

Integrated Technologies and Control Techniques for Trapped Ion Array Architectures

by

Jules Michael Stuart

B.S. Physics, The University of Texas at Austin (2014)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

© Massachusetts Institute of Technology 2021. All rights reserved.

Author
Department of Physics
June 29, 2021

Certified by
Isaac L. Chuang
Professor of Physics and EECS
Thesis Supervisor

Certified by
John Chiaverini
Senior Staff at MIT Lincoln Laboratory
Thesis Supervisor

Accepted by
Deepto Chakrabarty
Associate Department Head of Physics

Integrated Technologies and Control Techniques for Trapped Ion Array Architectures

by

Jules Michael Stuart

Submitted to the Department of Physics
on June 29, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

Over the past few decades, trapped ions have emerged as one of the top contenders for realizing large scale quantum information processing. To date, experiments with ions have reached the level of tens of ion qubits, but the current model of adding ions to long chains may not extend to the number required for some computations. An alternative architecture, which has been recently explored, is to arrange ions in a large array, such that they can be shuffled around to transmit quantum information around a chip. This approach promises greatly increased numbers of qubits, while maintaining speed, fidelity and connectivity, but as the scale of these arrays increases, the required density of control systems may become intractable with current methods. In this thesis, we explore the integration of classical control technologies into ion traps and investigate whether this can provide the level of control needed to establish the array architecture as a more viable path toward more complex trapped-ion quantum computers. We focus first on the integration of classical, cryogenic electronics into an ion trap, which are used to control an ion's trap frequency and demonstrate rudimentary motion. An integrated switch allows the ions to be isolated from the effects of voltage noise. Next, we demonstrate the operation of a stimulated Brillouin scattering (SBS) laser for addressing an ion in an atomic clock protocol. In our experiment, the SBS laser is shown to have a linewidth commensurate with a bulk-cavity-stabilized laser and may offer a path towards generating highly-coherent light within an ion trap package. Subsequently, we explore the integration of photonic waveguides and grating couplers which can route laser light around on-chip and focus light onto ions trapped above the chip. The effects of stray electric field are considered, and the benefits of integrated light sources are characterized. In an array architecture, it will be important to be able to transport ions around between zones without introducing excessive motional decoherence. We present a technique using circuit simulation to pre-distort voltage waveforms for fast transport and demonstrate the basic operation of a trap designed to rapidly split and join ion chains. The studies covered here help inform future ion-trap architecture decisions and set the stage for further analysis of tradeoffs between these different technologies.

Thesis Supervisor: Isaac L. Chuang
Title: Professor of Physics and EECS

Thesis Supervisor: John Chiaverini
Title: Senior Staff at MIT Lincoln Laboratory

Acknowledgments

My path through graduate school has been one long adiabatic transition. It's easy to look through this thesis now and think, "Yes, of course these are all of the things I've been fortunate enough to work on." But had you told me five years ago that this would be the current state, I would be stymied (and probably quite stressed!) to find a way from that point to this one. Undoubtedly, it's a path I could not have taken without some incredible mentors, lab mates and friends.

Firstly, I must thank Professor Isaac Chuang, my research advisor at MIT, who is indeed the reason that I chose to come study here at MIT. I remember on the open house days sitting in his office and having a blast talking about the different gadgets he kept pulling from his shelves and sharing our enthusiasm for Arduinos. Since then, I have come to understand that Ike has an amazing ability to see and guide the large scale direction of his research while also never losing his passion for the small little interesting details that hold the whole thing together. In the course of my time at MIT, Ike has guided my path through research, supplying a constant well of new ideas and always offering his help along the way. Going forward into my career as a scientist, I hope that I am able to channel at least a fraction of his creativity and curiosity.

The majority of my research was conducted at Lincoln Laboratory, working with my co-advisor John Chiaverini. I will always associate John with the mental picture I have of him coming down to the lab when none of the rest of us could trap ions and making the trapping seem effortless. (Scan the 422 of course!) To me, John is an amazing resource for when I'm trying to interpret some kind of weird behavior of our experiments or when I'm looking for some insight into prior art—no one can call forth obscure ion experiments like John. Throughout this process, he has also been a great friend; our conversations about experiments would inevitably turn into fun chats about the state of quantum computing or how to improve my photography skills. In the future, when I ever need recommendations on places to see and experience at a place John has visited, I can trust him to give me the definitive list.

