetry. In this location, there are actually two axial frequencies, one for the original linear array that the ion had been moving in and one for the perpendicular array that forms the junction.) Control potentials can be applied to break this degeneracy, but there is a relatively narrow range of permissible frequencies and attempts to exceed this will lead to the trap becoming unstable.

The chosen standard axial frequency of 3.6 MHz could not be realized at the center of the junction. Instead, the axial frequency is ramped up from 3.6 MHz to 5.7 MHz as the ion approaches the center. As long as this ramps occurs adiabatically, this does not result in excitation of the ion's




Figure 5.4: Predicted secular frequencies as a function of position corresponding to the waveform in Fig. 5.3. The axial frequency along \hat{z} is constrained to be 3.6 MHz during the majority of the transport while the radial frequencies are unconstrained. As the ion ascends the pseudopotential barrier, the frequency linearly ramps up to 4.2 MHz. Beyond the apex of the barrier, a second linear frequency ramp is performed up to 5.7 MHz. As the ion approaches the center of the junction, one of the radial frequencies becomes near-degenerate with the axial frequency.

motion. In fact, a higher secular frequency makes the ion less susceptible to noise that can excite that motion (see Ch. 7), so raising the frequency during the transport reduces the total motional excitation during this process.

The frequency is raised in two ramps and the three secular frequencies are shown in Fig. 5.4, as a function of the ion position. The first ramp occurs as the ion ascends the potential barrier while approaching the junction. The second ramp occurs as the ion descends the barrier and moves into the center. The first ramp takes the frequency from 3.6 MHz to 4.2 MHz, which was determined to be close to the highest axial frequency possible at the top of pseudopotential barrier while limiting $|V_N| \leq 10$ V. (A higher frequency would be used if possible, to reduce motional excitation.) Once the ion passes the apex of the barrier, the frequency can be further increased to 5.7 MHz. The choice was made to allow the two frequencies along \hat{z} and \hat{x} to be degenerate at the center. If transporting a pair of ions in the same trapping well, it would be better to keep a small difference between the frequencies so the string of ions always has a well defined direction. However, no

detrimental affects were seen when using this waveform with pairs of ions in actual experiments, suggesting that perhaps the two modes were sufficiently separated in the actual apparatus despite what the model predicts.

The orientation of the three principal axes rotates as the ion approaches the center of the junction. They change from approximately $(\hat{z}, \hat{x}+\hat{y}, \hat{x}-\hat{y})$ near the experiment zone to $(\hat{z}, \hat{x}, \hat{y})$ at the center.

The constraints used near the junction are the standard set given in Sec. 5.3.3, excluding the final constraint that electrodes on opposite sides of the trap axis be constrained to be equal. The solutions seem to roughly fulfill this excluded constraint without the need to explicitly apply it. In addition, the LLS inequality constraints described in the previous section were used. Through a process of trial and error, a set of appropriate frequency ramps was determined to provide a reasonable solution near the junction.

In the original waveforms (those used for $^{24}\text{Mg}^+$ and those first tried on $^9\text{Be}^+$), the computer model assumed the designed trap geometry was faithfully replicated in the constructed trap. In reality, there was some misalignment between the trap wafers due to manual assembly. Though the waveforms produced from the ideal model worked (but poorly) on $^{24}\text{Mg}^+$, they did not work on $^9\text{Be}^+$ and the ion was lost during every transport attempt. By measuring this misalignment using photographs, the model was updated to more accurately depict the geometry of the assembled trap. The new waveforms using this updated model successfully transported $^9\text{Be}^+$ through the junction and are what is reported above. (The updated model should also improve the results for $^{24}\text{Mg}^+$ transport, but this has yet to be verified.)

The method used to solve for waveforms treats the rf potential as a static pseudopotential using the secular approximation and did not consider the Mathieu stability parameters, a_i and q_i . Figure 5.5 depicts these stability parameters during the waveform for the three principal axes. At no point does the solution approach an instability region predicted by the Mathieu equation.

Another important issue to consider is the micromotion and any associated nonlinear behavior. Figure 5.6(a) shows the peak amplitude of the axial micromotion as the ion moves through

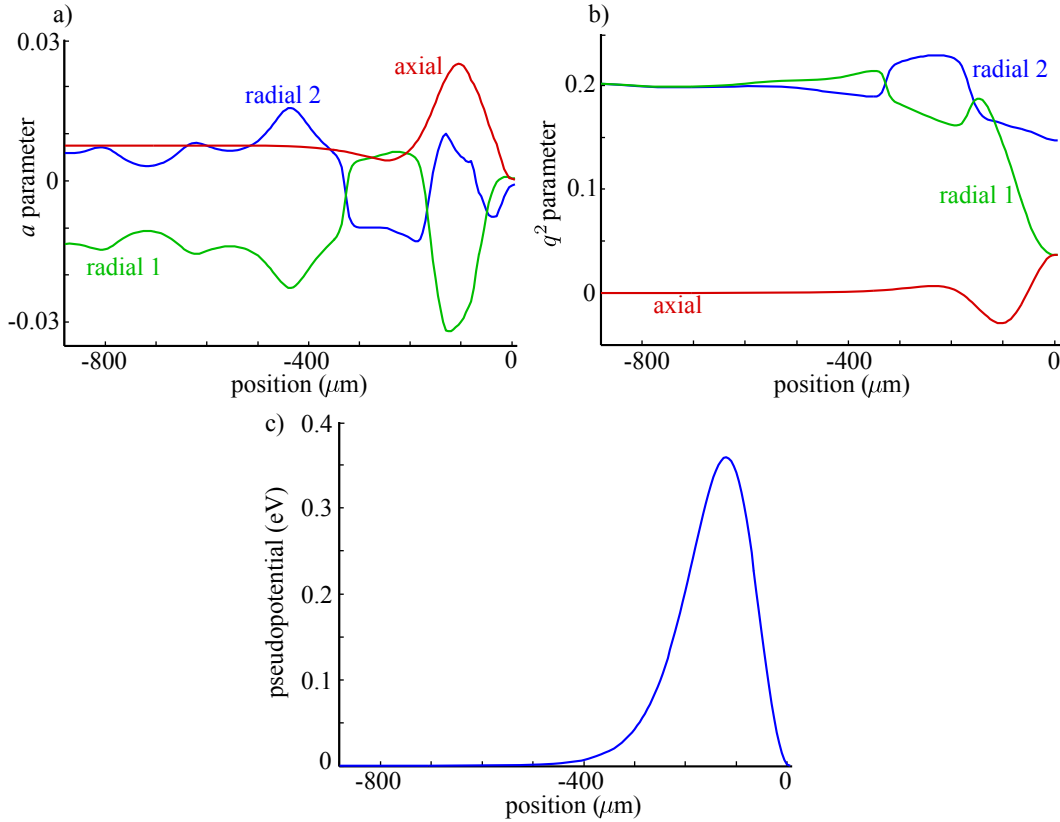


Figure 5.5: a) The Mathieu a_i stability parameter during the transport waveform in the X-junction array. The axial and two radial modes are shown, where ‘radial 1’ is the mode that becomes degenerate with the axial mode at the center of the junction. b) The Mathieu q_i^2 parameter during the waveform. c) The pseudopotential along the trap axis showing the location and height of the barrier. This assumes $V_{\text{rf}} = 204$ V (peak) and $\Omega_{\text{rf}} = 2\pi \times 83$ MHz and ${}^9\text{Be}^+$ is the ion being transported.

the waveform to the center of the junction. This amplitude becomes quite large, exceeding $5 \mu\text{m}$, as the ion approaches the apex of the pseudopotential barrier. The spatial extent of the harmonic wave packet, z_0 , is only 11 nm at the same position. Section 3.4.1 showed that a linear rf electric field, which is what the ion will experience (in the axial direction) along the barrier, will not lead to instability regardless of its magnitude. Even though the micromotion is large, it is not of immediate concern. Instead, concern should be focused on potential instabilities due to nonlinear contributions to the Mathieu equation while the ion is on the pseudopotential barrier. The adiabaticity parameter, η_i , given in Eq. (3.49), can be used to characterize this. Figure 5.6(b) shows that the axial η_z during transport into the junction never exceeds 0.14. From the discussion in Sec. 3.4.2,




Figure 5.6: a) The axial micromotion peak amplitude for various locations along the trajectory of the waveform that moves from the experiment zone into the center of the X-junction. b) The axial adiabaticity parameter, η_z , along the same path.

this is unlikely to lead to trap instability. From this, it is possible to conclude the ${}^9\text{Be}^+$ waveform that moves ions from the experiment zone into the X-junction should have no inherent stability problems. This conclusion is further supported by successful demonstration of this transport in the experiment.

The waveforms used to transport ${}^{24}\text{Mg}^+$ were similar to those discussed for ${}^9\text{Be}^+$. One difference was the lower secular frequencies of 1.8 MHz in the axial direction and 4.1 and 5.0 MHz in the radial directions. The ${}^{24}\text{Mg}^+$ waveforms were the first waveforms produced for the X-junction array and were produced with less sophisticated criteria than later Be waveforms. They included fewer frequency ramps and were calculated using the ideal computer model rather than the model that incorporated the measured trap misalignment. Many of the lessons learned perfecting ${}^9\text{Be}^+$ transport could be applied to ${}^{24}\text{Mg}^+$ transport and would likely improve the ${}^{24}\text{Mg}^+$ results significantly.⁵

⁵ The reason further testing of ${}^{24}\text{Mg}^+$ transport has not been performed is that the X-junction apparatus is no longer in a laboratory with the lasers needed for ${}^{24}\text{Mg}^+$. This trap was initially tested in a ${}^{24}\text{Mg}^+$ laboratory but was moved to a new laboratory where ${}^9\text{Be}^+$, which is a better qubit, could be used.

Chapter 6

Apparatus

The experimental apparatus will be explained in this chapter. This will begin with a description of the fabrication and assembly of the X-junction trap array. Then the systems for creating the potentials applied to the trap electrodes, including both the rf and control potentials, will be described. Emphasis was placed on filtering electronic noise to limit excitation of the ion's secular motion. This is followed by an explanation of the vacuum system in which the trap is installed and the components used in the process of loading ions into the trap. The next section will describe the laser sources and the imaging devices used. The chapter will end with a brief explanation of the computer system employed to control the experiment.

6.1 Ion trap

The X-junction trap array is constructed from laser-machined alumina wafers similar to previous two-layer traps [Rowe 02, Barrett 04, Wineland 05]. These wafers have gold deposited on them to produce the electrodes. The trap consists of five wafers in total:

- (1) The top 'trap' wafer provides one of the electrode layers in the two-layer electrode design. (125 μm thick.)
- (2) The bottom 'trap' wafer provides the other electrode layer. (125 μm thick.)
- (3) A spacer wafer is used to achieve the correct separation between the two trap wafers. (215 μm thick.)




Figure 6.1: The stack of five wafers seen in cross section along the axis of a linear section (not to scale). Each wafer has a channel cut through it to allow a place for the ion and to provide a path for the laser beams to pass through the wafer stack. The top and bottom wafers provide the electrode structure. The RC filters are placed on the filter board and the electrical connections to the rest of the apparatus originate from this wafer. The bias wafer is a single electrode structure used to null stray electric fields along \hat{y} that cannot be nulled using the control electrodes alone. The ion is trapped between the top and bottom wafers. Gold (represented in yellow) is coated on the top side of the trap wafers and wraps around to the bottom side, and vice-versa for the bias wafer. Gold is used to form electric traces and pads on both sides of the filter board. No gold is applied to the spacer wafer. Wire bonds form the electrical connections between the various wafers.

- (4) A filter board wafer provides both the overall structural support and a location to attach the internal RC filters used to reduce electronic noise. (500 μm thick.)
- (5) A bias wafer functions as an additional electrode for compensating stray electric fields that cannot be compensated using the control electrodes alone. (125 μm thick.)

These five wafers are stacked as seen in Fig. 6.1 and then clamped together using two screws. The top trap wafer, spacer wafer, and bias wafer are all square with 1.6 cm sides. The bottom trap wafer is 1.6 cm by 1.8 cm so that it extends farther out and can be wire-bonded to. The filter board is much larger, at 5.6 cm by 8 cm.

All five wafers have channels cut completely through the wafers during laser machining so that, when assembled, the control lasers can pass through the channel, interact with the ion, and then exit the other side of the wafer stack. The spacer wafer is recessed from the tips of the electrodes so that it is no closer than 2.25 mm to any location where ions would be trapped, which reduces the proximity of dielectrics that could potentially hold charges that cause stray fields.

6.1.1 Trap fabrication

The wafer substrate is a high-purity alumina, which is a ceramic consisting of 99.6% Al_2O_3 and a small amount of binder. Alumina is chosen for its low rf loss tangent (0.0001 at 1 MHz and 0.00006 at 6 GHz [CoorsTek]) and relative inexpense compared to sapphire or similar materials. Alumina has the additional advantage of having a high thermal conductivity to dissipate heat caused by resistive losses in the trap. Finally, alumina is inert and compatible with many of the cleaning and fabrication methods used.

The substrates were ground and polished to the correct thickness and then cut to be larger than the final wafer dimension. The next step was to laser-machine the features into the substrates, which was performed by [Resonetics]. Two different lasers were used for this, a CO_2 and Excimer laser. The CO_2 laser is used for larger features, such as the screw holes. The Excimer laser was used for all of the electrode features. The possible feature size is given by the minimum laser focus of 20 μm . As the laser cut progresses deeper into the wafer, the side-walls of the cut taper inward. The minimum feature size of 20 μm is only achievable on the side of the wafer the cutting laser will exit from. For the two trap wafers, the laser-exit side was chosen to be the side that would be closer to the ion once the trap was assembled.

Both trap wafers consist of a two main channels that cross at a right angle to form the X-junction. On one side of each channel is the rf electrode that extends the length of that channel, while on the other side is a series of control electrodes. Figure 6.2 shows the assembled trap with these channels clearly visible. At the center of the junction is 70 μm -wide bridge that crosses the channels and connects the rf electrodes. There was some concern as to whether this bridge would survive the laser-machining process but this did not prove to be an issue.

To define the control electrodes, a series of cuts, 20 μm wide and perpendicular to the main channels, were made through the wafer. These cuts physically separate each electrode structure by a vacuum gap. The cuts extend ~ 1 mm back from the main channels. The gold was placed on the wafer so that it did not extend as far back as the cuts, causing each electrode to be electrically




Figure 6.2: The assembled X-junction array. The white material is the alumina substrate, that has electroplated gold deposited on it to form the electrodes. The edges of bottom trap wafer can be seen along the side of the channel and small portions of the bias board are also visible even farther back, at the top and bottom of the vertical channel, as well as to the left of the load zone.

isolated from each other. Charges can build up on exposed insulating material close to the ion, which leads to slowly varying stray electric fields that can push the ion off of the trap axis. Not having alumina material between the electrodes helps mitigate this problem.

The gold on the trap and bias wafers consists of an evaporated layer that defines the pattern, and then an electroplated layer to achieve a greater thickness ($\sim 3.5 \mu\text{m}$). Prior to evaporating gold, the substrates are rigorously cleaned. Each wafer was rinsed in de-ionized (DI) water and then submerged in ‘piranha etch’ for 30 minutes. Piranha etch is a mixture of concentrated H_2SO_4 and H_2O_2 in a 4:1 to 2:1 ratio that is heated to 100°C and agitated by a magnetic stir bar.¹ Piranha etch aggressively attacks and removes many materials, including organics and many reactive metals. The substrates were then rinsed in a rapid succession of DI water, acetone, and isopropynol and dried with a jet of dry, clean nitrogen. Finally, the substrates were subjected to three minutes of oxygen plasma cleaning before being inserted into the gold evaporation machine.

¹ Extreme caution should be used when working with Piranha etch. If the concentration of H_2O_2 exceeds that of H_2SO_4 , the exothermic reaction between chemicals can become self-sustaining and lead to an explosion.

The gold was deposited using an electron-beam evaporator at $\sim 10^{-7}$ Torr. Prior to each gold evaporation, a titanium adhesion layer (~ 30 nm thick) was deposited and then the system was switched to gold without breaking vacuum (there was a two minute delay between Ti and Au deposition). The Ti was deposited at 0.15 nm/s and Au at 20 nm/s. The substrate was placed at the top of the vacuum chamber, while the metal source was located at the bottom. The evaporating metal travelled upward from the source and adhered to the first surface it came into contact with, producing a nearly uniform layer across the top of the chamber (and on the substrate).

Shadow masks were used to define the pattern of the evaporated gold on the substrates. These masks were 250 μm thick stainless steel sheets with a pattern cut through the steel by [SolderMask]. When placed flush against the surface of the substrate, the mask blocked the evaporating gold from depositing on the wafer except where the mask material had been cut away. The evaporated gold was applied in three steps per trap wafer, each step using a different shadow mask. The first mask had a single hole that allowed gold to be deposited on the ‘top’ side of the electrode structures, extending from the center channels back nearly 1 mm along both the rf and control electrodes. Once this gold had been applied, the mask and substrate were removed from the evaporation chamber, the mask was replaced, and then they were returned to the chamber for a second deposition. This second mask defined the traces that lead from the the edges of the wafer to the electrodes. In Fig. 6.2, it is possible to see where these two depositions overlap. The mask and substrate were removed a second time, the wafer was flipped over to the bottom side, and a third mask, identical to the first, was used to coat the backside of the electrodes.

When applying the electrode traces (using the second mask), the mask and substrate were oriented normal to the incidence of the evaporating metal. A 110 nm thick layer of Au was deposited on top of a 22 nm Ti adhesion layer.

When evaporating gold using the first and third masks, the goal was to coat the top and bottom surfaces of the electrodes, as well as the sidewalls of the main channels. By metalizing the sidewalls, the top of each electrode are electrically connected to its bottom and traces are only needed on one side of the wafer. In addition, coating the sidewalls reduces the amount of

exposed dielectric, which would otherwise accumulate charge during experiments and push the ion in unintended directions. By orienting the substrate and the mask (which is flush against the surface of the substrate) at 45° (i.e., $\hat{x} + \hat{y}$) relative to the metal flux, both the surfaces and sidewalls could be coated. However, the sidewalls do not all face the same direction and cannot be coated simultaneously. To address every sidewall, the mask and substrate were slowly rotated together about the $\hat{x} + \hat{y}$ axis while evaporating, exposing each sidewall during some phase of the rotation. Both the first and third evaporations were performed in this manner to ensure good sidewall coating from both sides of the wafer.

Ideally, the sidewalls of the slits between control electrodes should also be metalized to reduce exposed dielectric. However, due to the slits' aspect ratio ($\sim 20 \mu\text{m}$ wide and $125 \mu\text{m}$ deep), it was not possible to ensure a full-thickness coating deeper than $\sim 20 \mu\text{m}$ from the substrate's surface.

For the regions exposed during the first and third masks, the thickness of the evaporated metal was 30 nm of Ti and 400 nm of Au. These values take into account the fact that the depositions occurred at 45° , however they correspond only to the thickness deposited on the top surfaces (during the first evaporation) and the bottom surfaces (during the third evaporation) of the substrate. Those surfaces remained unobscured throughout the entire revolution during each respective evaporation. Less material was deposited on the sidewalls of the laser cuts because those faces are obscured during at least some portion of the rotation. However, most of these sidewall surfaces received two layers of metal, once during the first evaporation and again during the third evaporation after the wafer had been flipped.

After the evaporation was completed, the thickness of gold on each substrate was increased to $\sim 3 \mu\text{m}$ using electroplating. This was performed in a home-made plating setup depicted in Fig. 6.3. The plating solution was Metakem SF gold solution, which is a non-cyanide solution. One liter of this was heated to 60°C and slowly agitated by both a magnetic stir bar and a hand-held stirring rod. The hand stirring was done to ensure the solution did not flow past the part in a well defined pattern, which can lead to variation in the deposition rate. A platinum anode was used while the evaporated gold on the substrate acted as the cathode. The proper electrical current




Figure 6.3: Electroplating setup. The wafer and anode sit in a 6" Pyrex tub and are electrically connected through a current supply. Both a magnetic stirring bar and a hand-held stirring rod are used to agitate the heated electroplating solution.

was calculated to achieve a current density of 10 mA/cm^2 for the surface area of cathode. At this current density, the electroplated gold would deposit at a rate of $38 \text{ } \mu\text{m}/\text{hour}$ and required 284 s for the desired $3 \text{ } \mu\text{m}$ thickness.

Between the evaporation and electroplating, the substrate was cleaned again using an oxygen plasma. Before the substrate was submerged, it was biased with a small negative voltage ($\sim 0.01 \text{ V}$) to ensure that a reverse bias would not occur as the part was submerged in the solution. Failure to apply such a bias led to problems with the surface quality and adhesion of the electroplated layer. The wafer was then submerged and the current rapidly increased to predetermined value for

the prescribed period of time. This current was then turned off and the wafer removed from the solution. The part was thoroughly rinsed with successively with DI water, acetone, and isopropynol and then dried with compressed nitrogen gas.

The substrates were cut to be larger than the final wafer so that a gold frame could be evaporated around the wafer to connect all the electrodes traces into a single contiguous gold pad. Thus a single electrical contact with the substrate connected all the electrodes during the electroplating process. After the plating was completed, the substrate was cut to the proper size using a dicing saw. This removed the gold frame and electrically isolated the electrodes.

During the evaporation, the shadow mask should sit against the surface of the substrate, but there was often a small unwanted gap between the two. This would allow a thin layer (10 nm) of metal to be deposited behind the mask, extending tens of micrometers out from any openings in the mask. However, this layer was not well adhered to the substrate, and during the electroplating process, it would increase in thickness and pull off of the substrate, forming long curls at the edges of the traces. To avoid this, the substrate was placed in an ion mill where 40 nm of gold was removed to destroy this thin layer. The ion mill will remove some alumina as well, but at a negligible rate compared to gold.

The total thickness of the metal layer on the substrates after electroplating was $\sim 3.5 \mu\text{m}$ with ± 10 nm surface roughness. The bias wafer was treated in the same manner as the trap wafers.

Gold was applied to the filter board in a different manner. There, DuPont 5062/5063 gold paste was applied by CoorsTek in a process similar to silk-screening. This paste is then baked to 800°C in an O_2 atmosphere, which cures the paste into a high-purity gold. The thickness of this gold is ~ 0.1 mm.

6.1.2 Assembly

The trap wafers were assembled into the trap following fabrication. The five wafers were stacked in the order seen in Fig. 6.1 and a top-down view can be seen in Fig. 6.5. The trap wafers, spacer wafer, and bias wafer are positioned at 45° to the rectangular shape of the filter board. A

set of two screw holes are cut through all five wafers and titanium screws are inserted to clamp the wafers together. Before tightening the screws, two 1 mm gauge pins were inserted into two gauge-pin holes that are also pass through all five wafers. The pins are used to align the wafers to each other. The screws were then tightened and the gauge pins removed.

Proper wafer alignment proved difficult to achieve. After multiple attempts, the wafers were aligned so that the lateral offset between the two trap wafers is minimized near the junction. The resulting offset is a lateral shift of the top wafer (relative to the bottom wafer), primarily in the $-\hat{x}$ direction, of $\sim 5 \mu\text{m}$ near \mathcal{F} . At \mathcal{E} , the shift is $\sim 10 \mu\text{m}$ in the same direction. At \mathcal{L} , the shift is $\sim 25 \mu\text{m}$, but this is the load zone, a location where precise trap alignment is not important. These shifts, which are measured using microscope images, correspond to a misalignment angle between the \hat{z} axis of the top trap wafer and the \hat{z} axis of the bottom trap wafer of 0.22° .

In addition to this angular shift, the distance between the top and bottom wafers was $250 \mu\text{m}$ rather than the intended $220 \mu\text{m}$. This may have been due to the wafers not sitting flush against each other due to warping. Neither this error in the wafer spacing nor the angular error proved to be a problem for transport and trap functionality, as long as the errors were included in the computer models of the trap.

The filter board (Fig. 6.4) has gold traces that connect contact pads on its perimeter to pads next to the trap wafers for routing the control potentials to the trap. At the perimeter, spring contacts press onto the pads to form an electrical connection between the pads and wires that carry the control potentials. At the trap, $50 \mu\text{m}$ diameter gold wire bonds connect the filter board to the trap wafers. The bottom trap wafer is made slightly larger so that it extends beyond the top wafer and can be bonded to.

Along the filter-board traces, between the perimeter contacts and the trap, are some gaps where components for an RC filter are placed (see Sec. 6.2.2). A resistor is used to bridge the gap and one side of a capacitor is bonded to the trace. A gold ribbon is strung across the tops of all the capacitors and then, suspended above the wafer, it connects to one of several grounded pads to produce the capacitive short between the trace and ground.