For most of my tenure at Lincoln, the ion group was co-led by Jeremy Sage, who in many ways was a co-co-advisor, even if his name isn't on the front page of this thesis. Jeremy has a kind of restless, enthusiastic approach to research which feels a lot like my own. He's worked on so many different projects and can describe each one of them to you like it's the next biggest thing in the field. It was always so much fun to brainstorm with Jeremy on our ideas and plans for experiments; the times I spent with him working on setting up our new cryosystem at MIT were truly some of my favorites throughout my graduate career. Though our shared love of music, we often went out to concerts together (seeing some amazing performances along the way) and hanging out afterwards, talking about music, life, science, and how I could convince my friends and family that I actually knew something about basketball. I don't know if we will ever find the optimal taco place + concert venue combination, but I hope to keep trying.

Early on in my time at MIT, a lot of my guidance came from fellow students, and for this I must thank Michael Gutierrez and Helena Zhang. Michael's ability to have a "full scope" picture of our lab was always inspiring to me. There was nothing in the lab that he couldn't explain or tweak up, and watching him build a new experiment from the ground up was awesome. I still refer to his notes from that period. I am thankful that he spent time during that period to help me out, and I have tried to emulate that behavior with the

students that I mentor now. Helena somehow wrote the control software for our whole lab and maintained it despite my attempts to break it with the clunky applets that I added. She was also always a good person to gripe with about the state of experiments or about the various hurdles we had to deal with in working at Lincoln.

Moving to Lincoln, I had close mentorship early on from Jonathon Sedlacek and William Loh. Jon's careful heating rates and tinfoil-hat noise-chasing remain the gold standard in my mind for obtaining solid anomalous heating results—not a small ask! I remember lots of evenings talking to Jon and speculating about the "many worlds" approach to surface treatment, discussing trashy pop music and gossiping about papers on the arXiv. I consider William to be my honorary academic advisor and also one of my best friends. In our long nights working on SBS lasers, we would talk about what it meant to have a career in academia or to be a good scientist...well, either that or we would discuss video games and speedruns. Outside of our discoveries in the lab, William and I discovered that the best pizza in Boston is the maple bacon at Dragon Pizza.

Colin Bruzewicz is the de facto chief of the ion lab at Lincoln, and he's also a great teacher. Many times in the course of my research we would go to the whiteboard or the notebook, and he always gave good insight. Later on, as I was more established in the lab, it was fun to talk with Colin about experiment plans or reconfigurations; he has a real knack for finding the "gotchas" and posing nice solutions. Earning Colin's approval and trust to work with his experiment felt like a significant step, like passing some kind of a lab qualifier. And while we may not have found the higher love quite yet, at least we can verify it's presence if we just think about it. Robert McConnell, the other permanent staff in the ion group during most of my time there, started up a new effort to develop a portable atomic clock, which was an effort I was always impressed to see and hear about from him, and one that I always liked to work on when I could. I really admire Robert as a physicist; he is not quick to refute points or interrupt with his opinions. Instead, he uses his intuition, makes some scribbles in the notebook in his lap, forms an argument and says, "Actually, there's no way that could be true." But he doesn't gloat about it and instead continues to wonder (or is it worry?) about what the real answer might be.

Towards the end, I got to work with another postdoc, Bob Niffenegger. Bob's readiness to go for the "real deal" ion trap quantum computer is inspiring. It's been cool chatting with him about the parallels he sees with his work at Intel and the work that we did together. Our conversations at NACTI 2019 were pretty unforgettable. In the clock effort, and in the push for lab automation, I also worked with Dave Reens, who is always amazing to watch in action. If it weren't for being worried about the timeline to graduation, I would have loved to have attended Dave's full-length seminar on how to design, machine, and clean a vacuum system from scratch.