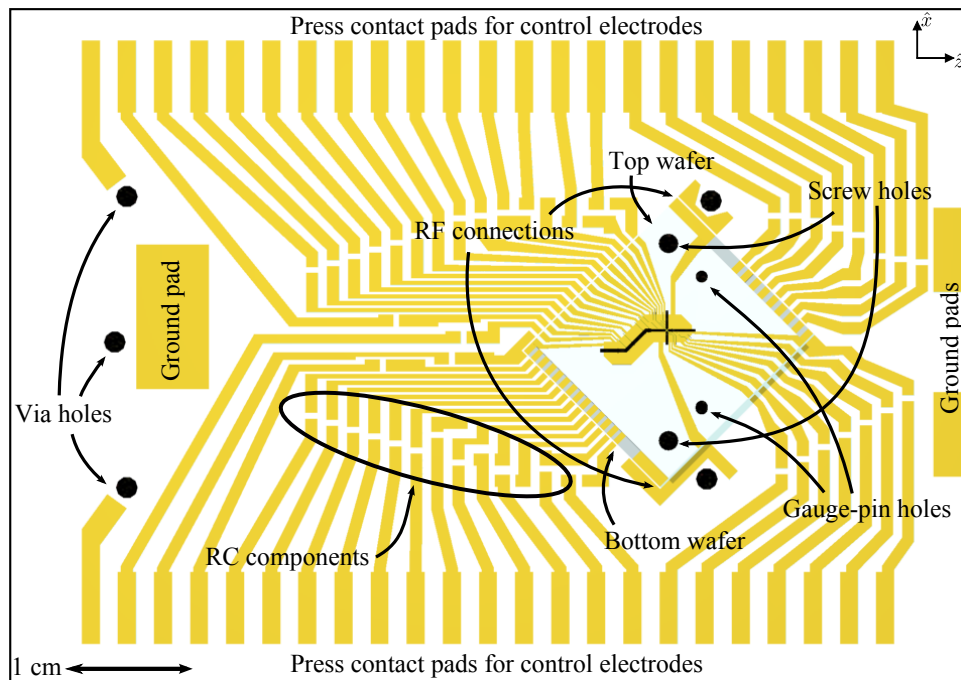


Figure 6.4: A diagram of the trap wafers and the filter board. The trap wafers are the ‘diamond’ just to the right of center. The bottom trap wafer extends 1 mm beyond the top wafer so that electrical connections can be made. Holes for screws, gauge pins (for alignment), and vias for wires to the backside can be seen. Each control-electrode trace has a gap and pad for a resistor and capacitor to form a low-pass filter (some of these locations are explicitly identified in the figure). The rf connections to the trap wafers are made at the top and bottom corners of the trap wafers and gold ribbons are fed through the holes next to these corners to contacts on the backside of the filter board.

Some electrical connections are made on the backside of the filter board, including the rf connections to the rf source and the connection to the bias wafer. Holes in the filter board act as vias, allowing gold ribbons to pass from the front of the filter board to the backside for these connections.

Many of the components mentioned above can be seen Fig. 6.5, which is a photograph of the assembled trap and filter board.

All the electrical connections were made with physical contacts, wire bonds, or resistive welds. No soldering, epoxies, or glues were used.




Figure 6.5: Photograph of the assembled trap. Three wires bonds were made per control-electrode trace, primarily for redundancy.

6.2 Trapping potential

The trapping potential consists of contributions from both the rf and control potentials. These two voltages are provided by very different systems: an rf resonant cavity and a series of digital-to-analog converters (DACs).

6.2.1 rf potential

The pseudopotential confinement in the X-junction array was created by an rf voltage of $V_{\text{rf}} \approx 200 \text{ V}$ and $\Omega_{\text{rf}} \approx 2\pi \times 83 \text{ MHz}$ applied to the rf electrodes. This large voltage is created with the aid of a quarter-wave step-up resonator [Jefferts 95]. The resonator is a coaxial design with a center rod (6.4 mm diameter) surrounded by an outer cylinder (inner diameter 3 cm). The volume in between the two conductors is vacuum. Both the center rod and outer conductor are constructed (primarily) from oxygen-free high-conductivity (OFHC) copper. The outer conductor, which is grounded to the vacuum system, can be seen in Fig. 6.11. The coaxial geometry acts as

a wave guide for a TEM-0,0 mode. The two conductors are electrically connected at one end by a conductive plane perpendicular to their shared axis. The other end of the inner conductor is free standing. This creates a resonant cavity where the wavelength λ is four times the length of center rod. A voltage node occurs at the shorted end of the center rod and an antinode at the other end. The large rf voltage required for trapping is achieved by attaching the trap to the free end of the rod.

The resonator is actually split into two sections by the wall of the vacuum system, with the majority located inside the vacuum but the shorted end of the resonator protruding into the air. The two sections of the center rod are connected to the two ends of an rf feed through, while the two sections of outer conductor bolt onto opposite sides of the conductive (steel) vacuum flange that holds the feed through.

Power is coupled into the resonator through a small wire loop at the closed end of the outer conductor (in air). By adjusting the size of the loop, critical coupling of $> 99\%$ can be achieved. When the loop is critically coupled to the resonator and the trap is not attached, the resonator exhibited a loaded Q of 789 with $\Omega_{\text{rf}} \approx 2\pi \times 107$ MHz. The filter board and trap are located at the antinode end of the center conductor with a gold ribbon connecting the rod to the rf leads on the filter board. The outer conducting tube was made longer than the inner conductor, which does not significantly affect the resonance frequency but allows the trap to be shielded within this conductor. Holes are cut in the outer conductor to allow laser access and a view of the trap for the light-collection optics.

With the trap attached, the resonant frequency dropped to $\Omega_{\text{rf}} \approx 2\pi \times 84$ MHz with a loaded $Q \approx 110$. When first attempting to load ${}^9\text{Be}^+$, these values further dropped to $\Omega_{\text{rf}} \approx 2\pi \times 83$ MHz and $Q = 42$. This additional drop is believed to be caused by neutral Be building up on the backside of the filter board to create a capacitor and leads to increased resistive losses. With this Q , achieving $V_{\text{rf}} \approx 200$ V requires about 2 W of rf power applied to the coupling loop.

In addition to producing the required large rf voltage, the resonator is also helpful for filtering noise on the rf line. As will be discussed in Sec. 7.6.1, noise on the rf electrodes at $\Omega_{\text{rf}} \pm \omega_z$ can lead to

excitation of the secular motion of the ion, especially when the ion is traversing the pseudopotential barriers near a junction. Since the secular frequency is large ($\omega_z = 2\pi \times 3.6$ MHz) the resonator will heavily filter noise at these two sidebands. However, with a relatively low Q , this resonator was not providing sufficient filtering and motional excitation due to this mechanism was observed during the transport experiments. In an effort to reduce this excitation, a second filter resonator tuned to Ω_{rf} was added in series.

This second resonator is similar to the first one except that it has a helical center conductor and is a half-wave ($\frac{\lambda}{2}$) resonator. The outer conductor is a copper tube and is sealed at both ends. The center conductor forms a helix that is coaxial with the outer conductor. Two small helical coupling coils are placed, coaxially, just inside the two ends of outer conductor. One coil couples power into the resonator and the other out of it. By adjusting the placement of these two coils along the axis, power can be efficiently passed through the resonator. A variable air-gap capacitor is attached between the anti-node location on the center conductor and the outer conductor. By adjusting that capacitance, the resonance can be tuned to better match the primary resonator. This ‘filter’ resonator has loaded $Q \approx 170$.

The output of the filter resonator is attached to the input of the primary resonator. The addition of a 3 dB attenuator between the resonators helped decouple them and allowed for improved overall filtering. The motional excitation of the ion during transport decreased by two orders-of-magnitude when this filter was introduced into the rf system. The noise at $\Omega_{\text{rf}} \pm \omega_z$ was estimated (using the ion’s heating rate) to be -177 dBc (see Sec. 7.6.1).

A schematic for the rf setup is seen in Fig. 6.6. The output of an Hewlett-Packard HP 8640 oscillator is amplified to drive the filter resonator. Approximately 15 W of power are needed at the input of the filter resonator to achieve $V_{\text{rf}} \approx 200$ V at the trap electrodes. It is important to have control over this voltage, which sets the pseudopotential amplitude, because the transports work best when the actual pseudopotential matches what is assumed when calculating the waveforms. The HP 8640 has an adjustable power level, but that adjustment does not provide sufficient precision to achieve optimal transport results. Instead the output is passed through a mixer which is used

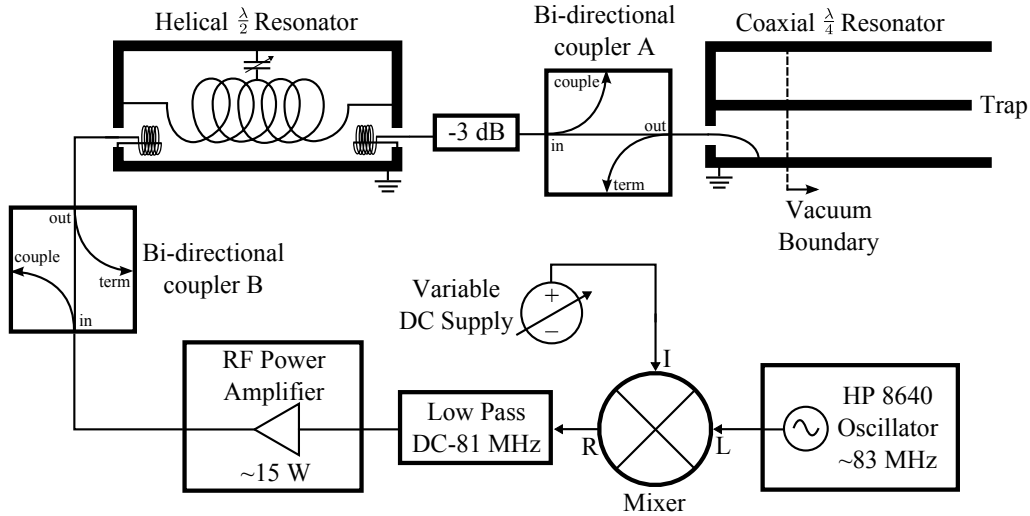


Figure 6.6: The rf setup consists of a coaxial $\frac{\lambda}{4}$ electrical resonator that produces a large voltage for the rf trap electrodes. A portion of this resonator is inside the vacuum system. A second filter resonator, which is a helical $\frac{\lambda}{2}$ design, is used to decrease the bandwidth of the filter for the rf to reduce motional excitation rates during transport. A 3 dB attenuator decouples the two resonators. The rf signal is created by an amplified HP 8640 oscillator. The rf power can be adjusted using a mixer. Two bi-directional couplers allow the incident and reflected powers to be measured so that the coupling into the resonators can be tuned. There is also a low pass filter to reduce higher-frequency noise coming from the source electronics.

as a variable attenuator by adjusting the DC voltage applied to the ‘I’ port of the mixer. This mixer is run above saturation (on the ‘I’ port) to make the attenuation less sensitive to the DC voltage and provide better control. This allows the power to be easily set to within 1% at the trap (which is measured by observing the radial secular frequencies).² Having accurate control over the rf amplitude is important because temperature changes cause the resonator parameters and the power levels to drift over several minutes, and the pseudopotential at the ion drifts accordingly. For best transport results, this power levels was adjusted every few minutes to ensure the secular frequencies did not drift.

Two bi-directional couplers have been inserted into the rf setup. These allow the incident and reflected rf power to be measured before both resonators to help achieve optimal coupling. Bi-directional coupler A can also be used to inject a small signal onto the rf electrodes for diagnostic

² Measuring the rf power at any point upstream from the ion does not provide an accurate measure of the pseudopotential amplitude because it will not take into account variations in the coupling or Q of the primary resonator.

tests. For example, a signal at $\Omega_{\text{rf}} \pm \omega$ can be injected to measure the ion's secular frequencies by coherently exciting that motion.

6.2.2 Control potentials

The control potentials are applied to the electrodes using digital-to-analog converters (DACs). Five National Instruments PXI-6733 DAC boards, each with eight channels (a total of 40 independent DACs), are used. These are 16-bit DACs that have a voltage range of ± 10 V, giving them 0.3 mV precision. When a waveform is being run, the DAC outputs are changed in discrete steps at a fixed time interval set by the update rate U_{DAC} , which can be set anywhere up to 500 kHz.

The X-junction array includes 46 control electrodes and one bias board, but there were only 40 available DAC channels. For this reason, several of the electrodes were shorted together so that only 40 independent channels were needed. The shorted electrode pairs were: 1/2, 5/14, 6/13, 7/8, 11/12, 31/32, and 37/38.

The ion is subject to slowly changing stray electric fields due to charges that build up on insulating surfaces. To null these fields and keep the ion on the trap axis, a small differential voltage can be applied to opposing electrodes, such as 17 and 18 at location \mathcal{E} , which will push the ion approximately in the $\hat{x} - \hat{y}$ direction. A voltage is applied to the bias wafer, which sits on the other side of the filter board, to produce a field that pushes roughly along \hat{y} . A combination of these 'shim' voltages will allow complete nulling of the stray field. This is most important in the region where data is taken, so, in practice, this nulling is usually only performed in \mathcal{E} .

The DAC outputs are filtered prior to being applied to the control electrodes to reduce electronic noise at the ion secular frequencies. There were two stages of filtering and they are shown in Fig. 6.7(a). An 'internal' set of RC filters, one per electrode, is located within vacuum on the filter board, as close to the electrodes as possible. An external set of filters is attached outside of the vacuum system. All control voltages (and the rf) are referenced to the grounded vacuum system.

The internal filters consist of a 240 Ω , 820 pF low-pass RC filter, which has a corner frequency

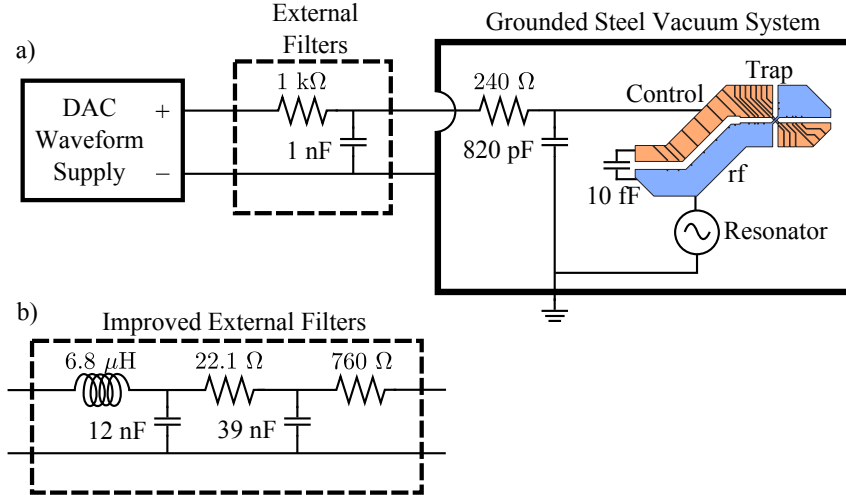


Figure 6.7: The control voltages are provided by 40 independent DACs (only one shown here). The DAC output is filtered prior to being applied to the trap electrodes, through a two stage filter seen in a). The control potentials are referenced to the grounded vacuum system, which serves as the rf ground as well. b) After DAC-update noise was observed to excite the secular motion of the ions, the external filters were replaced by the approximately-3rd-order ‘Butterworth’ filter shown here.

of 810 kHz. This uses vacuum compatible surface-mount components that are attached with gold ribbon that is resistively welded on one end to component and on the other end to the filter board. In addition to filtering the DAC output, the internal filters play another important role: they provide an rf short to ground for each control electrode. Each control electrode has a capacitance to the rf electrode of ~ 10 fF, so if not shorted to ground, these control electrodes will float at some finite rf potential. This would negatively affect the rf confinement and could lead to excess micromotion or other problems. The RC filter circuit is placed close to the trap electrodes to provide a route to ground for any capacitively coupled rf.

The standard low-pass RC filter has a frequency response profile given by

$$G(\omega) = \frac{1}{\sqrt{1 + (\omega RC)^2}} = \frac{1}{\sqrt{1 + (\omega/\omega_0)^2}}, \quad (6.1)$$

where ω_0 is the corner frequency. At $\omega = 0$, the gain is 1, and as the frequency exceeds ω_0 , the gain drops to zero as $G(\omega) \propto \omega^{-1}$ for $\omega \gg \omega_0$.

The original external filters were low-pass RC filters with 1 k Ω and 1 nF to give a 160 kHz corner frequency. However, when combined with the internal filters, the overall filter system behaves




Figure 6.8: This is a generic filter network. Each element is a resistor, capacitor, or inductor. By tuning the relative impedances, Z_i , different types of filter responses can be created.

as though it has a 83 kHz corner frequency. From 83 kHz to 1 MHz, the frequency response was $G(\omega) \propto \omega^{-1}$, but above 1 MHz this increased to $G(\omega) \propto \omega^{-2}$.

Though these filters provide two orders-of-magnitude suppression at the ω_z , noise due to the DAC output being discretized at U_{DAC} was sufficient to excite the ion's motion during transport. This mechanism is discussed in detail in Sec. 7.6.2. To combat this problem a new filter set was devised. Using a filter with a lower corner frequency could help, but would also limit how quickly the ions could be moved. Instead, it was decided to replace these simple RC filters with a 3rd-order Butterworth filter [Tietze 08].

The goal in using a new type of filter is to reduce $G(\omega)$ for $\omega > \omega_0$, but without reducing $G(\omega_0)$. This means the filter should cutoff faster as ω increases above ω_0 . A 'Butterworth' filter can accomplish this. The term Butterworth refers to the type of filter response and not the type or structure of the filter. A Butterworth filter has frequency response

$$G(\omega) = \frac{1}{|B_n(i\omega/\omega_0)|} = \frac{1}{\sqrt{1 + (\omega/\omega_0)^{2n}}}, \quad (6.2)$$

where $B_n(s)$ is the n^{th} -order Butterworth polynomial [Tietze 08]. If $n = 1$, the frequency response reduces to the RC frequency response Eq. (6.1). A 2nd-order Butterworth filter is difficult to physically realize. So we will jump directly to the 3rd-order case $n = 3$, where $B_3(s) = (1 + 2s + 2s^2 + s^3)$. This has a sharper cutoff than the standard RC filter while still having the appropriate ω_0 . The next step is to determine a filter network that will produce this desired frequency response.

Figure 6.8 shows a generic filter network where the different components, Z_i , are resistive, capacitive, or inductive elements. In the basic low pass RC filter, $Z_1 = R_1$ would be a resistor and $Z_2 = 1/i\omega C_2$ would be a capacitor, while the components are chosen such that $R_1 C_2 = 1/\omega_0$. The

rest of the odd components would be electrical shorts and the even components would be empty. To make a 3rd-order Butterworth filter, $Z_1 = i\omega L_1$ is an inductive element, $Z_3 = R_3$ is a resistor, while $Z_2 = C_2$ and $Z_4 = C_4$ are capacitors. The rest of the elements are either shorts or gaps, as before. The value of the circuit elements are then chosen such that the frequency response corresponds to Eq. (6.2). The appropriate values for the Butterworth filter are

$$\begin{aligned} R_3 C_4 &= 2/\omega_0 \\ C_2 &= C_4/3 \\ L_1 &= \frac{3}{8} R_3^2 C_4. \end{aligned} \tag{6.3}$$

Other values can be used, but will not result in the desired filter response Eq. (6.2) and would not be a proper ‘Butterworth’ filter. Also, circuit configurations other than the one shown in Fig. 6.8 could be used to achieve the same filter response.

The ‘order’ of a filter refers to the leading order in ω of the filter response. Generally the higher the order, the sharper the cutoff. The 3rd-order Butterworth filter has an ω^{-3} response. The standard RC filter is a 1st-order filter. By including more filter blocks in Fig. 6.8, filters with higher order can be created. In fact, the order can easily be calculated for Fig. 6.8 by counting the number of capacitors and inductors, as each introduces another term ω into the filter response. Resistive elements, however, do not increase the filter order.

The Butterworth filter response is not the only interesting filter. By adjusting the relative impedances of the components, a number of other standard filters can be produced with arbitrary order [Tietze 08]. Figure 6.9 depicts some commonly used filter responses. A 1st- and 3rd-order RC filter are shown in Fig. 6.9(a), where the specific 3rd-order RC filter depicted is for $Z_1 C_2 = Z_3 C_4 = Z_5 C_6 = 1/\omega_0$. Figure 6.9(b) is the 3rd-order Butterworth. A Bessel filter is similar to the Butterworth, but has a less sharp cutoff. What makes the Bessel filter interesting is that it has a flat group delay in the pass-band. The other filters will be dispersive at different frequencies. Two different Chebyshev filters are possible, both of which have a sharper cutoff at ω_0 than the previous filters. However, the trade off for the sharper cutoff is a ripple in $G(\omega)$ either below (type 1) or above (type 2) ω_0 (see Fig. 6.9(d) and (e)). The elliptic filter achieves an even sharper cutoff but

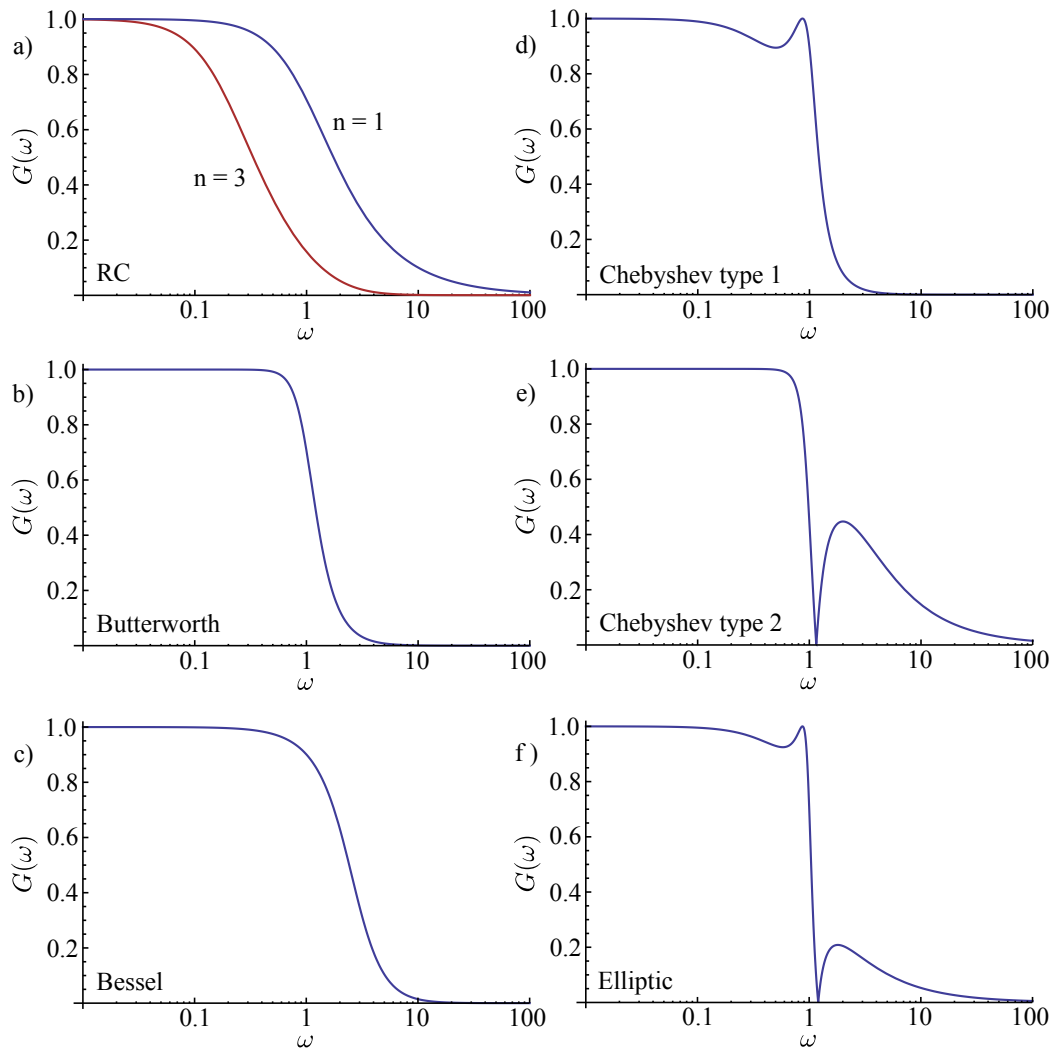


Figure 6.9: The gain $G(\omega)$ versus frequency ω for six different filter functions. These are all normalized to a corner frequency of $\omega_0 = 1$. a) An RC filter. Both a 1st-order and a 3rd-order RC filter are shown. b) A 3rd-order Butterworth filter. c) A 3rd-order Bessel filter. d) A 3rd-order Chebyshev type 1 filter. e) A 3rd-order Chebyshev type 2 filter. f) A 3rd-order elliptic filter.

has ripples in both locations. One distinguishing characteristic of the Butterworth filter is that it has the flattest possible gain $G(\omega)$ over the pass-band while still having a reasonably sharp cutoff.

For the purpose at hand, the Butterworth filter seems the most appropriate. The ripples in the Chebyshev and elliptic filters could lead to deformation of the waveform or, worse, passing of the electronic noise at the secular frequency, making them unattractive for filtering the DACs. For the time being, the sharper cutoff of the Butterworth filter is more important than obtaining




Figure 6.10: The gain $G(\omega)$ versus frequency for the old RC filter and the new approximately-‘Butterworth’ filter, where the new filter has a faster roll off at high frequency. Since the pass-band extends farther for the new filter, the transport speed can be increased while still providing more filtering at the secular frequency (3.6 MHz to 5.7 MHz).

the Bessel filter’s flat group delay. Waveform dispersion is expected to become more of an issue at higher transport speeds, at which point a Bessel filter may become more ideal. However, the current experiments in the X-junction array are not in that regime.

The external RC filter was replaced by what is shown in Fig. 6.7(b), which was designed to have the same corner frequency as the previous filters, $\omega_0 = 2\pi \times 160$ kHz. When the internal filter is included, this is technically a 4th-order filter. However, the components were designed so that Z_1 , Z_2 , Z_3 , and Z_4 formed an ideal 3rd-order Butterworth filter while the internal components Z_5 and Z_6 (which could not easily be changed) minimally contribute to the filter’s response. Attempting to design a 4th-order Butterworth filter that incorporates the internal capacitor, $C_6 = 820$ pF required an excessively large inductor at Z_1 to achieve the appropriate response. Instead, the filter is designed so that $C_4 \gg C_6$, which makes the second capacitor only a small perturbation to the overall filter. The internal resistor R_5 was effectively increased from 240Ω to $1 \text{ k}\Omega$ by adding a second resistor in series outside of the vacuum system, which improved the overall filtering without significantly affecting ω_0 .

The new composite filters increase the noise filtering over the old filters by a factor of 146

at 3.6 MHz and 380 at 5.7 MHz, as can be seen in Fig. 6.10. In addition, the increased pass-band will allow for faster transport and the ion was observed to settle within $3 \mu\text{s}$ after transport. The electric-field noise at the ion due to Johnson noise in the resistive elements of the new filters is less than that for the old filters for all frequencies. These new filters succeeded in improving the problem DAC-induced motional excitation (see Sec. 7.6.2).

6.3 Vacuum system

The assembled trap is placed inside of a vacuum system with $p < 5 \times 10^{-11}$ Torr = 7×10^{-9} Pa to minimize motional excitation and ion loss caused by collisions with background gas molecules. The vacuum system is pictured in Fig. 6.11 and is based around standard stainless steel Conflat vacuum parts. Many of the parts were baked, in air, to 400°C prior to being assembled to remove excess hydrogen inside the metal [Bernardini 98]. In addition to these standard parts, the system includes an ion pump and titanium sublimation pump (TSP). There is also an ion gauge and a series of electrical feed throughs to pass the signals to the electrodes and Be/Mg evaporation ovens.

The vacuum system was temporarily connected to a large external pump and then baked, under vacuum, to 250°C until the pressure stabilized (~ 1 week). Once the system cooled to room temperature, the external pump was valved off and removed. The internal ion pump was started and remains on indefinitely. The TSP was run several times after the bake to lay down a layer of Ti, which acts as a getter, on the internal walls of the vacuum system. This layer of Ti is refreshed every few weeks by running the TSP again. There is not a direct line of sight between the TSP and the ion trap so the layer of Ti does not extend to the trap.

The components inside of the vacuum system are intended for ultra-high vacuum (UHV) compatibility. The internal wiring consists of unshielded, Kapton-wrapped copper wires. All electrical contacts are either press contacts or resistive welds. The rest of the components are primarily copper, steel, aluminum, titanium, gold, fused-silica, alumina, macor, Kapton, or PEEK. These have all previously been used in UHV systems. No solders, glues, or epoxies are present inside the vacuum. Internal screws are vacuum vented where possible to reduce virtual leaks (either by a hole




Figure 6.11: Vacuum system used for the X-junction array. The trap is mounted in a copper rf resonator that provides the rf voltage. The other end of the resonator is bolted to a feed through (not visible) inside of a 6" Conflat cube. A glass envelope (in background but covered with Al foil) fits over the resonator and bolts to the vacuum system to allow the trap and resonator to be evacuated. Both the control and oven feed throughs are visible and are used to pass the control and oven signals to the trap, via Kapton-wrapped copper wires. An ion pump and TSP are used to reach $p < 5 \times 10^{-11}$ Torr, which can be measured with an ion gauge.

down screw axis or slot on the side through the threads).

The trap assembly is positioned at the rf antinode end of the copper resonator. The trap sits vertically in the resonator and there are holes cut into both sides of the resonator for optical access as seen in Fig. 6.12. The filter board is wider than the resonator tube and extends both above and below the resonator. The electrical connections between the filter board and the wires carrying the control voltages are made on the sections of the filter board that protrude beyond the resonator.

The end of the outer conductor of the resonator (beyond the end of the center rod) is cut into two pieces to form a half-cylinder section that can be removed from the rest of the resonator (see Fig. 6.16). The trap, the filter board, the loading ovens, and the control-voltage connections are all attached to this removable piece for easy access.




Figure 6.12: The trap as viewed through a hole cut in the copper resonator. The circular part of the hole is used for imaging the ion while the elongated side-lobes allow lasers to pass at 45° to the plane of the trap wafer. Kapton-wrapped copper wires, which carry the control voltages to the electrodes, can be seen above and below the resonator.

To complete the vacuum system, a large fused-silica tube, with a glass-to-metal seal on one end, is placed over the resonator/trap assembly and bolted to the steel vacuum system. A portion of the envelope can be seen in Fig. 6.13. Optical-quality fused-silica windows positioned on tubes at 45° angles allow laser access while a larger window is used by the imaging system. The envelope was designed to include a large flat section of glass that runs parallel to the envelope axis. This allows the imaging window to be positioned closer to the trap than it would be if attached directly to the main envelope tube. At the same time, the tube diameter in the vertical direction remains large enough to fit around the filter board. Figure 6.14 shows how the trap and resonator assembly fits inside of the glass envelope.

To achieve the best vacuum, a rigorous cleaning regimen was employed for all parts. The parts were first washed in Alconox, followed by a rinse in de-ionized (DI) water. The parts were then submerged in acetone while ultrasound was applied. This was rapidly followed by a rinse




Figure 6.13: Glass vacuum envelope seen from two angles. The large 1.5" window in the center is used for imaging the ions. The two smaller windows, at 45° , next to the imaging window, as well as the two windows on the other side of the envelope, allow the lasers to pass through the vacuum and slots in the trap assembly. The final small window can be used to pass a beam into the vertical leg of the X-junction. The side of the envelope with the imaging window was replaced with a glass flat to allow the imaging window to be placed close to the trap so that more ion fluorescence could be captured by the imaging system.




Figure 6.14: Figure showing the relative orientation of the trap, resonator, and glass envelope. The large window in the envelope is used for imaging ions in the trap while the smaller windows allow laser beams to pass through the trap.

in acetone, then isopropynol, and finally the parts were blown dry using clean, dry nitrogen. In between steps, parts were not allowed to dry in an attempt to keep material from condensing or collecting on the part. Particularly fragile or sensitive parts, such as the trap and the glass envelope, were not subjected to the entire cleaning protocol (especially the ultrasound) to keep from damaging the part. The vacuum bake at 250°C was used to remove residual contaminants from the system.

6.4 Loading

The loading procedure starts by locally heating a Be or Mg ‘oven’ with an electric current so that neutral atoms evaporate from the oven. These atoms then ballistically pass through the trapping potential, where they can be photo-ionized using a laser. Once ionized and trapped, the ions interact with the Doppler cooling laser and lose kinetic energy.

The Be ovens consist of 0.05 mm Be wire that is tightly wound around a 0.1 mm diameter tungsten wire (this is performed with the aid of a heat gun to warm the Be and make it more malleable). The tungsten is then wrapped into a ~ 2 mm diameter helix and bonded to an electrical contact. A current of ~ 0.8 A is passed through the oven, which causes it to glow weakly (dull orange visible to the human eye) and produces sufficient neutral Be flux to load ions.

The Mg ovens, which are described in Fig. 6.15, consist of a 0.5 mm long, 1.27 mm outer-diameter 316 stainless-steel tube (0.15 mm thick walls), with a small hole cut in the side. The tube is packed with Mg shavings and crimped at both ends. Tantalum wires are resistively welded to each end. A current of ~ 1.8 A will produce Mg flux and cause the oven to glow only very faintly when viewed through an IR viewer.

The ovens are placed in an oven assembly that is attached to the exterior of the removable portion of the resonator, as seen in Fig. 6.16. Two Be and two Mg ovens are connected so that they can be independently heated. An aluminum structure with four pockets fits over the ovens to keep them from depositing material on each other. The filter board and trap are also attached to this section of the resonator in such a way that the ovens have a direct line-of-sight to the load




Figure 6.15: Steps for creating a Mg oven. a) A 316 stainless steel tube is cut to length. b) A small slit is made in the side of the tube to create a small hole. c) One end of the tube is crimped and sealed with pliers. Mg shavings are then packed into the tube from the open end. d) The other tube end is crimped and sealed. e) Tantalum wires are bonded to the tube so that current can be passed through the tube. When heated, evaporated Mg will exit the small hole.

zone. The oven flux passes through the trap channel at an angle of $\sim 30^\circ$ to the normal of the trap wafer. As describe in Sec. 4.2.3, a shield blocks the experiment region of the trap from oven flux.

In addition to the ovens, the oven assembly also includes two electron guns (e-guns). These can be used to impact-ionize both Be and Mg. However, they are not used because the photo-ionization beam performs the same function without producing as many unwanted electrons that can charge up dielectric surfaces and lead to slow drifts in the trapping potential as these surfaces discharge. The e-guns also need to be run at high temperature, which will temporarily spoil the vacuum and reduce ion lifetime.

Photo-ionization of Be is accomplished with a two-photon process. The neutral Be is resonantly excited from $2s^2$ to $2s2p$ using a photon with $\lambda = 235$ nm. A second photon of equal energy can then excite to continuum. Alternately, a BD or BDD photon at 313 nm is also sufficient to excite from $2s2p$ to the continuum.

The 235 nm light is produced by a frequency-quadrupled Ti:Sapph laser. A Coherent Verdi V10 with 10 W of 532 nm light is used to pump a Ti:Sapph crystal and produce pulses of 940 nm

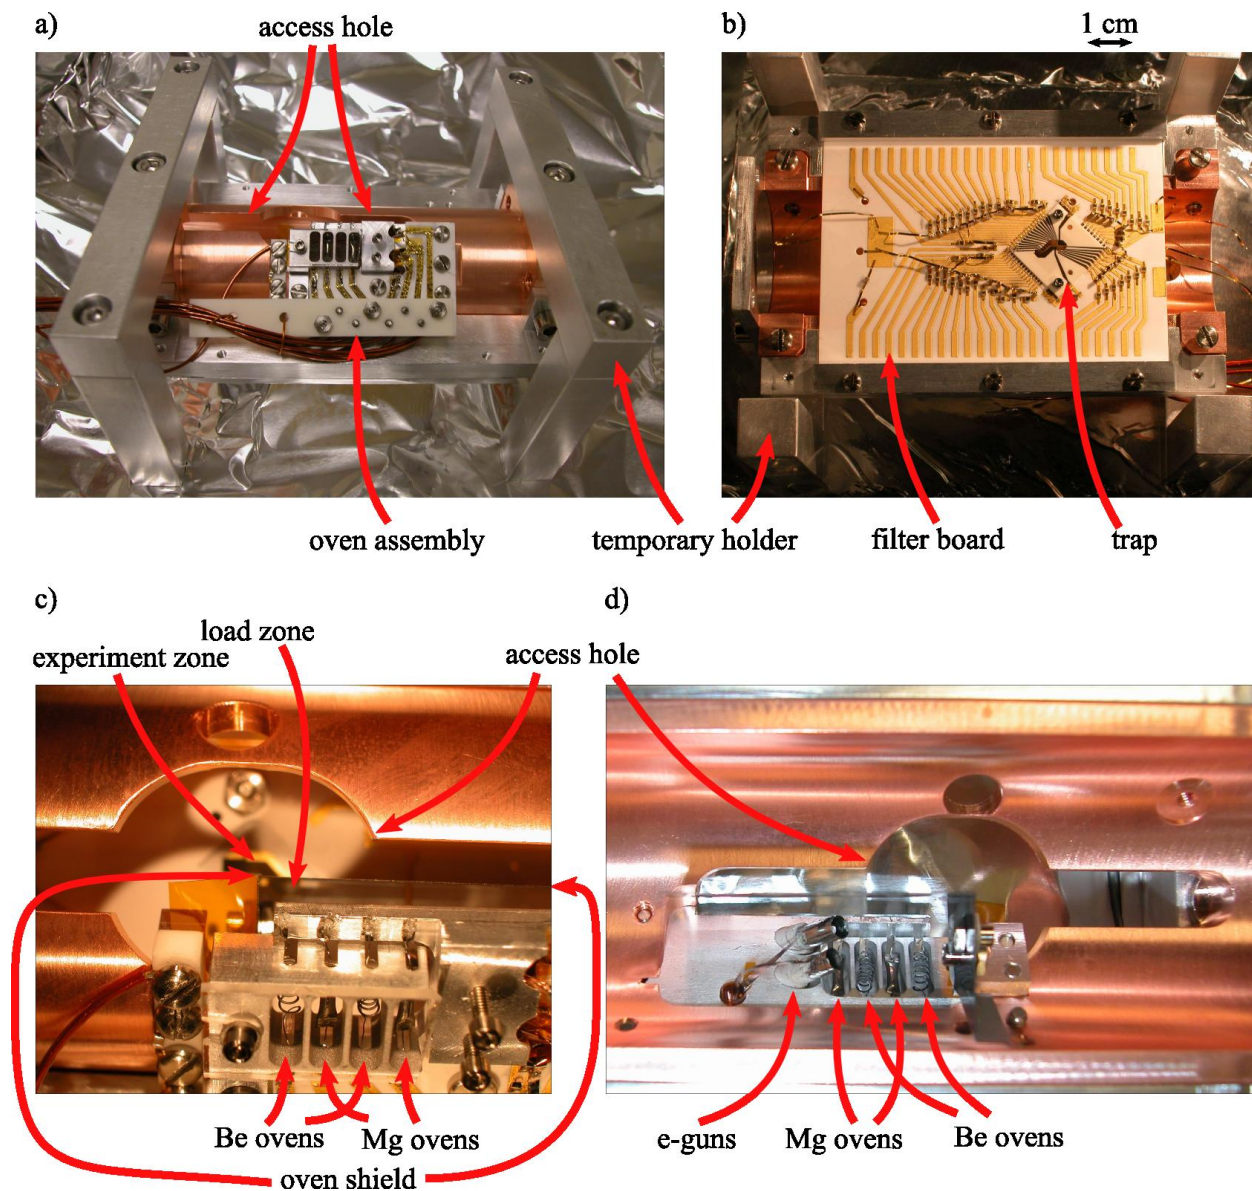


Figure 6.16: The end of the copper resonator, which can be removed from the resonator. This piece is half of a hollow cylinder, and has an access hole identical to the one seen in Fig. 6.12. a) The removable piece is seen here supported by a temporary aluminum stand. The oven assembly attaches to the outside of this piece at the edge of the access hole. b) The resonator piece is flipped over to reveal the filter board and trap, which are attached to this piece using aluminum bars. c) Zooming in on the oven assembly, it is possible to see the two Be and two Mg ovens as well as the trap. The ovens spray down onto the trap. The oven shield, which sits parallel to the viewer's line-of-sight in the picture, is barely visible. This shield allows oven flux to reach the trap's load zone while shadowing the experiment zone from the ovens. d) Flipping the part over and removing the trap and filter board, it is possible to see the backside of the oven assembly, as well as the two electron guns (that were not used).

light with a 100 MHz repetition rate. This is on the edge of the gain profile for Ti:Sapph, so only 0.6 W of averaged power is created at 940 nm, despite the large pump power. In addition, there are a number of water absorption lines near 940 nm and it is advantageous to purge the laser cavity with dry nitrogen.

The fourth harmonic of the 940 nm light is produced by a single pass through two BBO-crystals, which doubles the frequency with each pass. The high intensity pulses allow a reasonable conversion efficiency and 20 mW (average) of 235 nm light is created.

During the load process, this photo-ionization (PI) laser beam, along with the BD and/or BDD beams, are switched on and passed through the load zone of the trap. At the same time the current in the oven is switched on. An ion typically will load within a minute. Once an ion is loaded, it will scatter from the Doppler beams and can be observed using the camera. The PI light and oven currents are switched off at that point. The ion can then be transported from \mathcal{L} to the experiment zone \mathcal{E} .

6.5 Lasers and optics

Controlling the atomic state of ${}^9\text{Be}^+$ requires several ultraviolet (UV) laser wavelengths all near 313 nm. The BD and BDD beams are used for Doppler cooling and BD is also used for state detection. Collectively, these are referred to as the 'blue Doppler' beams. Another two wavelengths, RD1 and RD2, are used for optical pumping and are known as the 'red Doppler' beams. Ground state cooling and gate operations are performed with the Raman 90° and Raman co-carrier (CC) beams. And, finally, the PI beam is used for photo-ionization.

Each beam must be passed through the experiment zone \mathcal{E} and/or the load zone \mathcal{L} , as seen in Fig. 6.17. The BD, BDD, RD1, and RD2 light is arranged so that all beams copropagate along what is referred to as the Doppler beam line, though each of the four beams can be switched on and off independently. This light is passed through a beam splitter to create two beams (of equal power). One beam is sent to \mathcal{E} along $-\hat{y} + \hat{z}$ while the other beam goes to \mathcal{L} parallel to $-\hat{y} - \hat{z}$. The RD light is not needed or used in \mathcal{L} , so is not labeled in the figure along that beam. The PI




Figure 6.17: Laser beam orientations relative to the trap. The trap is depicted in both a front view (a) and a top view (b) (as seen in the lab). The PI beam and a combined BD/BDD beam both pass through the load zone along $-\hat{y} - \hat{z}$. The combined BD/RD/RD1/RD2 beam passes through the experiment zone along $-\hat{y} + \hat{z}$. One of the Raman 90° beams counter-propagates along the experiment BD beam line while the other Raman 90° beam and the Raman CC beams propagate along $+\hat{y} + \hat{z}$. A B -field, along $+\hat{y} - \hat{z}$, is applied to the entire trap to define an quantization axis for the ion.

beam is overlapped with the Doppler beam line at \mathcal{L} . The Raman CC and Raman 90° both consist of two beams with different frequencies. In the CC beam, the two beams are copropagating and pass through \mathcal{E} parallel to $\hat{y} + \hat{z}$. For the 90° beams, the beams of different frequency are split into two separate beams. One is roughly parallel to the CC beam while the other is at 90° and parallel to $\hat{y} - \hat{z}$.

A magnetic field of 13 G is applied to the entire trap array, using field coils, to provide a quantization axis for the ion. This field lies along $\hat{y} - \hat{z}$ and is parallel to the BD at \mathcal{E} (not at \mathcal{L}) and one of the Raman 90° beams.

An optical system, consisting of a compound lens and a camera/PMT can be positioned to view ions in either \mathcal{E} or \mathcal{L} , as well as any other location along the trap axis. This section will describe these laser and optical systems.

6.5.1 Doppler lasers

The BD and BDD light is derived from a single dye laser while the RD1 and RD2 light is derived from a second dye laser. The setup for both of these Doppler lasers is depicted in Fig. 6.18.

The blue Doppler laser uses Exciton Inc.'s Kyton Red laser dye, pumped with 4 W of 532 nm light from a Coherent Verdi laser, to produce 0.5 W of red light at 626.266 nm. This laser is locked to a cavity using a Hänsch-Couillaud lock [Hansch 80] for frequency stability. An acousto-optic modulator (AOM) picks off a fraction of the light, shifted by -351 MHz, and sends it to a setup for saturated-absorption spectroscopy, which allows the laser to be locked to a feature in the I_2 spectrum. The majority of the laser output is passed through an optical isolator and then coupled into an optical fiber. Once the light emerges from the fiber, it is frequency doubled in an optical cavity using a BBO crystal, resulting in 30 mW of 313.133 nm UV light. Lenses are used to collimate the beam as it exits the doubling cavity.

The blue Doppler UV laser intensity is stabilized using AOM 1, which deflects a portion of the main beam and sends it to beam dump. A glass plate picks off 5% of the main beam to measure the intensity and this is fed back onto the rf power applied to AOM 1, which deflects more or less power to ensure a constant power in the main beam.

Next, the blue Doppler light passes through two double-pass AOM setups. The first of these AOMs, AOM 2, deflects a portion of the main beam by adding 221 MHz. This will become the BD light. This deflected beam is focused onto right-angle retroreflecting mirror (RR) and passed a second time through AOM 2, where it picks up an additional shift of 221 MHz, for a total of 442 MHz. The retroreflected beam is displaced vertically from the incident beam so it can be picked off with a mirror after the double-pass setup. The AOM deflection efficiency can be controlled using the rf power supplied to AOM 2, to adjust the light power in the BD beam. When the rf power is switched off completely, the BD light will vanish.

The BDD beam is formed in a similar manner using AOM 3 in double-pass configuration to give a -160 MHz (total) shift to the BDD beam. AOM 3 is setup to deflect light from the main

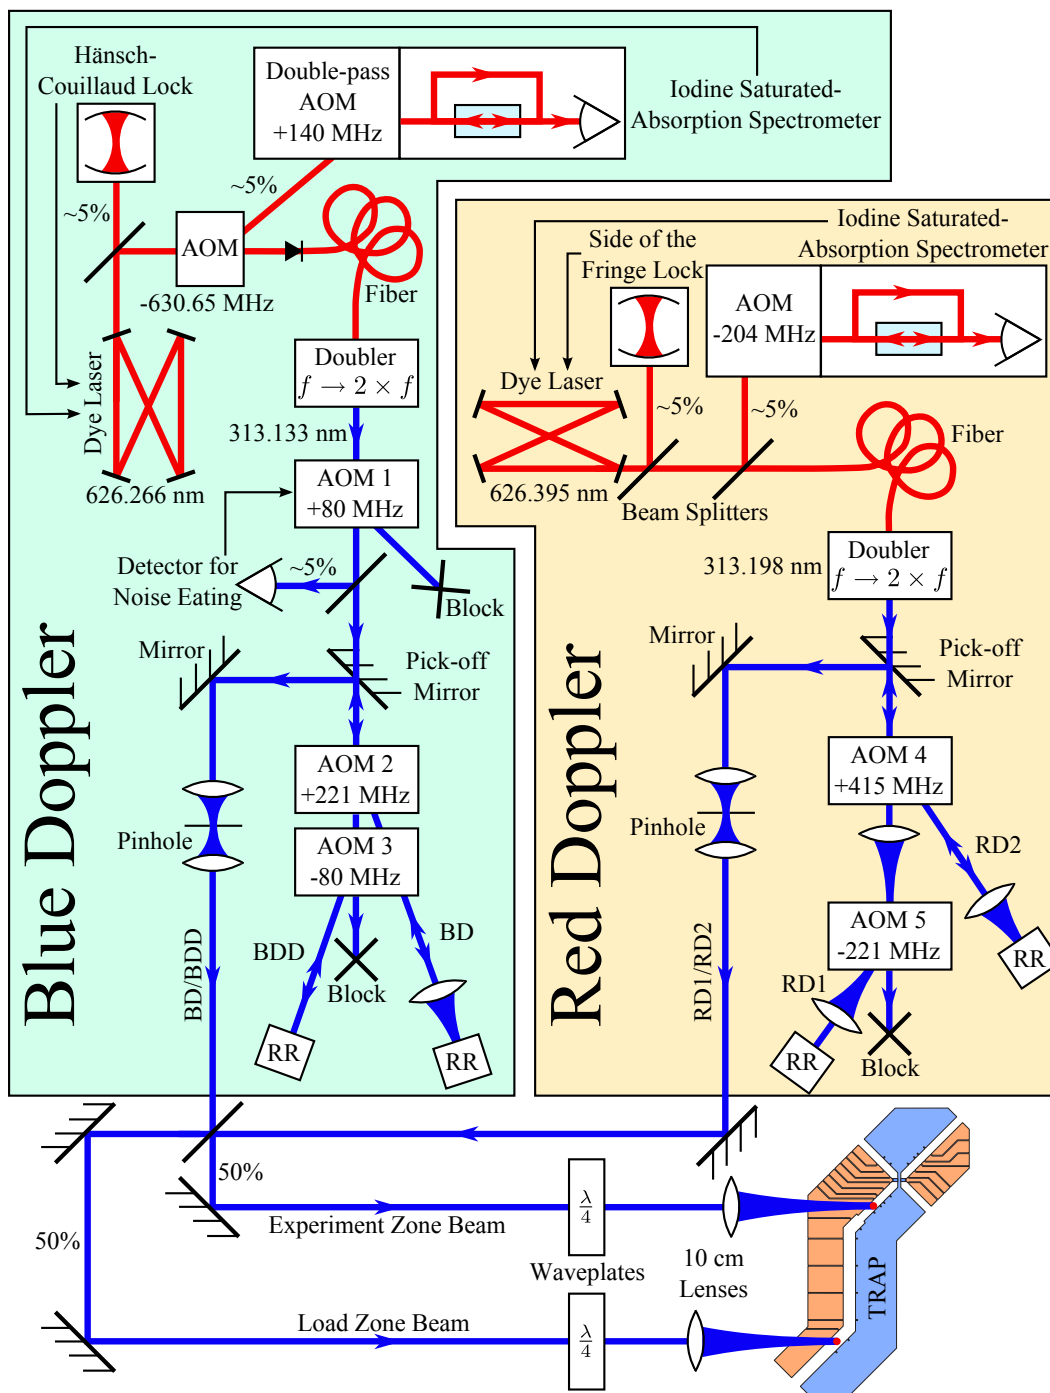


Figure 6.18: A schematic of both the blue Doppler and red Doppler beam lines (not to scale). The BD beam can be switched on and off by switching the rf power to AOM 2. The BDD beam is switched using AOM 3. Similarly AOM 4 and AOM 5 switch RD2 and RD1, respectively. If the power is switched off to these four AOMs, no light from this setup will arrive at the trap. If power is being applied to any of the four AOMs, the resulting light will be split between the load and experiment zones.

beam after it passes through AOM 2, but the third AOM does not affect the BD light deflected by AOM 2. Again, the power in BDD can be controlled by adjusting the rf power supplied to AOM 3. Instead of focusing onto a right-angle RR as is done with the BD beam, the BDD remains collimated but a RR with an adjustable angle between the mirrors is used.