Finally, for my MIT folks, I absolutely must thank Susanna Todaro who joined us at MIT just before this pandemic. It's been incredibly valuable to have had Susanna's friendship (virtually and in person!) throughout that last, strange year. I can't imagine what surviving the pandemic would have been like without her company. I also had the pleasure of mentoring Luke Qi, who was able to start with some hand-wavy explanations and a few citations and somehow make the both of us into ion motion experts. My weekly hackathons with Luke and our regular late night Slack conversations were always a treat. I'm looking forward to sharing some more tea with him soon. It is relieving to know that I am leaving the lab at MIT in truly capable hands, so I thank as well the new batch of graduate students, Jasmine Sinanan-Singh, Kyle DeBry, Xiaoyang Shi, Gabriel Mintzer, Felix Knollmann, and Trevor McCourt. I know you'll work together to get some barium

soon, and I can't wait to see Cheese in action.

Thanking all of the folks that made an impact on me during my time at MIT and Lincoln would take a whole second thesis, so I will try to keep the remainder short. Thanks to Professor Vladan Vuletić and Professor Will Oliver for serving on my committee. Thanks to my fellow students and friends Garrett Simon, Sam Buercklin, Jameson O'Reilly, Tony Zhang, Gavin West, and Phil Rich for their contributions and commiserations. Thanks to Zach Chin and Bianca Hanly for being amazing UROPs and frighteningly-quick studies. Thanks to Dave Kharas and Ryan Maxson for going to shows with me. Thanks to Alex Medeiros and Evan Golden for being around to chat between datasets. Thanks to John Barry, Adam Libson and Jeffery Birenbaum for their words of wisdom during experiments. Thanks to Cheri Sorace-Agaskar and Suraj Bramhavar and the rest of the photonics team for their amazing devices. Thanks to Rich Panock for his design work and for putting up with my pink PCBs. Thanks to May Kim for taking over computer duty and taking on the lab automation tasks at Lincoln. Thanks to Lee Mailhiot for some incredible soldering jobs. Thanks to Peter Murphy and Chris Thommouraj for some truly rapid packaging efforts. And of course, thanks to all the drivers of the MBTA.

Thanks to my undergraduate mentor, Greg Sitz, without whom I never would have made it to MIT. Thanks to my high school physics teacher, Rich Lines, who stoked the original spark of interest in physics. And thanks to my parents who were with me for every step of the way, even if it was sometimes virtual.

It's been real.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.

This material is based upon work supported by the Under Secretary of Defense for Research and Engineering and Department of Defense under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Under Secretary of Defense for Research and Engineering and Department of Defense.

© 2021 Massachusetts Institute of Technology.

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

For my mentors

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction to Trapped Ion Quantum Information Processing	23
1.1	A Brief History of Quantum Information with Trapped Ions	24
1.2	Why Choose Ions as Qubits?	25
1.2.1	Why Choose an Optical Qubit?	27
1.3	Looking Forward: Trapped Ion Array Architectures	27
1.4	Main Contributions of this Work	29
1.4.1	Integrated Voltage Control	29
1.4.2	Integrated Lasers for Controlling the Optical Qubit	31
1.4.3	Integrated Photonics for Laser Routing and Focusing	32
1.4.4	Techniques for Improving Ion Motion Protocols	32
1.5	Structure of this Thesis	33
1.6	Contributions of Co-workers and Published Papers	34
2	Trapping and Controlling Alkaline Earth Ions	37
2.1	Ionization	38
2.2	RF Paul Traps	38
2.2.1	Surface Electrode Traps	41
2.3	Controlling the DC Potential	43
2.4	Doppler Cooling	45
2.5	Quantum Operations	47
2.5.1	Quadrupole Interactions	47
2.5.2	Selection Rules	49
2.5.3	Sideband Transitions	50
2.5.4	Optical Pumping	51

2.5.5	Sideband Cooling	51
2.5.6	State Detection	52
2.5.7	Probing the Motional State	53
2.6	The Particulars of Strontium and Calcium	54
3	Experimental Apparatus: Lincoln Standard One	57
3.1	Cryogenic Vacuum Chamber	58
3.1.1	Feedthroughs and Wiring	62
3.1.2	Thermometry	62
3.2	Voltage Sources	64
3.2.1	DC Sources and Amplifiers	64
3.2.2	DC Filtering	65
3.2.3	RF Source and In-Vacuum Resonator	68
3.3	2D MOT	70
3.4	Laser Sources and Alignment	73
3.4.1	Atomic References	75
3.4.2	ULE Cavity System	76
3.4.3	Frequency Modulation and Focusing	79
3.5	Imaging System	81
3.6	Surface Electrode Trap Summary	84
3.6.1	"Lincoln Standard" Geometry	86
3.6.2	Voltage Simulation	87
3.7	The Brain: ARTIQ	88
4	Technical Challenges	89
4.1	Technical Electric Field Noise	89
4.1.1	Voltage Noise on DC Sources	91
4.1.2	Phase Noise on the RF Source	92
4.1.3	Stray DC Electric Fields	95
4.2	Anomalous Heating	96
4.2.1	Mitigation Techniques	97
4.2.2	Potential Sources	99
4.3	Qubit Decoherence	100