These AOM frequencies will produce a BD beam which is 5 MHz red of the ${}^9\text{Be}^+$ cycling transition and a BDD beam that is 607 MHz to the red of the same transition. The BD and BDD beams are aligned so that they copropagate once they emerge from the double-pass AOM setup. The BD/BDD beam is passed through a pinhole to clean the TEM-0,0 mode. It is then split on a 50/50 beam splitter and sent along two paths to \mathcal{E} and \mathcal{L} . The beams are focused to 30 μm spots in the trap using 10 cm lenses.

Through the majority of the setup, the BD/BDD light is vertically polarized but σ^- polarization is needed at the ion. To change the polarization, $\frac{\lambda}{4}$ wave plates are placed immediately before the focusing lenses. The orientation of these plates are fine-tuned to account for any birefringence in the glass envelope windows.

The RD1 and RD2 beams are created the red Doppler beam setup. This is very similar to the blue Doppler setup, with a Kyton Red dye laser locked to a reference cavity and I_2 line. This laser is tuned to 626.395 nm and produces 300 mW of red light. An AOM is used to shift the light sent to an iodine spectroscopy setup by -204 MHz and the light is locked to the same iodine feature as blue Doppler. RD2 is switched on and off using AOM 4 in double-pass configuration to give a 830 MHz shift to the light. RD1 is independently switched using AOM 5 to give a -442 MHz shift, so that the difference between RD1 and RD2 corresponds to the 1.272 GHz between the ${}^2S_{1/2} |F = 2, m_F = -1\rangle$ and ${}^2S_{1/2} |F = 1, m_F = -1\rangle$ atomic states. The RD1 and RD2 beams are also aligned to emerge copropagating from the double-pass setups.

The BD/BDD and RD1/RD2 beams are carefully overlapped on the 50/50 beam splitter so that two overlapped beams emerge, each potentially containing light from any or all of these frequencies. By switching the rf power to the four AOMs, the user has control over what light is incident on the ion in the trap. The light going to both zones \mathcal{E} and \mathcal{L} will be identical.

If light is needed in other zones of the trap (besides \mathcal{E} and \mathcal{L}), the beams can be collected after passing through the trap and directed back through a second location, though that is not the standard configuration.

6.5.2 Raman lasers

The Raman laser beams are tuned to two frequencies split by ~ 1260 MHz so that they can couple the two qubit states. To accommodate the qubit transition, which is $\Delta m_F = 1$, one of the beams is π polarized while the other is σ^+ . An equal amount of σ^- light is included with the second frequency to minimize Stark shifts [Wineland 03], though this additional light does not participate in any atomic transitions due to selection rules.

There are two sets of Raman beams, the co-carrier (CC) and the 90° beams. The CC beams are aligned so they copropagate, making the CC insensitive to the ion's motion. For the 90° beams, the two frequencies are split into separate beams and aligned so that they impinge on the ion at 90° and $\Delta \mathbf{k}$ points along the axial mode of motion. This geometrical difference (and the associated motional sensitivity) is the main characteristic that distinguishes the CC from the 90° beams.

The CC and 90° beams originate from the same laser setup, which is depicted in Fig. 6.19. Two Koheras infrared fiber lasers at 1550 nm and 1051 nm are combined using sum-frequency generation in a periodically poled LiNbO₃ crystal to form 626.340 nm light. This is doubled in a resonant cavity using a BBO crystal to produce 400 to 500 mW of 313.170 nm UV light.

This light is passed through two AOMs both of which shift the frequency by 210 MHz. One of the AOMs is used to switch the CC beams on and off, while the other controls the 90° beams. This is aligned so that the two resulting beam paths emerge parallel and offset by 1 mm from the two AOMs. However, only one of the two beams is switched on at any given time. For the time being, we will concentrate on one of these paths, the CC path.

A waveplate is used to create a mixture of vertical and horizontal polarization in the CC beam before the light enters the 'quad-pass' AOM setup. This quad-pass setup is actually two identical (but inverted) 315 MHz AOMs, followed by a mirror, aligned so that the beam passes

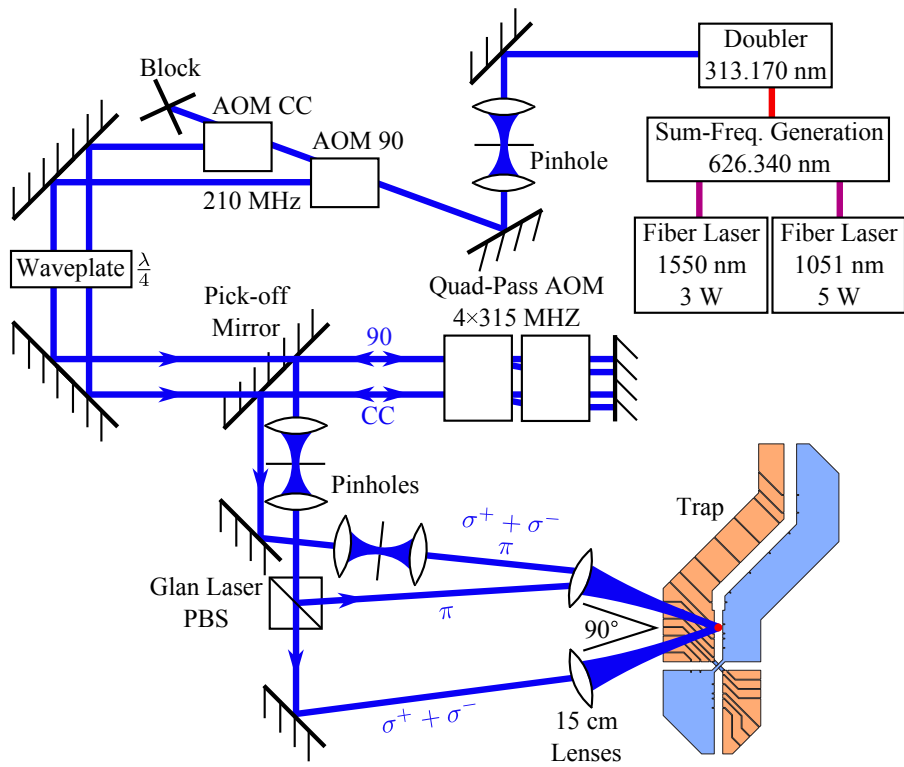


Figure 6.19: Setup to create the CC and 90° Raman beams for creating carrier and sideband transitions. Both AOM CC and AOM 90 shift the light by 210 MHz. See text for a detailed description of the setup.

twice through each AOM to produce four total passes. Ideally, each AOM will only deflect light that is vertically polarized and will not affect horizontally polarized light. In the quad-pass setup, the vertically polarized light is deflected in one direction by the first AOM and then again by the second AOM, but this time in the opposite direction so that the beam is parallel to the incident beam but shifted by ~ 0.5 mm and 630 MHz. It is then retroreflected by the mirror and passes again through the two AOMs where it picks up an additional 630 MHz frequency shift but the lateral position shift is removed. Thus the vertical polarization receives a total shift of 1260 MHz. The horizontal polarization, on the other hand, should not be deflected by any of the AOMs, so it passes straight through, reflects off of the mirror, and returns, picking up no shift. The vertically and horizontally polarized light paths will be colinear after the quad-pass setup, with the two polarizations having a relative frequency difference of 1260 MHz. This return light beam path is tilted vertically at a small angle so that it can be separated from the incident beam using a pick-off mirror.

For the CC beam, this combined beam is sent into the trap perpendicular to the B -field. This way the horizontal polarization is π light, while the vertical polarization is $\sigma^+ + \sigma^-$. The waveplate is used to adjust the relative intensities of the two frequencies at the ion.

Thus far, it has been assumed that an AOM only deflects vertical polarization. In reality, about 20% of the horizontal polarization is also deflected in each pass and about 20% of the vertical polarization is not deflected at each pass. This results in a small fraction of light that has the incorrect polarization for its frequency. This light will be present in the final beam but does not pose a significant problem for the experiments in this dissertation.

The 90° beams also pass through the quad-pass setup. The small horizontal displacement between CC and 90° beams allows efficient AOM defraction for both sets of beams. In contrast to the CC beams, a Glan Laser polarizing beam splitter is used to separate the two frequency components of the 90° beam into separate beams. The horizontal polarization is sent through the same vacuum window as the CC beams and produces π light on the ion, while the vertical polarization is sent through a different window so that it is perpendicular to the first beam and aligned with the B -field to produce $\sigma^+ + \sigma^-$.

By adjusting the rf frequency on the quad-pass AOMs, the splitting between the frequency components in the light can be made resonant either with the carrier transition, the red or blue sideband transitions, or the shelving transitions. This adjusts both the CC and 90° beams.

The typical power in the Raman beams is 10 mW (the sum of the two polarization components) and these are focused to $\sim 30 \mu\text{m}$ at the ion.

6.5.3 Light detection

The state of the ion is measured using state-dependent resonance fluorescence, where BD light is scattered from the ion. This scattered light is collected using a lens system that focuses onto a camera or photo-multiplier tube (PMT). The lens is an F-1.1 optic with a working distance of 4 cm. This lens is on a translation stage which allows the focus to be adjusted, as well as the lens to panned over the trap array to image different trap zones. The lens magnifies the image by a factor of 640.

The image of the ion is focused onto a Hamamatsu H6240 PMT, which has 25% quantum efficiency at 313 nm. An electronic TTL pulse is produced for each detected photon and counted by the experiment computer. During a typical detection for a bright ion, 1 photon is counted every $20 \mu\text{s}$.

Alternately, a ‘flipper’ mirror can be inserted to divert the focused light from the PMT to an Andor iXion DU885 EMCCD camera. This allows the ion to be imaged on a computer screen and is used primarily to see ions when they are first loaded into the trap. Two ions in the same trapping well, which will be separated by a few micrometers, are easily distinguishable by eye on the camera. The camera (and PMT) allows viewing of a $100 \mu\text{m}$ square region of the trap array, so it is not possible to simultaneously view zones on different sides of the X-junction (though the imaging system can pan from one side to the other).

Apertures are inserted in multiple locations along the imaging system to reduce laser light scattered from the electrodes.

6.6 Controlling the experiment

The experiment is controlled by a field programmable gate array (FPGA) similar to the setup described in [Langer 06]. A GUI computer program interfaces with the FPGA and allows the user to give instructions to the FPGA. The FPGA then sends out a series of pulses to different output devices to control the various aspects of the experiment and the FPGA collects and counts the TTL pulses coming from the PMT. The two types of output devices controlled by the FPGA are a series of direct digital synthesizers (DDS) and a series of TTL pulse generators. The DDSs are used as adjustable rf frequency sources for such applications as driving the AOMs. This allows the FPGA to control the rf frequency applied to, for example, the Raman quad-pass AOMs to set the transition the Raman beams will couple to.

The TTL outputs are used to switch different components of the experiment on and off. Several of the TTL channels are used to switch the rf power applied to the AOMs to turn the corresponding laser beams on and off. Another TTL pulse is used to signal the DACs to advance the waveform voltages to move the ion. The waveforms will have been pre-loaded into the PXI-6733 DAC cards by the computer GUI and the DACs wait for the TTL pulse from the FPGA to advance through the waveform.

During each detection period, the FPGA records photon counts from the PMT and sends that information to the computer GUI that interfaces with the FPGA, where the data is analyzed and displayed.

Chapter 7

Demonstration of transport through the X-junction

${}^9\text{Be}^+$ ions were transported through the X-junction array, to all three active legs of the X. The transports proved highly reliable with ion loss rates that appear to be limited by background gas collisions. The first transports through the junction resulted in the ion being excited by about 2000 quanta. After tracking down the main sources of motional excitation, it was possible to limit this motional excitation to ~ 1 quanta. The effect of junction transport on the internal atomic state coherence (the qubit coherence) was also tested and seen to be negligible. These results, in conjunction, demonstrate that junctions can be used in a multiplexed architecture to lead to scalable trap arrays.

In this chapter, these experimental demonstrations will be described, including the transport success rates, the motional excitation rates, and the test of qubit coherence. The chapter will then discuss the sources of motional excitation and how these were reduced. Finally, the chapter will end by considering what future improvements could be made.

7.1 Experiment procedure

The transport experiments all followed the same general procedure. An ion, or pair of ions, was loaded in zone \mathcal{L} by photo-ionizing neutral Be using the PI beam and simultaneously applying Doppler cooling light. The ion(s) were then transported to zone \mathcal{E} , which was the point of origin for all of the transport experiments. Once an ion was loaded, the same ion was used in all subsequent experiments until the ion was lost. Therefore, the load procedure was performed only when needed.




Figure 7.1: The ion begins and ends each transport experiment in \mathcal{E} . Three different waveforms, $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$, $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$, and $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$, were designed to move ions from one location to another. At each location \mathcal{C} , \mathcal{F} , or \mathcal{V} , the ion would pause (usually $20\ \mu\text{s}$) before continuing with the waveform.

The ion lifetime was typically hours, so it was common for the load procedure to be performed only once or twice in a day.

Once an ion was at location \mathcal{E} , the experiments began by cooling the ion to the Doppler limit using the BDD beam, followed by the BD beam. The ion was then optically pumped to the $|\downarrow\rangle$ state using RD1 and RD2. The ion was then further cooled to the ground state using 20 iterations of the ground state cooling pulses (90° Raman π -pulses on the red sideband followed by a pulse of RD1 and RD2). This completed the state initialization into $|\downarrow\rangle |n=0\rangle$ and the transport could begin. (In some experiments, ground state cooling was not needed and not performed.)

The ion was then transported into or through the junction and returned to \mathcal{E} using a waveform. A series of measurement pulses were performed to extract the desired data, typically including a $200\ \mu\text{s}$ detection pulse using BD. Then the experimental procedure was repeated, starting with the Doppler cooling and initialization. The experiment was repeated many (hundreds to thousands) times to build up sufficient statistics.

Three main transport waveforms were demonstrated and each transport will be named based on the destinations reached by that particular waveform. Figure 7.1 shows the four possible transports destinations. One transport starts at location \mathcal{E} , moves to the center of the junction \mathcal{C} , and

then returns to \mathcal{E} . This will be referred to as $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$. The other main transports are $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$ and $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$, which both move entirely through the junction (to different legs) and then return to \mathcal{E} . Due to the symmetry of the junction, $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ is analogous to moving entirely through the junction one time, while the other two transports are actually traversing the junction twice (forward and back). The $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transport moves the ion 1.76 mm, while $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$ and $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$ move the ion 3.52 mm and 2.84 mm, respectively.

7.2 Transport success rates

There are two means by which a transport can fail. The ion can either be lost entirely from the trap¹ or the ion can fail to reach the intended location. To check for both failure mechanisms, several different experiments were performed.

Since the imaging optics were only capable of viewing a $(100\ \mu\text{m})^2$ region of the trap, it was not possible to simultaneously view the all locations along the waveform trajectory to ensure the ion arrives at the intended locations. To test $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transport, two Doppler beams were passed through the trap. They both began centered on \mathcal{E} . Then the $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ waveform was run, but only partially, so that the ion moved some fraction of path toward \mathcal{C} before stopping and returning to \mathcal{E} . This fraction was gradually increased from 0 to 1, with one of the beams and the imaging optics translated to track this stopping position and ensure the ion was arriving at this location, while the other beam remained in \mathcal{E} . This proved that the waveform was moving the ion as intended.

In this experiment, it was not possible to track the ion with the laser beam and camera all the way from \mathcal{E} to \mathcal{C} . The ion was only visible from $z = -880\ \mu\text{m}$ to about $z = -300\ \mu\text{m}$. Beyond this point, the micromotion due to the pseudopotential barrier became too large and the scattering rate of the ion dropped to an undetectable level. Once the ion reached \mathcal{C} , the micromotion should have returned to zero and the ion should scatter once again, but the rf bridges obscure the camera's

¹ The term 'lost' refers to either having ions escape from the trapping potential or having ions undergo chemical reaction with a background gas molecule to form a molecular ion that does not scatter from the BD beam. Though a chemical interaction does not technically result in loss of the ion, it is difficult to distinguish this from actual ion loss, since we cannot detect or interact with the ion in either case.

view of the ion there.

To ensure the ion had indeed reached \mathcal{C} , a secondary waveform was included in the middle of the \mathcal{E} - \mathcal{C} - \mathcal{E} waveform, prior to returning the ion to \mathcal{E} . This inserted waveform is designed to eject any ion not at \mathcal{C} from the trap. The waveform starts by dropping all the control electrodes to 0 V. If the ion is at \mathcal{C} , the pseudopotential trapping at this location should keep it confined. Then each control voltages is raised, sequentially, starting with the electrodes nearest \mathcal{C} and continuing outward to more distant electrodes. This will sweep any ion not at \mathcal{C} toward the edges of the trap, where they will be ejected and lost. Any ion at \mathcal{C} will remain there due to the pseudopotential confinement. The control voltages are then dropped back to 0 V and the \mathcal{E} - \mathcal{C} - \mathcal{E} waveform is continued. This procedure did not result in loss of the ion and shows that the \mathcal{E} - \mathcal{C} - \mathcal{E} transport was transferring the ions to \mathcal{C} as intended.

To verify successful single-ion transport using the \mathcal{E} - \mathcal{C} - \mathcal{F} - \mathcal{C} - \mathcal{E} waveform, two experiments were performed. Doppler/detection beams were aligned both at both \mathcal{E} and \mathcal{F} , with the imaging optics and PMT focused at \mathcal{F} . The experiment began with the ion being prepared in location \mathcal{E} , after which the \mathcal{E} - \mathcal{C} - \mathcal{F} - \mathcal{C} - \mathcal{E} waveform was run halfway (\mathcal{E} - \mathcal{C} - \mathcal{F}). The detection beams were switched on in both zones with the expectation of observing the ion with the PMT at \mathcal{F} . The \mathcal{E} - \mathcal{C} - \mathcal{F} - \mathcal{C} - \mathcal{E} waveform was completed and a second detection was performed, now with the expectation of not observing an ion, as it should have been moved back to \mathcal{E} (and out of the PMT's view). Histogram pairs for 10,000 experiments (Fig. 7.2(b)) were analyzed and the lack of overlap implied the transport to \mathcal{F} had succeeded in every attempt.

The experiment was repeated but with the imaging optics now observing at \mathcal{E} . In this experiment, the ion should be detected during the second measurement and not during the first, as seen in Fig. 7.2(a). Again, 10,000 experiments showed the ion always returned to \mathcal{E} . Taken together, these experiments can be used to place an upper limit on the \mathcal{E} - \mathcal{C} - \mathcal{F} - \mathcal{C} - \mathcal{E} transport failure rate of 10^{-4} and show that the ion was moved to the correct location in each experiment. The histogram discrimination is slightly worse in Fig. 7.2(b) than in (a). This is because the B -field orientation and beam polarization was adjusted for maximal discrimination at zone \mathcal{E} , but the

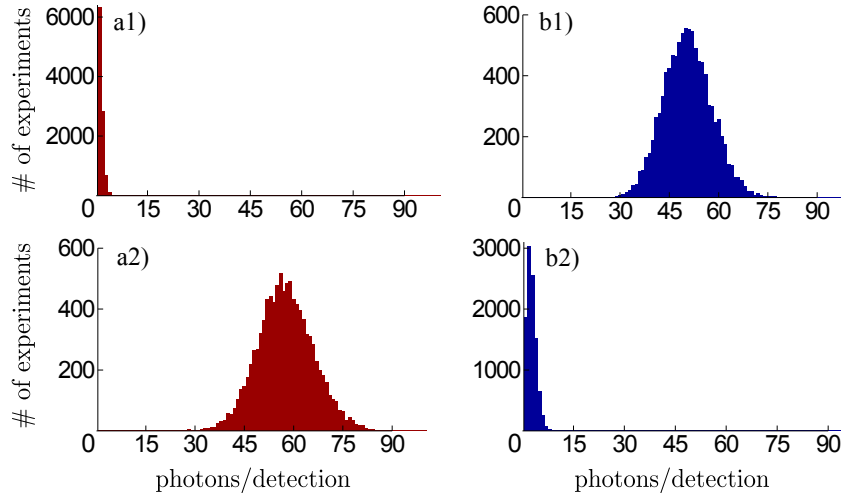


Figure 7.2: Fluorescence histogram pairs for experiments testing the success of $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transport. a) The histogram pairs when observing the ion at \mathcal{E} . In the first detection (a1), the ion should be in \mathcal{F} and the histogram should be dark. In the second histogram (a2) the ion should be in \mathcal{E} and the histogram should be bright. b) The same experiment but observing at \mathcal{F} . Here the first histogram should be bright and the second should be dark. The detection duration was $1200 \mu\text{s}$, which is longer than usual for state discrimination measurements, to reduce overlap of the bright and dark histograms.

same beam was also used in zone \mathcal{F} , by retro-reflecting it into this second zone after it had passed through \mathcal{E} . Thus the parameters were not nearly as optimal in \mathcal{F} , but they were still sufficient to determine the transport did not fail during these experiments.

The process was repeated for the $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$ waveform, with the same conclusion that the failure rate was below 10^{-4} . Here the discrimination was even worse at \mathcal{V} because the beam had to be sent through the trap at an angle orthogonal to the B -field, ruining the cycling transition. While using this non-ideal beam orientation, the BD was used ~ 50 MHz detuned to the red with the RD1/2 beams on simultaneously during the detection to achieve the maximum discrimination between an ion being present and not present.

The above experiments demonstrate that the ion is unlikely to be moved to an incorrect location during the transport. It is reasonable to assume that the waveform never fails in this manner and that, instead, transport failure would only result from ion loss. If the transports were to induce no loss at all, the ion lifetime would still be limited by background-gas collisions [Wineland 98],

which affect both moving and stationary ions. Though collisional losses occur too infrequently (~ 0.5 per hour) to build meaningful statistics, no increase in loss rates was observed for ions during transport over that for a stationary ion. This was true for all three transport waveforms. No experiments were performed specifically to test for ion loss. However, during the course of experiments measuring motional excitation, over 1 million $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$ and $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transports each were performed with less than 10 ions lost (consistent with collisional loss). Since transport comprised a small fraction (0.5 to 0.2) of the total experiment duration, many of these losses likely occurred when the ion was not transporting. Such collisional losses, though not related specifically to transport, limit the observed transport success probability for moving a single ion to 0.999995. In addition, over 10 million successive $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transports were recorded without any loss. For all three types of transport, given the low motional energy gain, we did not expect any failures due specifically to the transport procedure, which is compatible with the observations.

Thus, the single ion transport success probability was strictly demonstrated to exceed 0.9999 based on observations at the transport destinations, while it is likely to be ~ 0.999995 based on observed ion losses.

The loss rates for experiments involving transport of paired ions were higher than those for single ions but were still comparable to loss rates for two stationary ions (~ 2 per hour). This loss rate is higher than for a single stationary ion presumably due to multi-ion effects [Walther 93].

These experiments suggest that junction-transport can be highly reliable and limited by the overall lifetime of an ion in the trap.

7.3 Motional excitation rates

Motional excitation due to transport was determined using three different methods: Doppler recoiling [Epstein 07, Wesenberg 07], Rabi flopping [Meekhof 96], and sideband asymmetry [Turchette 00]. These three techniques are discussed in Appendix E and are each useful for different ranges of the average motional occupation \bar{n} . Which method was most applicable changed as the understanding of the sources of excitation increased and measures were incorporated into the

Transport	Energy Gain (quanta/trip)
$\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$	0.43 ± 0.1
$\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$	1.05 ± 0.2
$\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$	0.95 ± 0.2

Table 7.1: The axial-motion excitation, $\Delta\bar{n}$ for a single ${}^9\text{Be}^+$ ion for three different transports through the X-junction. This is measured using the method of sideband asymmetry (see Appendix E). Other measurement methods roughly agreed. The energy gain was measured after a single transit. The energy gain per trip is stated in units of quanta in a 3.6 MHz trapping well where $\Delta\bar{n} = 1$ quantum corresponds to 16 neV.

experiment to reduce \bar{n} . Whenever possible, multiple methods were used and compared to verify that the results roughly agreed, which they did in every case. Eventually, the single-ion excitation rates became low enough that sideband asymmetry method was the most practical to use. The two-ion excitation rates used the Rabi flopping method.

The procedure for measuring the motional excitation is as described in Sec. 7.1. The ion was initialized in the motional state $|0\rangle$. After performing either $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$, $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$, or $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transport, the ion was returned to \mathcal{E} where one of the three methods above was used to measure the resulting \bar{n} for the axial mode. This yields the average motional energy gain $\Delta\bar{n}$ for the transport protocol. The lowest recorded excitation rates are reported in Table 7.1. The main source of excitation is suspected to be noise on the rf field. These results are approximately an order of magnitude less than what was initially reported in [Blakestad 09], where the improvement was due to the addition of stronger filters to reduce noise coming from the DACs. Both sources of noise will be further discussed later in the chapter.

The transport between \mathcal{C} and \mathcal{E} or \mathcal{F} was performed in 140 μs and the distance moved was 880 μm , with 40 μs spent in transit through the pseudopotential barrier. The transport from \mathcal{C} to \mathcal{V} was performed in the same amount of time, 140 μs , despite the distance only being 540 μm .

Moving pairs of ions in the same trapping well would be useful for both sympathetic cooling and efficient ion manipulation [Kielinski 02]. This type of transport was demonstrated using pairs of ${}^9\text{Be}^+$ ions and the measured motional excitation is reported in Table 7.2. Both excitation of the

Transport	Energy Gain (quanta/trip)	
	COM	Stretch
$\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$	1.75 ± 0.25	0.6 ± 0.2
$\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$	2.0 ± 0.5	1.8 ± 0.8
$\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$	2.4 ± 0.6	2.0 ± 0.8

Table 7.2: The axial-motion excitation, $\Delta\bar{n}$ for a pair of ${}^9\text{Be}^+$ ions transported in the same trapping well. This is measured by fitting sideband Rabi flopping (see Appendix E). Values for both axial modes of motion (COM and stretch) are reported. The energy gain per trip is stated in units of quanta where the COM frequency is 3.6 MHz and the stretch frequency is 6.2 MHz.

COM and stretch modes were ground-state cooled and then measured after transport by fitting Rabi flopping on the sidebands. Additional heating mechanisms for multiple ions [Wineland 98, Walther 93] may explain the higher transport energy gain observed for the pair (Table 7.1). For $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{V}\text{-}\mathcal{C}\text{-}\mathcal{E}$ transport, the two-ion crystal must rotate from the \hat{z} axis to the \hat{x} axis before the ion begins to move from \mathcal{C} to \mathcal{V} . Since the waveform simply allowed these two directions to be degenerate at \mathcal{C} , rather than rotating in a controlled manner, it is possible that additional energy was gained during this rotation of the ion pair at \mathcal{C} , resulting in the larger energy gain seen versus $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{F}\text{-}\mathcal{C}\text{-}\mathcal{E}$.

Since the stretch mode frequency is higher than that of the COM mode ($\omega_{\text{STR}} = \sqrt{3}\omega_{\text{COM}}$), there will be less noise to excite this motion because the filters (both rf and control) suppress the higher frequency electronic noise more strongly. Furthermore, the stretch mode can only be excited by a differential force on the two ions. Given the proximity of the ions (a few μm), the size of any differential force that acts on the ion pair during transport is likely to be smaller than the symmetric force that will excite the COM mode. Therefore, it is expected that the stretch mode is less excited after transport, though the observed difference was not large.

7.4 Qubit coherence

For use in quantum computing, transport through a junction should preserve qubit coherence. To check this, a Ramsey experiment was performed and interleaved with the $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ junction

transport. The ion was first prepared at \mathcal{E} in the ground hyperfine state $|F = 2, m_F = -2\rangle$, which can be detected by fluorescence on a cycling transition. A pulse sequence (Fig. 7.3(a)) using coherent resonant Raman $\frac{\pi}{2}$ -pulses from the CC beam transferred the ion into an equal superposition of the $|2, -2\rangle$ (fluorescing) and $|1, -1\rangle$ (nonfluorescing) state. This was followed by a free-precession period $t_1 = 280 \mu\text{s}$, a spin-echo π -pulse (duration $1 \mu\text{s}$), and a second free-precession period $t_2 = t_1$. Finally, the phase ϕ of a second $\frac{\pi}{2}$ -pulse was varied and fluorescence was collected to generate a curve whose contrast is proportional to the ion's final coherence. The effect of transport on coherence was tested by inserting a single full \mathcal{E} - \mathcal{C} - \mathcal{E} transport during both t_1 and t_2 , during t_2 only, or not at all. The results are seen in Fig. 7.3(b). The contrast when not transporting was only $85.0 \pm 0.5 \%$ due to the presence of fluctuating magnetic fields and the length of the experiment. However, the contrast for one and both transports was $86.0 \pm 0.5 \%$ and $85.6 \pm 0.5 \%$, respectively, implying that transport does not affect coherence within the measurement errors. A difference in the B -field magnitude ($\frac{\Delta B}{B} \approx 0.4\%$) over the 880 nm path from \mathcal{E} to \mathcal{C} , which results in a variable transition frequency, induced a phase shift only when one transport was inserted, but the spin-echo effectively canceled this phase when the transport was performed twice.

B -field gradients will likely be present in any large-scale ion-trap array and will lead to this type of qubit phase shift. However, if the gradient is static (as is the case here), the induced phase shifts can be calibrated and included in the book-keeping for the algorithm being performed in the array, effectively eliminating the shift. Such shifts are not a fundamental problem for trapped-ion quantum computing but may pose a technical challenge when dealing with large-scale arrays.

7.5 Transporting $^{24}\text{Mg}^+$

In addition to moving both single and pairs of $^9\text{Be}^+$ ions through the junction, single $^{24}\text{Mg}^+$ ions were also successfully transferred along the same three transport paths. In fact, $^{24}\text{Mg}^+$ was the ion originally used for tests of transport in the X-junction array. However, shortly after the initial transports were demonstrated using $^{24}\text{Mg}^+$, the trap array was moved to a new laboratory that included the lasers needed for $^9\text{Be}^+$ instead of $^{24}\text{Mg}^+$, since $^9\text{Be}^+$ can be used as a qubit. Thus

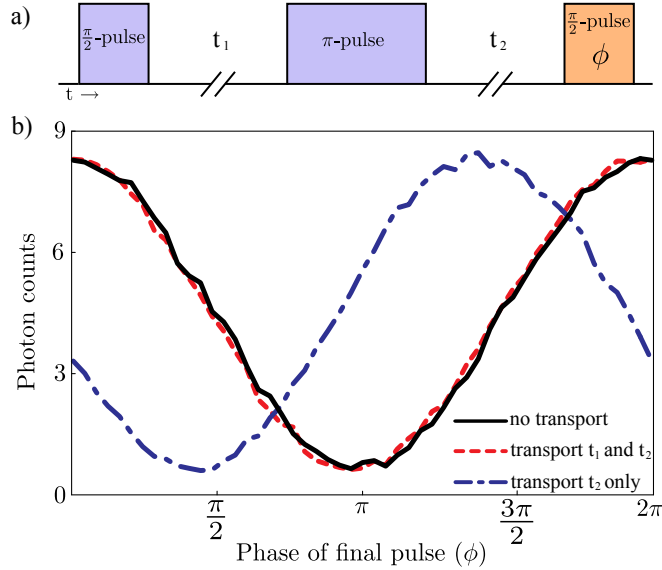


Figure 7.3: a) Pulse sequence for the Ramsey experiment showing the two $\frac{\pi}{2}$ -pulses with free-precession periods $t_1 = t_2$ and a spin-echo π -pulse. b) Fluorescence versus the final $\frac{\pi}{2}$ -pulse phase ϕ for three cases where a single full \mathcal{E} - \mathcal{C} - \mathcal{E} transport was inserted during both t_1 and t_2 , during t_2 only, or not at all.

significantly less time was spent perfecting $^{24}\text{Mg}^+$ transport.

The best results observed for doing the \mathcal{E} - \mathcal{C} - \mathcal{E} transport was 2.3×10^6 quanta. The failure rate, as measured by ion loss, was $\sim 1/1000$. The secular frequencies at \mathcal{E} were $\omega_z = 2\pi \times 1.8$ MHz, $\omega_x = 2\pi \times 4.1$ MHz, and $\omega_y = 2\pi \times 5.0$ MHz. The \mathcal{E} - \mathcal{C} - \mathcal{F} - \mathcal{C} - \mathcal{E} and \mathcal{E} - \mathcal{C} - \mathcal{V} - \mathcal{C} - \mathcal{E} transports were also successfully performed, but without quantitative results.

These results, though poor compared to those in $^9\text{Be}^+$, do demonstrate that $^{24}\text{Mg}^+$ can be moved through the junction. It is expected that the lessons learned while transporting $^9\text{Be}^+$, if applied to $^{24}\text{Mg}^+$, would produce similar results. One key improvement after the $^{24}\text{Mg}^+$ experiments was the computer model used for finding the potential and waveforms. When working with $^{24}\text{Mg}^+$, this model did not account for the alignment errors of the assembled trap. This was also true of the first $^9\text{Be}^+$ waveforms. For $^{24}\text{Mg}^+$, these waveforms worked, but poorly, while in $^9\text{Be}^+$ they did not work at all and the ion was lost on every transport attempt. Updating the model to account for these errors led to $^9\text{Be}^+$ waveforms that worked immediately with high reliability and excitation rates of ~ 2000 quanta/trip. Creating $^{24}\text{Mg}^+$ waveforms using the updated model should

result in several orders-of-magnitude reduction in the motional excitation. This, along with the noise-reduction techniques that will be discussed in the next section will likely reduce the $^{24}\text{Mg}^+$ excitation rates and failure rates to something comparable to the $^9\text{Be}^+$ results, if the array is returned to laboratory with $^{24}\text{Mg}^+$ lasers.

7.6 Mechanism that excite the secular motion

The axial mode of motion has the lowest secular frequency and is parallel to the direction of acceleration during the transport. Therefore, the majority of the motional excitation should occur along this axis, which is also the direction employed in demonstrated qubit gates. This is supported by the approximate agreement of energy gains estimated using Rabi flopping and sideband asymmetry (sensitive primarily to axial energy) versus Doppler recooling (sensitive to both axial and transverse energy; but in the analysis, all the excitation is attributed to the axial mode).

To obtain the low motional excitation rates observed during transport, careful attention was paid to sources of excitation. Two particular noise sources were identified that seem to cause the majority of this excitation: a particular rf-noise heating mechanism and DAC update noise.

7.6.1 rf-noise heating mechanism

Noise on the rf pseudopotential can cause ion heating [Wineland 98, Savard 97, Turchette 00], but this is negligible for an ion at the rf null on the axis of a standard linear Paul trap. However, similar heating occurs when the ion is placed in a pseudopotential gradient [Wineland 98]. Junction transports, which move over the pseudopotential barrier, will have such gradients, as can surface-electrode traps. The case of an axial pseudopotential gradient will be examined here, as this caused heating along the pseudopotential barriers near \mathcal{C} .

Consider an rf trap drive with a voltage noise component at $\Omega_{\text{rf}} + \omega$, where Ω_{rf} is the drive frequency and $|\omega| \ll \Omega_{\text{rf}}$. The trapping electric field will have the form

$$\mathbf{E}_{\text{rf}}(\mathbf{r}, t) = \mathbf{E}_0(\mathbf{r}) [\cos \Omega_{\text{rf}} t + \xi_N \cos (\Omega_{\text{rf}} + \omega)t], \quad (7.1)$$

where $\xi_N \ll 1$ is the noise amplitude relative to \mathbf{E}_0 , the ideal field. For $\xi_N = 0$, the above equation leads to the usual pseudopotential energy

$$\Phi_{\text{ps}} = \frac{q^2}{2m\Omega_{\text{rf}}^2} \langle \mathbf{E}_{\text{rf}}^2(\mathbf{r}, t) \rangle, \quad (7.2)$$

where the time average is performed over the period $2\pi/\Omega_{\text{rf}}$. The addition of the noise term, which beats with the carrier, introduces a noise force on the ion at frequency ω . For the transport between \mathcal{E} and \mathcal{C} , $\mathbf{E}_0(\mathbf{r})$ simplifies to $E_0(z)\hat{z}$, because the waveform restricts the ion to be on the trap axis ($x = 0, y = 0$), where the field is aligned with \hat{z} . This yields a noise force ($-\nabla\Phi_{\text{ps}}$) on the ion given by

$$\mathbf{F}_N \approx -\frac{q^2}{2m\Omega_{\text{rf}}^2} \left(\frac{\partial}{\partial z} E_0^2(z) \right) \xi_N \cos \omega t \hat{z}. \quad (7.3)$$

Equation (7.3) can be used to relate broad-spectrum voltage noise, given by a voltage spectral density S_{V_N} , to force noise spectral density S_{F_N} , by use of $S_{V_N}/V_0^2 = S_{E_N}/E_0^2$, where V_0 is the peak rf potential applied to the trap and S_{E_N} is the noise spectral density of electric field fluctuations $E_N = \xi_N E_0 \cos(\Omega_{\text{rf}} + \omega)t$. For ω near the axial-mode frequency ω_z , S_{F_N} heats the mode according to the rate $\dot{n} = S_{F_N}(\omega_z)/4m\hbar\omega_z$ [Turchette 00]. Therefore, the heating rate for noise applied around $\Omega_{\text{rf}} + \omega_z$ is

$$\dot{n} = \frac{q^4}{16m^3\Omega_{\text{rf}}^4\hbar\omega_z} \left[\frac{\partial}{\partial z} E_0^2(z) \right]^2 \frac{S_{V_N}(\Omega_{\text{rf}} + \omega_z)}{V_0^2}. \quad (7.4)$$

To test this, the ion was moved to a particular location along the axis between \mathcal{E} and \mathcal{C} and held there for a variable amount of time (≥ 1 ms) while various amounts of noise centered around the lower sideband $\Omega_{\text{rf}} - |\omega_z|$ was injected onto the trap rf drive. This was band-limited white noise (flat to 1 dB over 150 kHz) that was introduced using the ‘term’ port on bi-directional coupler A in Fig. 6.6. The ion was returned to \mathcal{E} and the increase in its motional excitation was measured using the recooling method. For each location along the axis, three different amplitudes of noise were used (all at least 40 dB above the background noise).

Figure 7.4 plots the ratio of measured heating rate to estimated injected S_{V_N} and theoretical values of this ratio according to Eq. (7.4) based on simulations of trap potentials, for the ion held at several positions between \mathcal{E} and \mathcal{C} . A plot including a scaling factor ($= 1.4$) is also included.




Figure 7.4: The ratio of heating rate \dot{n} to voltage noise spectral density $S_{V_N}(\Omega_{\text{rf}} - |\omega_z|)$ for various locations along the trap axis (\mathcal{C} is located at $0 \mu\text{m}$). The theoretical prediction used a pseudopotential modeled from electrode geometry and is shown both with and without a scaling parameter ($= 1.4$). The simulated pseudopotential is overlaid in the background, in units of eV. Since heating is gradient dependent, we see very little heating at the peak of the pseudopotential barrier, even though this is the point of maximum (axial) rf electric field. Nearly identical pseudopotential barriers occur on the other three legs of the junction.

The deviation of this factor from 1 is not unreasonable given the large number of experimental variables involved that are difficult to measure directly.

Two means were used to suppress this noise. First, extensive efforts were taken to reduce the amount of noise. The rf resonator will filter noise not at the resonant frequency, such as the noise at $\Omega_{\text{rf}} \pm \omega_z$. The resonator used has a Q of 42, which is low compared to typical resonator Q 's in the group (100 to 300 is more typical). By introducing a second resonator (see Sec. 6.2.1) with $Q = 170$, to filter the rf applied to the trap resonator, the noise in the rf sidebands was significantly reduced. With both resonators in place, the ambient $S_{V_N}(\Omega_{\text{rf}} \pm \omega_z)$ was estimated to be -177 dBc by performing the same heating measurements above, but without injecting noise.

The second method of reducing this form of motional heating is to move the ion over the barrier as fast as possible. Minimizing the duration the ion spends on the barrier limits the exposure to this noise. The motional excitation was observed to decrease as the ion speed was increased, up to some limit. At that maximum speed, the ion spent only $\sim 40 \mu\text{s}$ on each barrier. Given

this speed, the ambient S_{V_N} is expected to impart ~ 0.1 to 0.5 quanta per pass over an rf barrier. This is consistent with the values reported in Table 7.1. Above this maximum speed, the motional excitation increased and this is most likely due to DAC update noise, described in the next section.

7.6.2 DAC update noise

The other main source of motional excitation was traced to the digital-to-analog converters (DACs) that supply the waveform voltages to the electrodes. The DACs update their output voltage at a constant frequency U_{DAC} (≤ 500 kHz). This leads to discrete voltage steps on each electrode at this update frequency. If $2\pi \times U_{\text{DAC}} = \omega_z$, one would expect the DAC could excite the secular motion of the ion. However, the motion is also excited when $2\pi \times U_{\text{DAC}} = \omega_z/N$ for any integer N .

A specific experiment was used to test for this effect. The ion was prepared in the motional ground state in \mathcal{E} and then transported toward \mathcal{C} using the \mathcal{E} - \mathcal{C} - \mathcal{E} waveform. However, instead of traveling all the way to \mathcal{C} , the ion was stopped at $z = -300 \mu\text{m}$ and then returned to \mathcal{E} . As can be seen in Fig. 5.4, the axial secular frequency is constant along this portion of the waveform ($\omega_z = 2\pi \times 3.62$ MHz). Subsequently, a red sideband Raman π -pulse was applied and the fluorescence was measured. If the ion remained in the motional ground state, the π -pulse could not remove a quantum of motional energy and, thus, did not change the qubit state. In this case, the fluorescence was at full intensity. If, on the other hand, the motion was partially excited, the π -pulse could affect the qubit state to the ‘dark’ state and the fluorescence would drop. This experiment was performed for various U_{DAC} and the results are seen in Fig. 7.5.

The energy gain exhibited a resonance at several values for $2\pi \times U_{\text{DAC}} = \omega_z/N$ with $N = 8$ to 14. The number of update steps could be increased, while holding the update rate constant (which results in an increased transport duration). When this was done, the bandwidth of these resonances decreased, as would be expected from a coherent excitation.

Use of an update rate that was incommensurate with the motional frequency ($2\pi \times U_{\text{DAC}} \neq \omega_z/N$) minimized this energy gain. However, to increase the transport speed using the same update rate, the number of waveform steps has to be decreased and the dips in Fig. 7.5 broaden. At the

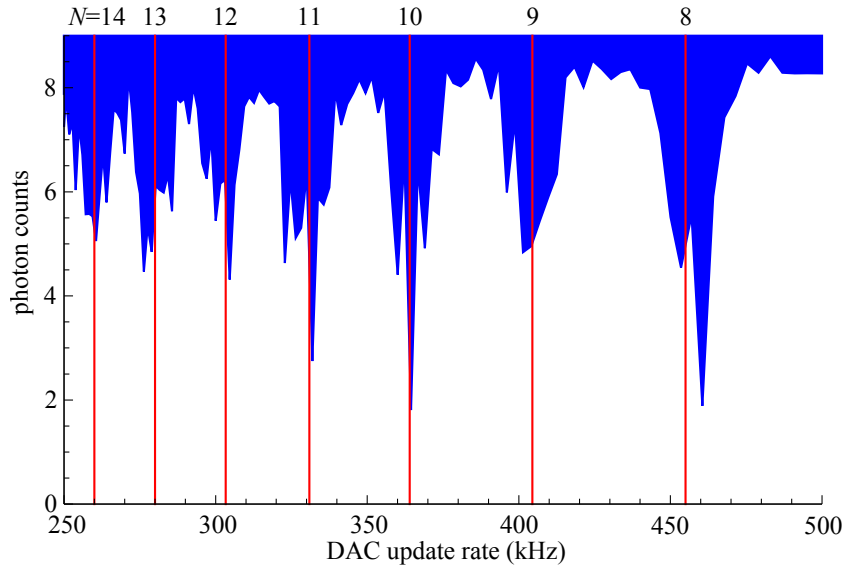


Figure 7.5: Plot showing the ion fluorescence following transport and a red sideband π -pulse, for various DAC update frequencies. The ion was prepared in the motional ground state and then transported through a specific waveform where $\omega_z = 2\pi \times 3.62$ MHz for the entire transport. If the ion remains in the ground state after transport, the ion should be at full fluorescence, but if the ion becomes motionally excited, the fluorescence will drop. As can be seen, the motion is excited at specific update frequencies that correspond to $\omega_z/2\pi N$ for $N = 8$ to 14 (marked by the vertical red lines).

same time, minimizing the rf-noise heating is required fast transport. At the speed that gave the lowest heating, the resonances were so broad that they overlap and there is no U_{DAC} that does not result in energy gain. This effect was further compounded by the fact that the axial frequency is not constant during a junction transport, making it impossible to achieve $2\pi \times U_{\text{DAC}} \neq \omega_z/N$ for any static U_{DAC} . Various U_{DAC} were tried during full \mathcal{E} - \mathcal{C} - \mathcal{E} transport. The update frequency $U_{\text{DAC}} = 479$ kHz appeared to minimize the total energy gain and was used for the results in Table 7.1 and in [Blakestad 09].

There are two possible routes to minimize the effects of DAC noise. A faster DAC that can achieve $2\pi \times U_{\text{DAC}} > \omega_z$ could be used, which would significantly suppress this source of motional excitation. Since the ion has several modes of motion and these secular frequencies can change during the course of a waveform, it may be necessary to use a U_{DAC} that is greater than all of the secular frequencies, including the radial modes. Furthermore, low-noise fast DACs are expensive,

especially if many channels are needed (the X-junction arrays has 40 independent electrodes and therefore requires up to 40 DACs). The second method is to more aggressively filter the DAC output. RC filters are generally used for this purpose. The RC constant must be large enough to substantially filter noise at ω_z but small enough so that the waveforms can be run quickly and the ion moved sufficiently fast. In [Blakestad 09], a RC corner frequency of 160 kHz was used, which still passed enough of the DAC-update noise at 3.6 MHz to observe ion excitation. For the improved motional excitation results given earlier in this chapter, the RC filter was replaced with an LCRC filter described in Sec. 6.2.2. This provided two orders-of-magnitude more filtering at ω_z while still allowing fast transport. The addition of these new filters has removed most of the DAC update noise. The remaining motional excitation seen in Tables 7.1 and 7.2 is likely limited by the rf-noise mechanism at this point.

This DAC noise acts like a coherent excitation while the rf-noise acts like a thermal excitation. When they are both simultaneously acting on the ion, the resulting state will be some complicated convolution of a thermal and coherent state. When using the techniques in Appendix E to measure \bar{n} , this complicated motional state will affect the results and should be considered. The exception to this is when $\bar{n} \lesssim 1$ quanta, in which case the thermal and coherent distributions are similar enough that assuming a thermal state is sufficient.

7.6.3 Other mechanisms

In addition to the two excitation mechanisms above, there are a number of other possible avenues for an ion to gain motional energy during transport. The issue of adiabaticity has already been discussed in Ch. 5. This is not likely to have resulted in excitation during the experiments described above, due to the relatively slow speed. The average speed, 6.3 m/s, was chosen as it seemed the best compromise between moving fast to minimize rf-noise excitation and moving slowly to minimize DAC-noise excitation. In addition to varying the average speed, different velocity profiles, as discussed in Sec. 5.4, were tested. Both a constant velocity and a sinusoidal velocity profile were used, but neither produced an appreciable difference in the excitation rate, suggesting

the transport was not near the limit of adiabaticity.

There existed an optimal rf power that produced the lowest motional excitation. This is not surprising, since the waveforms are calculated with an assumed rf power (leading to an assumed pseudopotential amplitude) and the actual trapping well produced by the rf and control electrodes will deviate from the intended well if the rf power changes. The trapping well will likely be shifted in one direction and the secular frequencies may not equal the intended values. On the barriers, the pseudopotential contributes significantly to the intended trapping well, making these locations the most sensitive to the rf power. Since the rf power is prone to slow drifts over many minutes (probably due to temperature effects on the resonator and coupling into the resonator), it was necessary to occasionally adjust this power and hold it to within $< 1\%$ to achieve the best transport excitation rates. Every five or ten minutes, the rf power was adjusted so that the radial secular frequencies at \mathcal{E} were within a desired range.

An additional form of motional heating that excites the motion of the ions is so called ‘anomalous heating’ [Wineland 98, Turchette 00]. This is thought to arise from noisy electric potentials on the surface of the electrodes. In this model, the heating scales as R^{-4} [Turchette 00, Deslauriers 06], where R is the distance from the ion to the nearest electrode. In the X-junction array, this heating rate was measured to be 40 quanta/s, which would only contribute 0.01 quanta for \mathcal{E} - \mathcal{C} - \mathcal{E} transport. Thus this form of heating does not contribute significantly to the results given earlier in this chapter.

One useful method for mitigating motional excitation is to increase the secular frequency of the ion. The noise spectral density is typically observed to be lower at higher frequencies, which will reduce both the rf-noise excitation mechanism and the anomalous heating. In addition, increasing the secular frequency helps with DAC heating by forcing the different resonance conditions for U_{DAC} to be farther spaced in frequency and increases the N needed to reach those conditions. When testing different versions of the \mathcal{E} - \mathcal{C} - \mathcal{E} waveform that had different constraints on the secular frequency, it was found that higher secular frequencies, especially in locations on the pseudopotential barriers, resulted in lower excitation rates, as expected. The waveform with the lowest excitation

rate used the axial secular frequency ramp shown in Fig. 5.4. One of the advantages of the ${}^9\text{Be}^+$ ion is its light mass, which allows for higher secular frequencies and less motional excitation.

7.7 Possible improvements

The experiments above have demonstrated reliable ion transport through a two-dimensional junction array with low motional-energy gain while maintaining internal-state qubit coherence. However, there is still some room left for improvement. Though the motional energy gain for a single pass through the junction is low, many passes would likely be needed in an actual algorithm, especially as the number of qubits increases. Motional excitation would accrue during each pass and would soon exceed the limits needed for good gate fidelity.

At these energies, sympathetic cooling could be employed for recooling an excited qubit ion in a reasonable period. For the parameters of current experiments, cooling from motional state $n = 1$ back to $n = 0$ would require approximately $20 \mu\text{s}$ and sympathetic cooling could be performed in a similar period.

Still, it would be helpful to further reduce the motional excitation during transport. Since rf noise is likely the current limit on this motional excitation, further reducing the noise at $\Omega_{\text{rf}} \pm \omega_z$ on the rf drive using a higher Q resonator or better electronics, could help. Future trap designs could attempt to change the shape of the pseudopotential barrier to reduce the slope and resulting noise.

In addition, faster DACs will help reduce the DAC update excitation. Such DACs would also allow the ion to be moved faster over the rf bump to decrease susceptibility to the rf noise while reducing the time needed for transport.

Chapter 8

Conclusion

Though an enormous amount of work remains before a large-scale quantum computer will be realized, the work described here should provide a path toward transfer of information in a larger-scale ion-based quantum processor and the potential to employ an increased number of qubits in quantum algorithm experiments.

In future work, junction architectures can be exploited to allow control over larger numbers of qubits than is currently possible. To do this, the waveforms must be expanded to incorporate multiple trapping wells to accommodate the different qubits, but this should be straight-forward. In addition, it is likely that a sympathetic cooling ion species will need to be co-trapped with the ${}^9\text{Be}^+$ so that the motional excitation from electronic noise or from multiple junction transports or separating and combining wells can be removed. If sympathetic-cooling ions are present, it may be advantageous to move trapping wells that contain ions of both species through a junction. Since the pseudopotential (and micromotion) are mass dependent, these different ions will see different potentials, which could lead to motional excitation of the ions. Given the large pseudopotential barriers near a junction, it is not clear that multiple species can be simultaneously transported through a junction without significant heating. However, it should still be possible to separate the ions by species and pass the different species through the junction separately. They could then be re-combined on the other side of the junction and sympathetically cooled.

The work here highlights a number of issues that future trap arrays must consider, and in some cases provides solutions to those issues. The DAC-update noise can be mitigated using

Butterworth filters and/or faster DACs. Noise on the trap rf at the motional sideband can result in motional excitation, especially on the pseudopotential barriers near a junction. This work shows that the slope of the pseudopotential barrier is more important than its amplitude for this particular issue. High- Q rf resonators can reduce such noise, but as trap arrays become more complex, the capacitance of the trap will likely increase and achieving large rf Q 's may become more difficult.

The success of the ion-transport protocols reported is an important step towards demonstrating scalability in the multiplexed architecture approach to trapped-ion based quantum computing described in [Wineland 98, Kielpinski 02]. We have seen that transport through a two-dimensional trap array incorporating a junction can be highly-reliable and excite the ion's harmonic motion by less than one quantum, marking a significant improvement over prior work with junction arrays [Hensinger 06]. At the current motional excitation, sympathetic cooling could be employed. Alternately, some quantum gates are robust against small energy gain [Mølmer 99, Milburn 00, Solano 99, Ospelkaus 08] and may require only minimal recooling. In addition, the internal state qubit coherence is not adversely affected. Thus, the viability of trap arrays incorporating junctions for use in large-scale ion-based quantum computing has been demonstrated.

Bibliography

- [Abramowitz 64] M. Abramowitz & I.A. Stegun. Handbook of mathematical functions with formulas, graphs, and mathematical tables. U.S. Govt. Print. Off., Washington, 1964.
- [Alheit 96] R. Alheit, S. Kleineidam, F. Vedel, M. Vedel & G. Werth. Higher order non-linear resonances in a Paul trap. International Journal of Mass Spectrometry and Ion Processes, vol. 154, pages 155–169, 1996.
- [Alheit 97] R. Alheit, X.Z. Chu, M. Hoefler, M. Holzki, G. Werth & R. Blümel. Nonlinear collective oscillations of an ion cloud in a Paul trap. Phys. Rev. A, vol. 56, pages 4023–4031, 1997.
- [Amini 09] J.M. Amini. private communication, 2009.
- [Amini 10] J.M. Amini, H. Uys, J.H. Wesenberg, S. Seidelin, J. Britton, J.J. Bollinger, D. Leibfried, C. Ospelkaus, A.P. VanDevender & D.J. Wineland. Toward scalable ion traps for quantum information processing. New J. Phys., vol. 12, page 033031, 2010.
- [ANSYS] ANSYS. ANSYS, Inc., www.ansys.com. Canonsburg, PA.
- [Bardroff 96] P.J. Bardroff, C. Leichtle, G. Schrade & W.P. Schleich. Endoscopy in the Paul trap: measurement of the vibratory quantum state of a single ion. Phys. Rev. Lett., vol. 77, pages 2198 – 2201, 1996.
- [Barrett 03] M.D. Barrett, B. DeMarco, T. Schaetz, V. Meyer, D. Leibfried, J. Britton, J. Chiaverini, W.M. Itano, B. Jelenkovic, J.D. Jost, C. Langer, T. Rosenband & D.J. Wineland. Sympathetic cooling of $^9\text{Be}^+$ and $^{24}\text{Mg}^+$ for quantum logic. Phys. Rev. A, vol. 68, page 042302, 2003.
- [Barrett 04] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri & D. J. Wineland. Deterministic quantum teleportation of atomic qubits. Nature, vol. 429, pages 737–739, 2004.
- [Baseia 93] B. Baseia, R. Vyas & V.S. Bagnato. Particle trapping by oscillating fields: squeezing effects. Quant. Opt., vol. 5, pages 155–159, 1993.

- [Bernardini 98] M. Bernardini, S. Braccini, R. De Salvo, A. Di Virgilio, A. Gaddi, A. Gennai, G. Genuini, A. Giazotto, G. Losurdo, H. B. Pan, A. Pasqualetti, D. Passuello, P. Popolizio, F. Raffaelli, G. Torelli, Z. Zhang, C. Bradaschia, R. Del Fabbro, I. Ferrante, F. Fidecaro, P. La Penna, S. Mancini, R. Poggiani, P. Narducci, A. Solina & R. Valentini. Air bake-out to reduce hydrogen outgassing from stainless steel. *J. Vac. Sci. Technol. A*, vol. 16, page 188, 1998.
- [Blain 08] M.G. Blain, A. Allerman, M.H. Crawford, W.W. Chow, M.P. Lilly, L.G. Pierson, B.R. Hamlet, J.R. Hamlet, O.B. Spahn, G.N. Nielson, P.J. Resnick, E.J. Heller, C.P. Tiggles, B. Jokiel Jr & B.E. Jakaboski. Developing Key Capabilities for Quantum Computing: Trapped Ion and GaAs Approaches. Rapport technique, SAND2008-6189, Sandia National Laboratories, Albuquerque, NM, 2008.
- [Blakestad 09] R.B. Blakestad, C. Ospelkaus, A.P. VanDevender, J.M. Amini, J. Britton, D. Leibfried & D.J. Wineland. High-Fidelity Transport of Trapped-Ion Qubits through an X-Junction Trap Array. *Phys. Rev. Lett.*, vol. 102, no. 15, page 153002, 2009.
- [Blatt 08] R. Blatt & D. J. Wineland. Entangled states of trapped atomic ions. *Nature*, vol. 453, no. 7198, pages 1008–1015, JUN 19 2008.
- [Bollinger 85] J.J. Bollinger, J.S. Wells, D.J. Wineland & W.M. Itano. Hyperfine structure of the $2p^2P_{1/2}$ state in $^9\text{Be}^+$. *Phys. Rev. A*, vol. 31, no. 4, pages 2711–2714, Apr 1985.
- [Bollinger 91] J.J. Bollinger, D.J. Heinzen, W.M. Itano, S.L. Gilbert & D.J. Wineland. A 303-MHz Frequency Standard Based on Trapped Be^+ Ions. *IEEE Trans. Instr. Meas.*, vol. 40, page 126, 1991.
- [Britton 08] Joe Britton. Microfabricated Ion Traps for Quantum Computing. PhD thesis, University of Colorado, Boulder, CO, USA, 2008.
- [Cetina 07] M. Cetina, A. Grier, J. Campbell, I. Chuang & V. Vuletic. Bright source of cold ions for surface-electrode traps. *Phys. Rev. A*, vol. 76, page 041401(R), 2007.
- [Chiaverini 05] J. Chiaverini, R.B. Blakestad, J. Britton, J.D. Jost, C. Langer, D. Leibfried, R. Ozeri & D.J. Wineland. Surface-Electrode Architecture for Ion-Trap Quantum Information Processing. *Quant. Inf. Comput.*, vol. 5, pages 419–439, 2005.
- [Chiaverini 08] J. Chiaverini & W.E. Lybarger Jr. Laserless trapped-ion quantum simulations without spontaneous scattering using microtrap arrays. *Phys. Rev. A*, vol. 77, page 022324, 2008.
- [Chou 10] C.W. Chou, D.B. Hume, J.C.J. Koelemeij, D.J. Wineland & T. Rosenband. Frequency Comparison of Two High-Accuracy Al^+ Optical Clocks. *Phys. Rev. Lett.*, vol. 104, page 070802, 2010.
- [Cirac 95] J. I. Cirac & P. Zoller. Quantum Computations with Cold Trapped Ions. *Phys. Rev. Lett.*, vol. 74, no. 20, pages 4091–4094, May 1995.

- [CoorsTek] CoorsTek. Inc., www.coorstek.com. Golden, CO, data sheet.
- [Couvert 08] A. Couvert, T. Kawalec, G. Reinaudi & D. Guery-Odelin. Optimal transport of ultracold atoms in the non-adiabatic regime. Europhysics Letters, vol. 83, no. 1, page 13001, 2008.
- [CPO] CPO. Scientific Instrument Service, Inc., www.simion.com/cpo. Ringoes, NJ.
- [Cubric 99] D. Cubric, B. Lencova, F.H. Read & J. Zlamal. Comparison of FDM, FEM and BEM for electrostatic charged particle optics. Nucl. Instr. Meth. A, vol. 427, pages 357–362, 1999.
- [Dawson 69] P.H. Dawson & N.R. Whetten. Non-Linear Resonances in Quadrupole Mass Spectrometers due to Imperfect Fields I. The Quadrupole Ion Trap. International Journal of Mass Spectrometry and Ion Physics, vol. 2, pages 45–59, 1969.
- [Dawson 95] P.H. Dawson. Quadrupole mass spectrometry and its applications. Springer, 1995.
- [Dehmelt 90] Hans Dehmelt. Experiments with an isolated subatomic particle at rest. Rev. Mod. Phys., vol. 62, no. 3, pages 525–530, Jul 1990.
- [Deslauriers 06] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk & C. Monroe. Scaling and suppression of anomalous heating in ion traps. Phys. Rev. Lett., vol. 97, page 103007, 2006.
- [DiVincenzo 98] D.P. DiVincenzo & D. Loss. Quantum information is physical. Superlattices Microstruct., vol. 23, pages 419–432, 1998.
- [Drakoudis 06] A. Drakoudis, M. Söllner & G. Werth. Instabilities of ion motion in a linear Paul trap. International Journal of Mass Spectrometry, vol. 252, pages 61–68, 2006.
- [Duan 10] L.-M. Duan & C. Monroe. Quantum networks with trapped ions. Rev. Mod. Phys., page accepted, 2010.
- [El-Dib 01] Y.O. El-Dib. Nonlinear Mathieu equation and coupled resonance mechanism. Chaos, Solitons & Fractals, vol. 12, pages 705–720, 2001.
- [Epstein 07] R. J. Epstein, S. Seidelin, D. Leibfried, J. H. Wesenberg, J. J. Bollinger, J. M. Amini, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, N. Shiga & D. J. Wineland. Simplified motional heating rate measurements of trapped ions. Phys. Rev. A, vol. 76, page 033411, 2007.
- [Feynman 82] R.P. Feynman. Simulating physics with computers. Int. J. Th. Phys., vol. 21, pages 467–488, 1982.
- [Folman 05] R. Folman. private communication, 2005.
- [Franzen 91] J. Franzen. Simulation study of an ion cage with superimposed multipole fields. International Journal of Mass Spectrometry and Ion Processes, vol. 106, pages 63–78, 1991.

- [Gabrielse 05] G. Gabrielse. Atoms Made Entirely of Antimatter: Two Methods Produce Slow Antihydrogen. *Adv. At. Mol. Opt. Phys.*, vol. 50, pages 155 – 217, 2005.
- [Gerlich 92] D. Gerlich. Inhomogenous RF Fields: A Versatile Tool for the Study of Processes with Slow Ions. *Adv. Chem. Phys.*, vol. 82, pages 1–176, 1992.
- [Ghosh 95] P.K. Ghosh. Ion traps, volume 90 of The International series of monographs on physics. Clarendon Press, Oxford, 1995.
- [Gill 81] P.E. Gill, W. Murray & M.H. Wright. Practical optimization. Academic Press, London, 1981.
- [Golub 65] G.H. Golub & W. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *J. Soc. Ind. Appl. Math. B*, vol. 2, pages 205–224, 1965.
- [Gottesman 99] D. Gottesman & I.L. Chuang. Demonstrating the Viability of Universal Quantum Computation Using Teleportation and Single-Qubit Operations. *Nature*, vol. 402, pages 390–393, 1999.
- [Grover 97] L.K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.*, vol. 79, page 325, 1997.
- [Gudjons 97] T. Gudjons, P. Seibert & G. Werth. Influence of anharmonicities of a Paul trap potential on the motion of stored ions. *Appl. Phys. B*, vol. 65, pages 57–62, 1997.
- [Häffner 08] H. Häffner, C.F. Roos & R. Blatt. Quantum Computing with Trapped Ions. *Phys. Rep.*, vol. 469, page 155, 2008.
- [Hanneke 08] D. Hanneke, S. Fogwell & G. Gabrielse. New Measurement of the Electron Magnetic Moment and the Fine Structure Constant. *Phys. Rev. Lett.*, vol. 100, page 120801, 2008.
- [Hansch 80] T W Hansch & B Couillaud. Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity. *Opt. Comm.*, vol. 35, pages 441–444, 1980.
- [Haroche 06] S. Haroche & J.-M. Raimond. Exploring the quantum : atoms, cavities and photons. Oxford Univ. Press, 2006.
- [Hasegawa 05] T. Hasegawa & J.J. Bollinger. Rotating-radio-frequency ion traps. *Phys. Rev. A*, vol. 72, page 043403, 2005.
- [Hensinger 06] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe & J. Rabchuk. T-junction ion trap array for two-dimensional ion shuttling, storage and manipulation. *Appl. Phys. Lett.*, vol. 88, no. 3, page 034101, 2006.
- [Herskind 09] P.F. Herskind, A. Dantan, M. Albert, Marler J.P & M. Drewsen. Positioning of the rf potential minimum line of a linear Paul trap with micrometer precision. *J. Phys. B: At. Mol. Opt. Phys.*, vol. 42, page 154008, 2009.

- [Hoffnagle 95] J.A. Hoffnagle & R.G. Brewer. Stability of two-ion crystals in the Paul trap: A comparison of exact and pseudopotential calculations. *Physica Scripta*, vol. 159, page 380, 1995.
- [Home 06] J.P. Home & A.M. Steane. Electrode Configurations for Fast Separation of Trapped Ions. *Quant. Inf. Comput.*, vol. 6, no. 4-5, pages 289–325, 2006.
- [Home 09] J.P. Home, D. Hanneke, J.D. Jost, J.M. Amini, D. Leibfried & D.J. Wineland. Complete methods set for scalable ion trap quantum information processing. *Science*, vol. 325, page 1227, 2009.
- [Home 10] J.P. Home. private communication, 2010.
- [Huber 08] G. Huber, T. Deuschle, W. Schnitzler, R. Reichle, K. Singer & F. Schmidt-Kaler. Transport of ions in a segmented linear Paul trap in printed-circuit-board technology. *New J. Phys.*, vol. 10, page 013004, JAN 14 2008.
- [Hucul 06] D. Hucul. Operation of a two-dimensional ion trap array for scalable quantum computing. Master's thesis, University of Michigan, 2006.
- [Hucul 08] D. Hucul, M. Yeo, S. Olmschenk, Monroe, W. K. C.Hensinger & J. Rabchuk. On the Transport of Atomic Ions in Linear and Multidimensional Ion Trap Arrays. *Quant. Inf. Comput.*, vol. 8, pages 501–578, 2008.
- [Hughes 96] R.J. Hughes, D.F.V. James, E.H. Knill, R. Laflamme & A.G. Petschek. Decoherence bounds on quantum computation with trapped ions. *Phys. Rev. Lett.*, vol. 77, pages 3240–3243, 1996.
- [Hume 07] D.B. Hume, T. Rosenband & D.J. Wineland. High-fidelity adaptive qubit detection through repetitive quantum nondemolition measurements. *Phys. Rev. Lett.*, vol. 99, page 120502, 2007.
- [Itano 82] W.M. Itano & D.J. Wineland. Laser cooling of ions stored in harmonic and penning traps. *Phys. Rev. A*, vol. 25, pages 35–54, 1982.
- [Jackson 99] J.D. Jackson. *Classical electrodynamics.* Wiley, New York, 3rd edition, 1999.
- [Jefferts 95] S.R. Jefferts, C. Monroe, E.W. Bell & D.J. Wineland. Coaxial-resonator-driven RF (Paul) trap for strong confinement. *Phys. Rev. A*, vol. 51, no. 4, pages 3112–3116, 1995.
- [Jost 09] J.D. Jost, J.P. Home, J.M. Amini, D. Hanneke, R. Ozeri, C. Langer, J.J. Bollinger, D. Leibfried & D.J. Wineland. Entangled Mechanical Oscillators. *Nature*, vol. 459, page 683, 2009.
- [Jost 10] J.D. Jost. Entangled Mechanical Oscillators. PhD thesis, University of Colorado, 2010.
- [Kielinski 02] D. Kielinski, C. Monroe & D.J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, vol. 417, pages 709–711, 2002.

- [King 98] B.E. King, C.S. Wood, C.J. Myatt, Q.A. Turchette, D. Leibfried, W.M. Itano, C. Monroe & D.J. Wineland. Cooling the collective motion of trapped ions to initialize a quantum register. *Phys. Rev. Lett.*, vol. 81, pages 1525–1528, 1998.
- [King 99] B.E. King. Quantum State Engineering and Information Processing with Trapped Ions. PhD thesis, University of Colorado, 1999.
- [Konenkov 02] N.V. Konenkov, M. Sudakov & D.J. Douglas. Matrix Methods for the Calculation of Stability Diagrams in Quadrupole Mass Spectroscopy. *J. Am. Soc. Mass. Spectrom.*, vol. 13, pages 597–613, 2002.
- [Labaziewicz 08] J. Labaziewicz, Y. Ge, D.R. Leibbrandt, S.X. Wang, R. Shewmon & I.L. Chuang. Temperature Dependence of Electric Field Noise above Gold Surfaces. *Phys. Rev. Lett.*, vol. 101, page 180602, 2008.
- [Langer 05] C. Langer, R. Ozeri, J.D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R.B. Blakestad, J. Britton, D.B. Hume, W.M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P.O. Schmidt & D.J. Wineland. Long-lived qubit memory using atomic ions. *Phys. Rev. Lett.*, vol. 95, 2005.
- [Langer 06] C. Langer. High fidelity quantum information processing with trapped ions. PhD thesis, University of Colorado, Boulder, 2006.
- [Leibfried 03] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband & D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, vol. 422, pages 412–415, 2003.
- [Leibfried 07] D. Leibfried, E. Knill, C. Ospelkaus & D. J. Wineland. Transport quantum logic gates for trapped ions. *Phys. Rev. A*, vol. 76, no. 3, page 0323241, 2007.
- [Lin 09] G.-D. Lin, S.-L. Zhu, R. Islam, K. Kim, M.-S. Chang, S. Korenblit, C. Monroe & L.-M. Duan. Large-scale quantum computation in an anharmonic linear ion trap. *Euro. Phys. Lett.*, vol. 86, page 60004, 2009.
- [Madsen 04] M.J. Madsen, W.K. Hensinger, D. Stick, J.A. Rabchuk & C. Monroe. Planar ion trap geometry for microfabrication. *Appl. Phys. B*, vol. 78, pages 639–651, 2004.
- [Marquet 03] C. Marquet, F. Schmidt-Kaler & D.F.V. James. Phonon-phonon interactions due to non-linear effects in a linear ion trap. *Appl. Phys. B*, vol. 76, pages 199–208, 2003.
- [Mathews 70] J. Mathews & R.L. Walker. *Mathematical method of physics*. Addison-Wesley Pub. Co, 1970.
- [MATLAB 07] MATLAB. MathWorks, 2007. Natick, MA.
- [Meekhof 96] D.M. Meekhof, C. Monroe, B.E. King, W.M. Itano & D.J. Wineland. Generation of Nonclassical Motional States of a Trapped Atom. *Phys. Rev. Lett.*, vol. 76, no. 11, pages 1796–1799, 1996. This has an erratum.

- [Metcalf 99] H.J. Metcalf & P. Van der Straten. Laser cooling and trapping. Springer-Verlag, New York, 1999.
- [Michaud 05] A.L. Michaud, A.J. Frank, C. Ding, XianZhen Zhao & D.J. Douglas. Ion Excitation in a Linear Quadrupole Ion Trap with an Added Octopole Field. Journal of the American Society for Mass Spectrometry, vol. 16, pages 835–849, 2005.
- [Milburn 00] G.J. Milburn, S. Schneider & D.F. James. Ion Trap Quantum Computing with Warm Ions. Fortschr. Physik, vol. 48, no. 9-11, pages 801–810, 2000.
- [Moerhing 07] D.L. Moerhing, M.J. Madsen, K.C. Younge, R.N. Kohn Jr, P. Maunz, L.-M. Duan, C. Monroe & B. Blinov. Quantum Networking with Photons and Trapped Atoms. J. Opt. Soc. Am. B, vol. B24, pages 300–315, 2007.
- [Mølmer 99] K. Mølmer & A. Sørensen. Multiparticle Entanglement of Hot Trapped Ions. Phys. Rev. Lett., vol. 82, no. 9, pages 1835–1838, 1999.
- [Monroe 95] C. Monroe, D.M. Meekhof, B.E. King, S.R. Jefferts, W.M. Itano, D.J. Wineland & P. Gould. Resolved-sideband raman cooling of a bound atom to the 3D zero-point energy. Phys. Rev. Lett., vol. 75, no. 22, pages 4011–4014, 1995.
- [Monroe 08] C. Monroe & M. D. Lukin. Remapping the quantum frontier. Physics World, vol. 21, no. 8, pages 32–39, 2008.
- [Moore 20] E.H. Moore. On the reciprocal of the general algebraic matrix. Bull. Am. Math. Soc., vol. 26, pages 394–395, 1920.
- [Nägerl 99] H.C. Nägerl, D. Leibfried, H. Rohde, G. Thalhammer, J. Eschner, F. Schmidt-Kaler & R. Blatt. Laser addressing of individual ions in a linear ion trap. Phys. Rev. A, vol. 60, page 145, 1999.
- [Neuhauser 78] W. Neuhauser, M. Hohenstatt, P. Toschek & H. Dehmelt. Optical-Sideband Cooling of Visible Atom Cloud Confined in Parabolic Well. Phys. Rev. Lett., vol. 41, no. 4, pages 233–236, Jul 1978.
- [Nie 09] X.R. Nie, C.F. Roos & D.F.V. James. Theory of cross phase modulation for the vibrational modes of trapped ions. Phys. Lett. A, vol. 373, pages 422–425, 2009.
- [Ospelkaus 08] C. Ospelkaus, C.E. Langer, J.M. Amini, K.R. Brown, D. Leibfried & D.J. Wineland. Trapped-Ion Quantum Logic Gates Based on Oscillating Magnetic Fields. Phys. Rev. Lett., vol. 101, no. 9, page 090502, 2008.
- [Ozeri 07] R. Ozeri, W.M. Itano, R.B. Blakestad, J. Britton, J. Chiaverini, J.D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J.H. Wesenberg & D.J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. Phys. Rev. A, vol. 75, 2007.
- [Paul 90] W. Paul. Electromagnetic traps for charged and neutral particles. Rev. Mod. Phys., vol. 62, pages 531–540, 1990.

- [Paul 92] W. Paul. Laser Manipulation of Atoms and Ions. In E. Arimondo, W.D. Phillips & Strumia F., editeurs, Proceedings of the International School of Physics, Enrico Fermi, Course 118, 1991. Elsevier, Amsterdam, 1992.
- [Penrose 55] R. Penrose. A generalized inverse for matrices. Proc. Camb. Philos. Soc., vol. 51, pages 406–413, 1955.
- [Poulsen 75] O. Poulsen, T. Andersen & 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: Atom. Mol. Phys., vol. 8, pages 1393–1405, 1975.
- [Rainville 04] S. Rainville, J.K. Thompson & D.E. Pritchard. An Ion Balance for Ultra-High-Precision Atomic Mass Measurements. Science, vol. 303, pages 334 – 338, 2004.
- [Reichle 06] R. Reichle, D. Leibfried, R.B. Blakestad, J. Britton, J.D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin & D.J. Wineland. Transport dynamics of single ions in segmented microstructured Paul trap arrays. Fortschr. Physik, vol. 54, no. 8-10, pages 666–685, 2006.
- [Renau 82] A. Renau, F.H. Read & J.N.H. Brunt. The charge-density method of solving electrostatic problems with and without the inclusion of space-charge. J. Phys. E.: Sci. Instrum., vol. 15, pages 347–354, 1982.
- [Resonetics] Resonetics. www.resonetics.com. Nashua, NH.
- [Roos 08] C.F. Roos, T. Monz, K. Kim, M. Riebe, H. Häffner, D.F.V. James & R. Blatt. Nonlinear coupling of continuous variables at the single quantum level. Phys. Rev. Lett., vol. 77, page 040302, 2008.
- [Rosenband 10] T. Rosenband. private communication, 2010.
- [Rowe 02] M.A. Rowe, A. Ben-Kish, B. Demarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W.M. Itano, B. Jelenkovic, C. Langer, T. Rosenband & D.J. Wineland. Transport of quantum states and separation of ions in a dual RF ion trap. Quant. Inf. Comput., vol. 2, pages 257–271, 2002.
- [Sadiku 09] M.N.O. Sadiku. Numerical techniques in electromagnetics. CRC Press, Boca Raton, 3rd edition, 2009.
- [Savard 97] T.A. Savard, K.M. OHara & J.E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. Phys. Rev. A, vol. 56, pages R1095–R1098, 1997.
- [Schulz 06] S. Schulz, U. Poschinger, K. Singer & F. Schmidt-Kaler. Optimization of segmented linear Paul traps and transport of stored particles. Fortschr. Physik, vol. 54, no. 5-10, pages 648–665, 2006.
- [Seidelin 06] S. Seidelin, J. Chiaverini, R. Reichle, J.J. Bollinger, D. Leibfried, J. Britton, J.H. Wesenberg, R.B. Blakestad, R.J. Epstein, D.B. Hume, W.M. Itano, J.D. Jost, C. Langer, R. Ozeri, N. Shiga & D.J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. Phys. Rev. Lett., vol. 96, page 253003, 2006.

- [Shor 94] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In S. Goldwasser, editeur, Proceedings of the 35th Annual Symposium on the Foundations of Computer Science, volume 35, page 124. IEEE Computer Society Press, 1994.
- [SIMION] SIMION. Scientific Instrument Service, Inc., www.simion.com. Ringoes, NJ.
- [Solano 99] E. Solano, R. L. de Matos Filho & N. Zagury. Deterministic Bell states and measurement of the motional state of two trapped ions. Phys. Rev. A, vol. 59, no. 6, pages R2539–R2543, 1999.
- [Solano 01] E. Solano, R.L. de Matos Filho & N. Zagury. Mesoscopic Superpositions of Vibronic Collective States of N Trapped Ions. Phys. Rev. Lett., vol. 87, no. 6, page 60402, 2001.
- [SolderMask] SolderMask. Inc., www.soldermask.com. Huntington Beach, CA.
- [Splatt 09] F. Splatt, M. Harlander, M. Brownnutt, F. Zähringer, R. Blatt & W. Hänsel. Deterministic reordering of $^{40}\text{Ca}^+$ ions in a linear segmented Paul trap. New J. Phys., vol. 11, page 103008, 2009.
- [Stenholm 86] Stig Stenholm. The semiclassical theory of laser cooling. Rev. Mod. Phys., vol. 58, no. 3, pages 699–739, Jul 1986.
- [Stick 06] D. Stick, W. K. Hensinger, S. Olmschenk, M. J. Madsen, K. Schwab & C. Monroe. Ion trap in a semiconductor chip. Nature Phys., vol. 2, pages 36–39, 2006.
- [Tietze 08] U. Tietze & Ch. Schenk. Electronic circuits: Handbook for design and applications. Springer, Berlin, 2008.
- [Turchette 00] 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 & D.J. Wineland. Heating of trapped ions from the quantum ground state. Phys. Rev. A, vol. 61, page 063418, 2000.
- [Van Dyck Jr 01] R.S. Van Dyck Jr, S.L. Zafonte & P.B. Schwinberg. Ultra-precise mass measurements using the UW-PTMS. Hyperfine Interactions, vol. 132, pages 163–175, 2001.
- [VanDevender 10] A.P. VanDevender, Y. Colombe, J. Amini, D Leibfried & D.J. Wineland. Efficient Fiber Optic Detection of Trapped Ion Fluorescence. arXiv:quant-ph/10040668v1, 2010.
- [von Busch 61] F.V. von Busch & W. Paul. Isotope Separation with the Electrical Mass Filter. Z. Phys., vol. 164, page 581, 1961.
- [Walther 93] H. Walther. Phase-transitions of stored laser-cooled ions. Adv. At. Mol. Phys., vol. 31, pages 137–182, 1993.
- [Wang 93] Y. Wang, J. Franzen & K.P. Wanczek. The non-linear resonance ion trap, Part 2. International Journal of Mass Spectrometry and Ion Processes, vol. 124, pages 125–144, 1993.

- [Wesenberg 07] J.H. Wesenberg, R.J. Epstein, D. Leibfried, R.B. Blakestad, J. Britton, J.P. Home, W.M. Itano, J.D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin & D.J. Wineland. Fluorescence during Doppler cooling of a single trapped atom. Phys. Rev. A, vol. 76, page 053416, 2007.
- [Wesenberg 08] J.H. Wesenberg. Electrostatics of surface-electrode ion traps. arXiv:0808.1623, 2008.
- [Wesenberg 09] J.H. Wesenberg. Ideal intersections for radio-frequency trap networks. arXiv:0802.3162v1, vol. 79, page 013416, 2009.
- [Wineland 78] D.J. Wineland, R.E. Drullinger & F.L. Walls. Radiation-pressure cooling of bound resonant absorbers. Phys. Rev. Lett., vol. 40, pages 1639–1642, 1978.
- [Wineland 79] D.J. Wineland & W.M. Itano. Laser cooling of atoms. Phys. Rev. A, vol. 20, pages 1521–1540, 1979.
- [Wineland 83] D.J. Wineland, J.J. Bollinger & W.M. Itano. Laser-Fluorescence Mass Spectroscopy. Phys. Rev. Lett., vol. 50, pages 628–631, 1983.
- [Wineland 98] D.J. Wineland, C. Monroe, W.M. Itano, D. Leibfried, B.E. King & D.M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. J. Res. Natl. Inst. Stand. Technol., vol. 103, pages 259–328, 1998.
- [Wineland 03] D.J. Wineland, M. Barrett, J. Britton, J. Chiaverini, B.L. DeMarco, W.M. Itano, B.M. Jelenkovic, C. Langer, D. Leibfried, V. Meyer, T. Rosenband & T. Schaetz. Quantum information processing with trapped ions. Phil. Trans. Royal Soc. London A, vol. 361, no. 1808, pages 1349–1361, 2003.
- [Wineland 05] D.J. Wineland, D. Leibfried, M.D. Barrett, A. Ben-kish, J.C. Bergquist, R.B. Blakestad, J.J. Bollinger, J. Britton, J. Chiaverini, B. Demarco, D. Hume, W.M. Itano, M. Jensen, J.D. Jost, E. Knill, J. Koelemeij, C. Langer, W. Os-kay, R. Ozeri, R. Reichle, T. Rosenband, T. Schaetz, P.O. Schmidt & S. Seidelin. Quantum control, quantum information processing, and quantum-limited metrology with trapped ions. Proc. Int. Conf. on Laser Spect., pages 393–402, 2005.

Appendix A

Glossary of acronyms and symbols

$ \uparrow\rangle$	${}^2S_{1/2} F = 1, m_F = -1\rangle$ state of ${}^9\text{Be}^+$
$ \downarrow\rangle$	${}^2S_{1/2} F = 2, m_F = -2\rangle$ state of ${}^9\text{Be}^+$
\doteq	An equality that will be constrained to be true
a_i	Mathieu equation a parameter that relates to static confinement
\hat{a}_i	harmonic motion annihilation operator along i^{th} direction
\hat{a}_i^\dagger	harmonic motion creation operator along i^{th} direction
$\mathbf{C}_1, \mathbf{C}_2$	a matrix and a vector, respectively, that define the constraints placed the trapping potential when solving for waveforms
COM	center-of-mass motional mode of multiple ions, where every ions' motion is in phase
η	the Lamb-Dicke parameter, the ratio of the size of the atom's wave packet to the wave length of light ($k_z z_0$)
η_i	trapping adiabaticity parameter
\hat{H}	total ion Hamiltonian
$\hat{H}_{\uparrow\downarrow}$	internal atomic state Hamiltonian
\hat{H}_M	harmonic motion Hamiltonian
\hat{H}_{int}	electric-dipole interaction Hamiltonian
\hbar	Planck's constant divided 2π
\mathcal{H}	the Hessian matrix of second derivatives, Eq. (3.6)

\mathbf{k}	laser wave vector
k	laser wave number
m	ion mass
LLS	Linear least-squares method for solving for trap voltages
n	harmonic-vibration quantum number
n_i	harmonic-vibration quantum number along the i^{th} direction
\bar{n}	average harmonic-vibration quantum number
Ω	the Rabi frequency
$\Omega_{n,m}$	the adjusted Rabi frequency for motional sideband flopping between n and m
Ω_{rf}	the frequency of the RF trap drive
ω_0	splitting between the $ \downarrow\rangle$ and $ \uparrow\rangle$ internal atomic states
ω_{L}	laser frequency
ω_{rf}	the pseudopotential contribution to the radial secular frequencies in a linear Paul trap (Eq. (3.17))
$\omega_x, \omega_y, \omega_z$	vibrational motional frequency along a given axis
$\check{\phi}_n(i, j, k)$	an array of potential values, resulting from 1 V being applied to the n^{th} electrode
$\check{\Phi}$	the total potential, including the $\check{\phi}_n(i, j, k)$ arrays
q	charge of an electron
q_i	Mathieu equation q parameter that relates to rf confinement
QIP	quantum information processing
rf	radio frequency
\hat{S}_i	spin (or analogous spin) operator in the i^{th} direction
\hat{S}_{\pm}	spin (or analogous spin) ladder operators
SVD	Singular value decomposition method for solving for trap voltages

- V** the vector of control voltages applied to the electrodes to produce a desired potential at the ion
- z_0 the spatial extent of the atom's ground-state wave function ($\sqrt{\frac{\hbar}{2m\omega_s}}$)

Appendix B

Methods for modeling trap potentials

In order to perform experiments with an ion trap, especially if transport protocols are to be used, it is imperative to have an accurate model of the potentials within the trap. This is also critical when designing the ion trap. In the most basic geometries, such as a perfect quadrupole trap, it may be possible to determine an analytical model. Also, in the more complicated surface-electrode geometries, a variation of the Biot-Savart law can be used to determine an approximate analytical solution [Wesenberg 08]. However, in general, generating a model of the potentials requires a computational electrostatic modeling program.

Though the rf Paul trap uses an inherently time-dependent potential, it is not necessary to use a time-dependent modeling program. Since the wavelength of the rf is much greater than the size of the ion trap, the dynamics of the rf potential due to retardation effects can be neglected. Therefore, it is easier to use a static potential modeling program. The modeled potentials can be directly used when working with control electrodes. In the case of the rf potential, the model assumes a static potential is being applied to the rf electrodes. Once the static potential is determined, the time dependence can be added by multiplying the static potential by $\cos \Omega_{\text{rf}} t$. However, for the purposes here, it is more convenient to approximate the rf potential as a static pseudopotential rather than a true time-dependent potential. Thus, the rf potential is typically converted to a pseudopotential using Eq. (5.2).

Three common methods for finding static electric potentials are: the relaxation method, the finite element method, and the boundary element method. All three were used at various points

during the work in this dissertation. When evaluating the modeling method, three key requirements become important. First, the method must be sufficiently accurate. Second, the method must be efficient (or at least fast). And, finally, the method must be user-friendly. There may be some additional factors that make one method better than another for a particular problem. Below, I will discuss these three methods and how well suited they were for the task at hand.

Relaxation Method

This is a technique [Jackson 99, Sadiku 09] that has traditionally been used for modeling ion traps, primarily due to its simplicity and the availability of commercial programs intended for ion optics, such as [SIMION].

The geometry of the trap is entered into the software and the first step is to divide the relevant spatial region into a cubic grid. The potential will be calculated at each point on that grid. All points start at 0 V (in the standard routine). The user then defines what voltage will be applied to each electrode. Any grid points that lie within an electrode will be fixed to the assigned voltage. The calculation proceeds point by point through the grid. For every point that is not constrained to a specific voltage, the potential at that point is set to be equal to the average of the potential at that point's six nearest neighbors. When the algorithm reaches the last point in the grid, the process is repeated using the new values of the potentials. These iterations proceed until the potentials converge, asymptotically, to a steady solution. This procedure produces a solution that fulfill the equations of electrostatics. The algorithm is terminated when the marginal change between iterations falls below some desired threshold. The algorithm can be tweaked to provide faster convergence.

This is called the relaxation method because it starts with a potential where the difference between neighboring points is un-physically large. Then the algorithm relaxes that potential until the lowest possible potential gradients are found.

This method has some significant drawbacks. First, some geometries are poorly suited for this approach and converge extremely slowly. For example, inside a hollow cube held at 1 V, the actual potential will also be uniformly at 1 V. However, as the algorithm progresses a significant number

of iterations are needed for the 1 V to propagate from the cube walls to the center. The marginal change between iterations at the center can be small. This can lead to premature termination of the calculation and the associated inaccuracies. Another problem is this method maps the geometry onto a rectangular grid. If the geometry includes curves or non-perpendicular angles, this mapping can lead to errors. The errors can be alleviated by using a sufficiently small grid but that requires additional computational resources and time.

Finite Element Method

The finite element method (FEM) [Jackson 99, Sadiku 09] is an extremely powerful and widely used method. This technique can be applied to a large variety of problems, well beyond electrostatics. These include mechanical stresses, fluid flow, thermal dissipation, and even simultaneous combinations of different effects. Given the wide applicability of FEM, it is a common tool for engineers. There are a number of commercially available FEM software packages, such as [ANSYS], that can cost tens of thousands of dollars and take years to master. Of course, it can also be used to solve our relatively simple electrostatics problem.

The FEM algorithm requires the volume of the model to be divided into small cells in a process called meshing. These cells are often either tetrahedrons or hexahedrons. The cells are not typically uniform in shape nor size, though the cells typically have the same number of faces. Each cell consists of vertices and edges. The potentials will be calculated at each vertex. In addition to the potential at each vertex, a linear equation will be defined along each edge of every cell. These equations depend on the cell geometry and the physical laws governing the system at hand (in the trap potential problem that would be Laplace's equation). These equations define the relationship between the vertex potentials located at the ends of the given edge. Dirichlet boundary conditions are defined for vertices that lie on electrode surfaces. Then a large matrix equation is generated using the vertex values and edge equations. This matrix equation is solved to yield the potentials at each vertex.

If one has access to one of the expensive commercial software packages, this method can be a good choice for finding the potentials. FEM has many strengths. It easily conforms to electrode

surfaces that contain odd angles or curved surfaces. This technique can also handle dielectrics and even electrodynamics problems. In fact, FEM can even solve hybrid problems that involve multiple physical phenomena, such as finding the potentials in a trap that is flexing under thermal stress.

However FEM also has a number of drawbacks. The entire validity of the solution hinges on how well the space was meshed. This meshing process is an art form and requires significant experience to achieve optimal results. For example, for the current problem of modeling an ion trap, the user is only interested in the potentials near the trap axis, where the ions are located. This may entice the user to mesh this region with a smaller mesh to obtain more accurate results in this area, while meshing the rest of the volume with a more coarse mesh. However, because the potentials originate at the electrodes, they must be propagated through this coarse region before they affect the region near the axis. Thus, if the user succumbs to this temptation, the accuracy near the axis will be no better than it is in the coarse region. Care must be taken to make sure a good mesh is achieved throughout the volume.

This is also not the most efficient method, as it must use resources to solve for the potential in regions far from the ion (and, thus, of little interest to the user). If a good mesh is needed over a large volume, the required amount of computer memory, as well as the time needed, can be quite large. These problems can easily become unwieldy and take hours or days to solve on a standard PC. In theory, FEM is a strong tool for solving for potentials in a trap, but in practice it can be very challenging.

Boundary Element Method

The boundary element method (BEM) [Sadiku 09] is another tool for solving electrostatics problems. Like FEM, this method can be used for a wide variety of physics problems. In general it can only be used on problems for which Green's functions can be calculated. The electrostatic problem happens to be one of the problems for which BEM is well suited.

Instead of meshing the entire volume of the model, BEM meshes only the surfaces of the electrodes, using two-dimensional cells (that are located in three-dimensional space). Each vertex N is described by three parameters: the position \mathbf{r}_N , the potential ϕ_N , and an effective point

charge q_N . The user defines the potential on each electrode surface, which is then mapped on to these vertices. The BEM algorithm will use the positions and the potentials to solve for the point charges. Coulomb's law relates the potential at a given vertex, ϕ_N , to the charge at every other vertex:

$$\phi_N = \frac{1}{4\pi\epsilon_0} \sum_{n \neq N} \frac{q_n}{|\mathbf{r}_N - \mathbf{r}_n|}, \quad (\text{B.1})$$

where the interaction of a charge with itself is neglected. Here, ϵ_0 is the vacuum permittivity. Using Eq. (B.1) at each vertex, a large matrix equation can be constructed and solved to find the set of q_n that satisfy the user-defined ϕ_n 's.

The potential in the vicinity of the ion is calculated in a second step after the set of q_n has been determined. This begins with the user defining a grid of points where the potential will be calculated. Coulomb's law is used again to find the potential at each grid point by summing the contributions from each effective charge q_n given the distance between the charge and the grid point.

This leads to an accurate model of the potential near the ion, as long as the mesh vertices on the electrode surfaces are sufficiently dense. However, since the charges are defined at the vertices only, and not across the entire surfaces of the electrodes, solving for the potential near an electrode surface can lead to large errors.

Variations on this method can be used. One particularly useful variation, that was used here, is to replace the point charges (calculated on mesh vertices) with charge patches (calculated in the area of the mesh cells) [Renau 82]. This method is essentially the same except Coulomb's law is updated for a charge patch and the algorithm solves for the uniform charge density within each patch using the potentials defined at points. For a similar-sized matrix, this leads to more accurate solutions, especially near the electrode surfaces.

BEM is an attractive alternative to FEM. Like FEM, it can be used on electrodes with odd angles or curved surfaces, as long as the mesh is fine enough. However, since only the electrode surfaces are being meshed, the computational resources are much smaller than FEM. Furthermore,

the likelihood of a poor mesh occurring is much smaller and this issue requires significantly less attention in BEM than it did in FEM. Commercial packages, such as [CPO], are available and are much simpler than their FEM counterparts. However, introducing dielectrics is more difficult in BEM. Also FEM is better suited for many hybrid problems involving multiple physical effects. A comparison of the accuracy of the methods can be found in [Cubric 99].

The BEM method proved to be the best suited of the three methods for the problem of finding trap potentials. It was efficient, accurate, and easy to use.

Appendix C

Possible linear constraints

In Sec. 5.3, the SVD technique is used to solve Eq. (5.14) for a set of voltages that would produce a potential satisfying some applied constraints. These constraints are introduced into Eq. (5.14) by including one row per constraint into \mathbf{C}_1 and \mathbf{C}_2 . It was shown how Eqs. (5.15) and (5.16) placed constraints on the trapping position while Eq. (5.18) constrained the motional frequencies and forced the principal axes to lie along \hat{x} , \hat{y} and \hat{z} . These are not the only possible constraints that can be realized using \mathbf{C}_1 and \mathbf{C}_2 . This appendix will cover select additional constraints, some of which can be useful in certain situations while others have little practical applicability (but are instructive to cover, nonetheless).

Perhaps the simplest constraint is implemented by including the row

$$\mathbf{C}_1 = \left[\begin{array}{cccccccccccc} 1 & \vdots & 0 & 0 & 0 & \vdots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right] \text{ and } \mathbf{C}_2 = [a], \quad (\text{C.1})$$

which enforces the constraint $\Phi(\mathbf{r}_0) \doteq a$. This particular constraint is valuable primarily as an example, since the value of $\Phi(\mathbf{r}_0)$ does not have a direct role in the ion dynamics nor play role in transport.

A more functional constraint would be to set the difference between (the square of) two

secular frequencies to $\omega_x^2 - \omega_z^2 \doteq w^2$, by using

$$\mathbf{C}_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{C}_2 = \begin{bmatrix} \frac{m}{q} w^2 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (\text{C.2})$$

The first line constrains the squares of the two frequencies. The last three lines force the principal axes to align along the directions that the derivatives are taken. Without the last three lines, the principal axes are free to rotate to a different orientation and the first line may no longer be constraining independent secular modes but would, instead, be constraining some arbitrary mixture of modes.

The above constraint has been helpful in a few instances. In cases where Eq. (5.14) is poorly conditioned it is helpful to introduce additional constraints that direct the solution to a more stable parameter space, and this constraint can do that in certain circumstances. One such example is at the center of the X-junction. At this location, the trap pseudopotential is identical along the axes (\hat{x} and \hat{z}) of the two intersecting linear arrays (assuming no misalignment). If the control electrodes are grounded, then $\omega_x = \omega_z$ by symmetry. If control potentials are included, it is not possible to split these two frequencies by more than ~ 1 MHz due to geometry constraints. At the same time, it is important that $\omega_x \neq \omega_z$ so that the trap is not degenerate in these two directions. This defines a narrow band of possible solutions. It proved convenient to use the above constraint to ensure the two secular frequencies were close, but not too close ($w \doteq 2\pi \times 0.5$ MHz).

In addition, a huge variety of rather odd, and mostly unhelpful, constraints can be constructed by arbitrarily adjusting other sets of coefficients in \mathbf{C}_1 and \mathbf{C}_2 . For example,

$$\mathbf{C}_1 = \begin{bmatrix} 1 & 0 & 0 & -1 & 0 & 0 & 0 & \alpha & 0 & \beta & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{C}_2 = [d], \quad (\text{C.3})$$

would set up the constraint

$$\left(1 - \frac{\partial}{\partial z} + \alpha \frac{\partial^2}{\partial y \partial x} + \beta \frac{\partial^2}{\partial y \partial z} \right) \Phi(\mathbf{r}_0) \doteq d. \quad (\text{C.4})$$

This constraint is likely of no practical importance but does demonstrate the types of constraints that can be easily implemented using \mathbf{C}_1 and \mathbf{C}_2 .

Instead of the simple constraints considered above, more complex constraints can be constructed by carefully considering the system.

In Sec. 5.3, Eq. (5.18) simultaneously placed six constraints on the harmonic potential: three constrained the secular frequencies and another three constrained the orientation of the principal axes to lie along \hat{x} , \hat{y} and \hat{z} . However, this orientation may not be convenient and the ability to constrain the principal axes to lie along some alternate set of orthogonal directions can be helpful. The proper set of constraints can be determined by rotating the Hessian constraints of Eq. (5.8). Any arbitrary rotation can be defined by three separate rotations: one by angle θ about \hat{z} , one by angle ϕ about \hat{x} , and again a rotation by angle ψ about \hat{z} :

$$\begin{aligned} \mathbf{R} &\equiv \begin{bmatrix} \cos \psi & \sin \psi & 0 \\ -\sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \phi & \sin \phi \\ 0 & -\sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} \cos \theta \cos \psi - \sin \theta \cos \phi \sin \psi & \sin \theta \cos \psi + \cos \theta \cos \phi \sin \psi & \sin \phi \sin \psi \\ -\cos \theta \sin \psi - \sin \theta \cos \phi \cos \psi & -\sin \theta \sin \psi + \cos \theta \cos \phi \cos \psi & \sin \phi \cos \psi \\ \sin \theta \sin \phi & -\cos \theta \sin \phi & \cos \phi \end{bmatrix}. \end{aligned} \quad (\text{C.5})$$

By setting these three angles properly, \mathbf{R} will rotate from the basis in which \mathcal{H} is calculated to the desired basis $(\hat{x}', \hat{y}', \hat{z}')$ in which \mathcal{H} is diagonal. Using this, the six constraints on the harmonic potential can be written as

$$\mathcal{H}(\mathbf{r}_0) \doteq m\mathbf{R}^{-1} \begin{bmatrix} \omega_{x'}^2 & 0 & 0 \\ 0 & \omega_{y'}^2 & 0 \\ 0 & 0 & \omega_{z'}^2 \end{bmatrix} \mathbf{R}. \quad (\text{C.6})$$

As in Eq. (5.8), the entries in matrix \mathbf{C}_1 are defined in the left hand side of the above equation. Expanding the right-hand side of Eq. (C.6) will give the entries for \mathbf{C}_2 . For three arbitrary angles, θ , ϕ , and ψ , these constraints can be rather long and unwieldy.

Since we are working in a linear Paul trap, one of the principal axes is almost always chosen to lie along the trap axis \hat{z} . In this case, two of the rotation angles can set to $\phi = \psi = 0$ and the rotation matrix reduces to

$$\mathbf{R} = \begin{bmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (\text{C.7})$$

This significantly simplifies Eq. (C.6),

$$\mathcal{H}(\mathbf{r}_0) \doteq m \begin{bmatrix} \omega_{x'}^2 \cos^2 \theta + \omega_{y'}^2 \sin^2 \theta & (\omega_{x'}^2 - \omega_{y'}^2) \cos \theta \sin \theta & 0 \\ (\omega_{x'}^2 - \omega_{y'}^2) \cos \theta \sin \theta & \omega_{y'}^2 \cos^2 \theta + \omega_{x'}^2 \sin^2 \theta & 0 \\ 0 & 0 & \omega_z^2 \end{bmatrix}. \quad (\text{C.8})$$

The rows of \mathbf{C}_1 do not change from what is reported in Eq. (5.18). The rows of \mathbf{C}_2 for this rotated set of principal axes become

$$\mathbf{C}_2 = \begin{bmatrix} (m/q) (\omega_{x'}^2 \cos^2 \theta + \omega_{y'}^2 \sin^2 \theta) \\ (m/q) (\omega_{y'}^2 \cos^2 \theta + \omega_{x'}^2 \sin^2 \theta) \\ \dots\dots\dots (m/q) \omega_z^2 \\ (m/q) (\omega_{x'}^2 - \omega_{y'}^2) \cos \theta \sin \theta \\ 0 \\ 0 \end{bmatrix}. \quad (\text{C.9})$$

The last two rows of \mathbf{C}_1 and \mathbf{C}_2 constrain the angles ϕ and ψ so that one of the principal axes points along \hat{z} . The third row constrains ω_z . Since the Hessian is being calculated in a basis that does not align with the principal axes, the other two secular frequencies are not separable in the Hessian's frame. Thus, these other three constraints are simultaneously and collectively constraining $\omega_{x'}$, $\omega_{y'}$, and the third angle θ .

In practice, this set of six constraints is too rigid, especially because physical and geometry constraints can also implicitly constrain the system. It may be desirable reduce these six constraints to a more flexible subset of constraints. One possibility is to constrain the orientation of the principal axis without constraining the secular frequencies. Assume again that one of the principal

axes should lie along the \hat{z} of \mathcal{H} . The last two rows in Eq. (5.18) and Eq. (C.9) are still useful for achieving this. The trick is finding a third constraint that will depend on θ but is independent of the secular frequencies.

The desired potential will still satisfy Eq. (C.8) even if the secular frequencies are not being directly defined. Three of the equations can be written explicitly as

$$\frac{\partial^2}{\partial x^2} \Phi(\mathbf{r}_0) = \frac{m}{q} (\omega_{x'}^2 \cos^2 \theta + \omega_{y'}^2 \sin^2 \theta) \quad (\text{C.10})$$

$$\frac{\partial^2}{\partial y^2} \Phi(\mathbf{r}_0) = \frac{m}{q} (\omega_{y'}^2 \cos^2 \theta + \omega_{x'}^2 \sin^2 \theta) \quad (\text{C.11})$$

$$\frac{\partial^2}{\partial x \partial y} \Phi(\mathbf{r}_0) = \frac{m}{q} (\omega_{x'}^2 - \omega_{y'}^2) \cos \theta \sin \theta. \quad (\text{C.12})$$

These three equations can be combined in the following manner to completely remove the dependence on the secular frequencies,

$$\left[\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} - \frac{2 \cos 2\theta}{\sin 2\theta} \frac{\partial^2}{\partial x \partial y} \right] \Phi(\mathbf{r}_0) = 0. \quad (\text{C.13})$$

Thus, using the constraint matrices

$$\mathbf{C}_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & \vdots & 1 & -\frac{2 \cos 2\theta}{\sin 2\theta} & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \vdots & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \vdots & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{C}_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (\text{C.14})$$

will constrain the principal axes to lie along $(\cos \theta)\hat{x} + (\sin \theta)\hat{y}$, $(\cos \theta)\hat{y} - (\sin \theta)\hat{x}$, and \hat{z} , without constraining the secular frequencies. Unfortunately, it is more difficult to do the opposite and constrain the frequencies without constraining the axes.

There are many constraints that can be realized using rows in \mathbf{C}_1 and \mathbf{C}_2 , but it may take some effort to determine exactly how to implement the constraints.

Appendix D

Maximum waveform step size

When solving for transport waveforms, a solution to the waveform constraints is found for several positions along the ion trajectory, following the trap axis. When the ion is between these calculated points, we assume that the potentials are still well behaved (i.e., the trapping well is located where it is intended to be and is harmonic). But this raises the question, what is the maximum possible distance between these calculated points for this assumption to hold?

To determine the maximum step size, assume a two-step waveform $\{\mathbf{V}_\alpha, \mathbf{V}_\beta\}$, with the waveform constraints satisfied at both points. In addition, assume the voltage applied to each electrode will ramp from the first value to the second value, with the same functional form. Here, we will examine the potential at an arbitrary point between the endpoints to determine how well that potential satisfies the appropriate constraints. For simplicity, the discussion will be restricted to one-dimensional transport and constraints, though the argument can be scaled to higher dimension.

From Eq. (5.3), the potential the ion experiences along the trap axis, z , is written as

$$\Phi(z, \tau) = \phi_{\text{ps}}(z) + \sum_{i=1}^N V_i(\tau) \phi_i(z), \quad (\text{D.1})$$

where τ refers to the step of the waveform sequence and is restricted to just two values, α and β . All τ dependence is contained in the voltages applied to each electrode, V_i . The first term, ϕ_{ps} , is the properly scaled pseudopotential. This term can be treated as a fixed electrode potential and combined into the sum as the zeroth term, where $V_0 = 1$.

Any linear combination of the two endpoint potentials can be written as a combination of

their respective V_i 's. For example,

$$(1 - \delta)\Phi(z, \alpha) + \delta\Phi(z, \beta) = \sum_{i=0}^N [(1 - \delta)V_i(\alpha) + \delta V_i(\beta)] \phi_i(z). \quad (\text{D.2})$$

Here, δ is assumed to be some fraction between 0 and 1.

We will apply the basic set of constraints: the slope of Φ is zero at both endpoints and the frequency is constrained. These constraints are written as:

$$\begin{aligned} \frac{d}{dz}\Phi(z_\alpha, \alpha) &= \sum_{i=0}^N V_i(\alpha) \frac{d}{dz}\phi_i(z_\alpha) = 0 \\ \frac{d^2}{dz^2}\Phi(z_\alpha, \alpha) &= m\omega_\alpha^2/q \\ \frac{d}{dz}\Phi(z_\beta, \beta) &= \sum_{i=0}^N V_i(\beta) \frac{d}{dz}\phi_i(z_\beta) = 0 \\ \frac{d^2}{dz^2}\Phi(z_\beta, \beta) &= m\omega_\beta^2/q. \end{aligned} \quad (\text{D.3})$$

Here, we reserve the option to have the harmonic frequency change from step to step in the waveform.

As the waveform steps from α to β , the voltage on each electrode will transition from the start value to the end value following some ramp function, $\kappa(t)$,

$$V_i(t) = V_i(\alpha) + (V_i(\beta) - V_i(\alpha))\kappa(t). \quad (\text{D.4})$$

This ramp function is ideally defined such that $\kappa(\alpha) = 0$ and $\kappa(\beta) = 1$. One possible ramp function is a simple linear ramp,

$$\kappa(t) = \begin{cases} 0 & \text{for } t < \alpha \\ \frac{t-\alpha}{\beta-\alpha} & \text{for } \alpha \leq t \leq \beta \\ 1 & \text{for } t > \beta \end{cases}. \quad (\text{D.5})$$

A more realistic choice for $\kappa(t)$ is an exponential ramp,

$$\kappa(t) = \begin{cases} 0 & \text{for } t < \alpha \\ 1 - e^{-\frac{t-\alpha}{RC}} & \text{for } t \geq \alpha \end{cases}, \quad (\text{D.6})$$

which is the ramp function that would result from placing a simple RC filter network applied to a step change in the control voltages. Ramp functions due to more complicated filter networks, as described in Sec. 6.2.2 would also work, as would an unfiltered step function which is the idealized DAC output.

Since these ramp functions are applied individually to the control electrodes, one might wonder if having each electrode ramp independently will cause the overall potential to be distorted in an undesirable way. Introducing Eq. (D.4) into Eq. (D.1) gives

$$\Phi(z, t) = \sum_{i=0}^N [V_i(\alpha) + (V_i(\beta) - V_i(\alpha)) \kappa(t)] \phi_i(z). \quad (\text{D.7})$$

Rearranging terms simplifies this to

$$\Phi(z, t) = \Phi(z, \alpha) + [\Phi(z, \beta) - \Phi(z, \alpha)] \kappa(t). \quad (\text{D.8})$$

Thus, applying the ramp to the individual electrodes results in the total potential ramping with $\kappa(t)$, as desired. Of course, if the ramp is not identical on each electrode (e.g., the RC filter constants differ between electrodes), this condition breaks down. In the X-junction array, the filter networks on each electrode were identical to within the components' tolerances.

To examine the potential between the two steps of the waveform, choose $\alpha \leq t \leq \beta$. In the ideal case, the trapping well should be located at $z_\kappa = z_\alpha + (z_\beta - z_\alpha)\kappa(t)$ and that the frequency should be $\omega_\kappa^2 = \omega_\alpha^2 + (\omega_\beta^2 - \omega_\alpha^2)\kappa(t)$, for any t . In this ideal situation, the waveform constraints in Eq. (D.3) should hold at the ideal ion location z_κ :

$$\begin{aligned} \frac{d}{dz} \Phi(z_\kappa, t) &= 0 \\ \frac{d^2}{dz^2} \Phi(z_\kappa, t) &= \frac{m}{q} \left[\omega_\alpha^2 + (\omega_\beta^2 - \omega_\alpha^2)\kappa(t) \right]. \end{aligned} \quad (\text{D.9})$$

The question becomes, to what extent do the above ideal equations hold?

It will be instructive to expand $\Phi(z, \alpha)$ and $\Phi(z, \beta)$ into their Taylor series components around the trapping minimum,

$$\begin{aligned} \Phi(z, \alpha) &= \psi_0^\alpha + \psi_1^\alpha(z - z_\alpha) + \psi_2^\alpha(z - z_\alpha)^2 + \psi_3^\alpha(z - z_\alpha)^3 + \dots \\ \Phi(z, \beta) &= \psi_0^\beta + \psi_1^\beta(z - z_\beta) + \psi_2^\beta(z - z_\beta)^2 + \psi_3^\beta(z - z_\beta)^3 + \dots \end{aligned} \quad (\text{D.10})$$

where ψ_i^τ is the coefficient for the i^{th} order term in the expansion. From Eq. (D.3), we know $\psi_1^\alpha = \psi_1^\beta = 0$ and the frequency constraint specifies both ψ_2^α and ψ_2^β .

To determine whether or not Eq. (D.9) holds, one must examine the various terms in $\Phi(z, t)$. It will help to define the distance between steps, $\Delta \equiv z_\beta - z_\alpha$ and change coordinates to $\chi \equiv z - z_\kappa = z - (z_\alpha + \kappa\Delta)$. Here $\kappa(t) \rightarrow \kappa$ because only one point in time is being examined. Equation (D.10) can be rewritten as

$$\begin{aligned}\Phi(z, \alpha) &= \sum_{i=0}^n \psi_i^\alpha (\chi + \kappa\Delta)^i \\ \Phi(z, \beta) &= \sum_{i=0}^n \psi_i^\beta (\chi - (1 - \kappa)\Delta)^i\end{aligned}\tag{D.11}$$

Equations (D.11) and (D.8) can now be combined to find $\Phi(\chi) = \sum_n \psi_n^\kappa \chi^n$ and determine the nature of the actual trapping well. Grouping terms of equal order in χ gives

$$\psi_n^\kappa = \sum_{j=0}^{\infty} \binom{j+n}{n} \left[(1 - \kappa)\kappa^j \psi_{j+n}^\alpha + (-1)^j \kappa(1 - \kappa)^j \psi_{j+n}^\beta \right] \Delta^j.\tag{D.12}$$

The leading order behavior for each coefficient is just $\left[(1 - \kappa)\psi_n^\alpha + \kappa\psi_n^\beta \right]$. This is an average of ψ_n^α and ψ_n^β weighted by the fraction of the ramp that has occurred and is exactly what would be desired in the ideal case. However, higher order corrections are also present, which cause deviations from that ideal case.

Notice that the leading order correction to ψ_n^κ , which is the $j = 1$ term in Eq. (D.12), is only non-zero if $\psi_{n+1}^\alpha \neq \psi_{n+1}^\beta$. Thus the leading order correction to the ion's position will be $\mathcal{O}(\Delta^2)$ if the frequency is held constant between waveform steps. The same is true of the correction to the frequency (squared) if the cubic term is constant.

To minimize all of the correction terms and achieve the most ideal situation, the first term of Eq. (D.12) must dominate the rest. A set of sufficient (though not necessary) conditions for this is

$$\begin{aligned}|\psi_n^\alpha| &\gg \left| \sum_{j=1}^{\infty} \binom{j+n}{n} \psi_{j+n}^\alpha \Delta^j \right| \\ |\psi_n^\beta| &\gg \left| \sum_{j=1}^{\infty} \binom{j+n}{n} \psi_{j+n}^\beta \Delta^j \right|.\end{aligned}\tag{D.13}$$

This raises the question of whether or not these sums converge, allowing these conditions to be satisfied at all. For trapping ions, the goal usually is to produce a strongly quadratic potential. In this case, there is some distance, D , from the minimum of the trapping potential and across this distance, the harmonic term will dominate shape of the potential. Using this, limits can be placed on the ψ_n^τ terms. Recasting Eq. (D.13) using D yields

$$|\psi_n^\alpha| \gg \left| \sum_{j=1}^{\infty} \binom{j+n}{n} \psi_{j+n}^\alpha D^j \left(\frac{\Delta}{D}\right)^j \right| \quad (\text{D.14})$$

$$|\psi_n^\beta| \gg \left| \sum_{j=1}^{\infty} \binom{j+n}{n} \psi_{j+n}^\beta D^j \left(\frac{\Delta}{D}\right)^j \right|.$$

Now one can see that two conditions are desired to minimize the corrections to the ideal potential between steps of the waveform: $\Delta < D$ and $\psi_2^\tau \gg \psi_{j+2}^\tau D^j$. The second condition is exactly the condition that the well be harmonic over distance D at both endpoints of the waveform, so it should already be fulfilled assuming the endpoint potentials were determined to make nice harmonic potentials. The first condition suggests a proper step size for the waveform: the step size should be on the order of, or smaller than, the extent across which the wells are dominantly quadratic. If the higher-order $|\psi_j^\tau D^j|$ fall off quickly with increasing j , it is permissible for Δ to exceed D without significant degradation of the waveform between steps. In the end, the exact limit on Δ will depend on the exact ψ_j^τ 's and on the extent to which the user is willing to accept anharmonicities. For the X-junction waveforms, $\Delta = 5 \mu\text{m} < D \approx 20 \mu\text{m}$, so the criterion holds.

Since the cubic and higher order terms should always be small, corrections to the secular frequencies will usually be negligible. On the other hand, the ψ_2^κ term will be large and there are waveforms for the X-junction array that involve a secular frequency constrained to change. This suggest the corrections to ψ_1^κ may be noticeable. The case $\psi_1^\kappa \neq 0$ will result in an unwanted displacement of the harmonic well from the desired location, along with a small change in well frequency due to mixing in of the higher-order terms in the potential. For a dominantly harmonic

well, the amplitude of that displacement, z' , is given by

$$z' = \frac{\psi_1^\kappa}{2\psi_2^\kappa} \approx \frac{(1-\kappa)\kappa [\psi_2^\alpha - \psi_2^\beta]}{[(1-\kappa)\psi_2^\alpha + \kappa\psi_2^\beta]} \Delta + \frac{3(1-\kappa)\kappa [\kappa\psi_3^\alpha + (1-\kappa)\psi_3^\beta]}{2[(1-\kappa)\psi_2^\alpha + \kappa\psi_2^\beta]} \Delta^2. \quad (\text{D.15})$$

If the secular frequency is held constant at both waveform endpoints, the first term vanishes and the second term should be small. In the case where the frequency does change, $|z'|$ will be largest when $\kappa = \frac{1}{2}$. A fractional curvature change can be defined, $p \equiv \frac{\psi_2^\beta - \psi_2^\alpha}{\psi_2^\alpha}$, and the maximum position offset is

$$z' \approx -\frac{p}{2p+4} \Delta. \quad (\text{D.16})$$

For small changes in p , the maximum offset from the intended position is $-\frac{p}{4} \Delta$. Converting this to small changes in the secular frequency, the offset is $\frac{\omega_\alpha - \omega_\beta}{2\omega_\alpha} \Delta$. The maximum offset in the X-junction waveforms would not have exceeded $0.4 \mu\text{m}$ and would have generally been much smaller.

Appendix E

Measuring motional excitation

Measuring the motional excitation of the ions is an important component to determining the quality of the transport protocol. The ions are initialized to a known motional state and then transported. Afterwards, the final motional state is measured using one of three methods: Doppler recoiling, Rabi flopping, and sideband asymmetry. This appendix will compare these three methods.

Doppler recoiling

The Doppler recoiling method [Epstein 07, Wesenberg 07] involves measuring the fluorescence of an ion illuminated with a Doppler beam tuned close to but slightly to the red of resonance. The scattering rate for a highly excited ion will be low, as it spends the majority of its time Doppler shifted from the laser beam's frequency, resulting in reduced scattering. However, with each scattering event, the ion loses energy and slows. As time progresses, the scattering rate will increase as the ion cools and the beam's frequency becomes better matched to the velocity of the ion. The scattering rate approaches a steady state as the ion nears the Doppler limit. By measuring the change in fluorescence as a function of time while the Doppler beam is applied, and averaging over many experiments, a fluorescence profile can be created. Assuming some motional distribution (e.g., a thermal distribution or a coherent distribution), this fluorescence curve can be fit to find the average initial motional occupation \bar{n} of the ion. If an incorrect motional distribution is assumed, the excitation determined may deviate from the correct value. A thermal distribution was always assumed when measuring transport excitation because the shape of the fluorescence curves

more closely match those predicted for a thermal distribution than for a coherent or Fock state.

The theory for fitting these fluorescence curves could only be calculated in one dimension [Wesenberg 07]. Therefore it must be assumed that one mode of motion (here the axial mode) is more excited than the others (the radial modes), so that the ion behaves like a one-dimensional oscillator. As discussed above this assumption seems reasonable. This method of measuring \bar{n} is not extremely accurate and agrees with more accurate methods within about a factor of two [Epstein 07].

In addition to a fluorescence curve, information about the on-resonance saturation parameter s_0 , the Doppler detuning δ , and ω_z is needed to perform the fit. The saturation s_0 is found by measuring the fluorescence as a function of different laser powers and extracting s_0 for the power that will be used. For the experiments performed on ${}^9\text{Be}^+$, these three values were typically $s_0 = 0.3$, $\delta = 2\pi \times 5$ MHz, and $\omega_z = 2\pi \times 3.6$ MHz.

Recooling is convenient because it only requires a single Doppler cooling beam. It can also be used to measure relatively high \bar{n} . For ${}^9\text{Be}^+$, this method provided reasonable fits for $75 \leq \bar{n} \leq 50,000$ quanta. This was the range of \bar{n} after transport for the initial transport experiments through the X-junction. However, as the transport results improved and \bar{n} dipped below 75 quanta, the fluorescence would return to steady-state too quickly to be measured. For transport heating rates below 75 quanta, multiple transports can be performed before measurement to accrue enough excitation to be measured.

When using the Doppler recooling method on a pair of ions in the same trapping well, the ions have twice the cooling rate. At the same time, there are two modes of motion (COM and stretch), with similar frequencies, so there is roughly twice the excitation to cool. These should roughly cancel, allowing the user to fit the curve as if it were a single ion with one mode of motion. This was done when fitting the motional excitation for ion pairs, but these assumptions will likely lead to a decrease in the accuracy of the measurement when performed on two ions, since the excitation may not be equal for both modes.

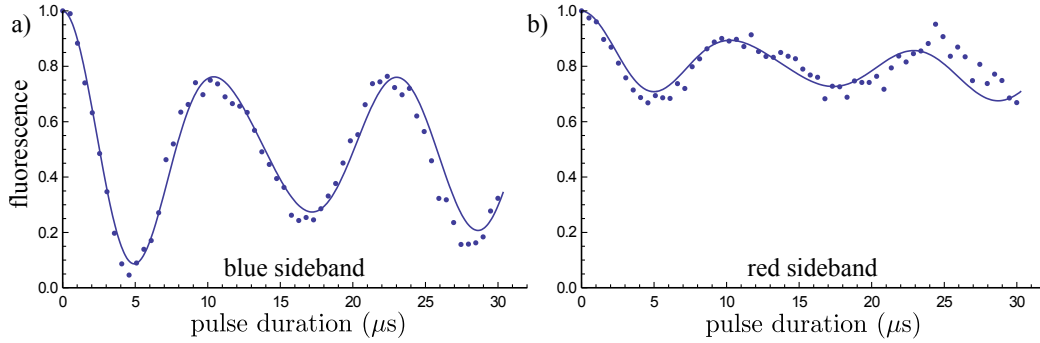


Figure E.1: Fits to a) blue sideband Raman flopping and b) red sideband Raman flopping with a single ion, after performing transport $\mathcal{E}\text{-}\mathcal{C}\text{-}\mathcal{E}$ using the new DAC filters. These fits assume a thermal distribution of $|n\rangle$. For the blue sideband, the extracted value is $\bar{n} = 0.42$ and for the red sideband it is $\bar{n} = 0.43$.

Rabi flopping

In the Rabi flopping method [Meekhof 96], the ions were first cooled to the ground state and then transported. Using the 90° Raman beams, two-state flopping curves on the “red” and “blue” sidebands or “carrier” of the $|F = 2, m_F = -2\rangle \leftrightarrow |1, -1\rangle$ ground state hyperfine transition were obtained. Since the 90° Raman beam transitions are sensitive to the motional state of the ion, the Rabi rates for each $|n\rangle$ will differ. By fitting the flopping curves with calibrated Rabi rates for the different motional Fock states, \bar{n} can be extracted [Meekhof 96, Rowe 02]. If the type of motional distribution is known (e.g., thermal or coherent state), this distribution can be assumed while fitting. However, the motional distribution after transport through a junction is not known, so several approaches were taken. The distribution was first assumed to be a thermal state, then a coherent state, and then some weighted sum of the two. Prior to installing the new filters to reduce DAC noise, none of these fits accurately reproduced the flopping features, though they all returned similar best-fit values of \bar{n} for the transports. For excitations of $\bar{n} \leq 2$, all distributions contain a significant fraction of the ground state $|0\rangle$, and the assumed distributions become difficult to distinguish, making the choice of distribution when fitting less important. After installing the new filters, the excitation was small enough that all distributions returned reasonable fits. Figure E.1 shows a sample fit.

The flopping curves were also fit using an arbitrary superposition of Fock states, which resulted in a reasonable reproduction of the data both before and after the new filters. Unfortunately, this method of fitting depends significantly on assumed ranges for n , since the Rabi rates for a given Fock state are not monotonic in n , nor unique [Meekhof 96, Wineland 98]. Thus population in a low $|n\rangle$ state can be incorrectly attributed to a larger $|n\rangle$ with the same Rabi rate, or vice versa. Attempts to fit fabricated test data using this method show that these fits likely over-estimate the actual ion temperature by up to a factor of 2. Nevertheless the reasonable agreement of the results based on the recoiling and sideband methods, lends confidence to the overall results.

The method of fitting Rabi flopping is best suited when the motional excitation is within the range of $0 \leq \bar{n} \leq 100$ quanta. It is not strongly affected by excitation in the other modes of motion, so can be used to measure axial excitation when the radial modes are also excited. It can also be used to measure the energy in a single mode of motion when using ion pairs.

Sideband asymmetry

The final method used for measuring the motional state of the ions after transport is described in [Turchette 00]. If an ion is in the state $|\downarrow\rangle \sum_n c_n |n\rangle$, with arbitrary c_n , and a red-sideband Raman pulse is applied, the probability P^{rsb} of driving the $|\downarrow\rangle |n\rangle \leftrightarrow |\uparrow\rangle |n-1\rangle$ transition will depend on the motional state $|n\rangle$. This probability can be measured experimentally by applying this pulse and then measuring the fluorescence of the ion in the BD light. If the probability is zero, the ion will remain in $|\downarrow\rangle$ and will fluoresce. If the probability is one, the ion will be driven to $|\uparrow\rangle$ and will not fluoresce. Thus the amount of fluorescence is a measure of P^{rsb} . Similarly, a blue-sideband Raman pulse, followed by detection, can measure the probability P^{bsb} of driving the $|\downarrow\rangle |n\rangle \leftrightarrow |\uparrow\rangle |n+1\rangle$ transition. However, since the red-sideband does not affect the $|\downarrow\rangle |0\rangle$ state and the blue-sideband does, these two probabilities are not symmetric. In fact, the ratio of these probabilities contains information about the motional state of the ion. In the specific case of a thermal distribution, where $c_n = \frac{\bar{n}^n}{(\bar{n}+1)^{n+1}}$ and $\bar{n} = (e^{\hbar\omega_z/k_B T} - 1)^{-1}$, the average excitation is given by

$$\bar{n} = \frac{P^{\text{rsb}}}{P^{\text{bsb}} - P^{\text{rsb}}}. \quad (\text{E.1})$$

method	Doppler recooling	Rabi flopping	Sideband asymmetry
useful range	$75 \leq \bar{n} \leq 50000$	$0 \leq \bar{n} \leq 100$	$0 \leq \bar{n} \leq 10$
requires fitting of data	yes	yes	no
assumes a distribution	yes	not necessary	yes
accuracy	ok	best	better

Table E.1: A brief comparison of the different techniques for measuring the motional excitation of an ion.

This method can be used to measure the average excitation of a thermal distribution within the range of $0 \leq \bar{n} \leq 10$ quanta. For a coherent distribution, this technique will give reasonable results for $\bar{n} \leq 1$ quanta, since the coherent and thermal distributions have significant overlap at low \bar{n} .