4.3.1	Laser Frequency Noise	101
4.3.2	Magnetic Field Noise	102
5	An Ion Trap with Integrated Voltage Sources	105
5.1	Integrated DAC Design and Testing	106
5.1.1	DAC Description	106
5.1.2	Testing Results	108
5.1.3	Design Revisions	109
5.2	Analog Switches and Voltage Noise	110
5.2.1	Analog Switch Design and Operation	110
5.2.2	Switch Characterization	111
5.2.3	Voltage Noise Measurements	113
5.3	Experiments with Ions	114
5.3.1	Rudimentary Motion and Frequency Control	116
5.3.2	Heating Rates	117
5.4	Trap Temperature	119
5.4.1	Thermometry Results	119
5.4.2	In-Situ Power Down	120
5.5	Output Bandwidth	121
6	A Compact Narrow-Linewidth Qubit Laser	123
6.1	Stimulated Brillouin Scattering Laser	124
6.1.1	Setup for Generating SBS Laser Light	126
6.1.2	Vibration Isolation	127
6.2	Temperature Feedforward Scheme	128
6.2.1	Extracting the Temperature Signal and Applying the Correction	129
6.2.2	Feedforward Simulation	130
6.2.3	Feedforward Results	132
6.3	Clock Operation	134
6.3.1	Measurement Process	134
6.3.2	Finding the Line from Zero	136
6.3.3	Results	138

7	Ion Traps with Integrated Photonics	141
7.1	Integrated Photonics Toolkit	142
7.1.1	Waveguides	143
7.1.2	Grating Couplers	144
7.1.3	Combining into an Ion Trap	146
7.2	Installing a Photonics Chip into a Cryogenic Apparatus	148
7.2.1	Attachment Procedure	150
7.2.2	Strain Relief	152
7.3	Technical Challenges	152
7.3.1	Trap Temperature	153
7.3.2	Charging and Technical Noise	153
7.3.3	Surface Treatment	154
7.4	Results	155
7.4.1	Single Qubit Operations	155
7.4.2	Improved Coherence with Integrated Photonics	158
7.5	Multiplexed Operations in Two Zones	160
8	Ion Motion and Planar Trap Designs for Array Architectures	163
8.1	Generating Voltage Waveforms	164
8.1.1	Ion Motion Primitives	164
8.1.2	Determining Optimal Trajectories	166
8.1.3	Creating and Simulating Voltage Waveforms	168
8.1.4	Predistortion	171
8.2	Optimizing Trap Designs	172
8.2.1	Choosing the Right Pitch	173
8.2.2	Tilt Considerations	173
8.2.3	The MaxBeta Trap	175
8.3	Experimental Results	177
8.3.1	Baseline Characterization	177
8.3.2	Linear Transport	179
8.3.3	Ion Behavior during Splits	182

9	Conclusions and Future Directions	185
9.1	Future Experiments	186
9.1.1	Junction Design and Operation	186
9.1.2	Faster Transport	187
9.1.3	Architecture Studies	188
9.2	Final Thoughts	189
A	High-Voltage Amplifier Design	191
A.1	Schematic and Layout	191
A.2	Testing Results	193
B	Ion Trap Voltage Generation	195
B.1	Generating Grid Files	195
B.2	Determining Taylor Coefficients	196
B.3	Voltage Optimization	197
B.4	Calculating the Total Potential and Generating Tilt	198
B.5	Translating and Applying Voltages	199

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

1-1	Ion trap array and unit cell concept	28
1-2	Illustration of prior work toward integrating control technology into ion traps	30
2-1	Standard linear Paul trap	39
2-2	Concept of a planar Paul trap	42
2-3	Concept of resolved sideband cooling	52
2-4	Level diagrams for strontium and calcium	54
3-1	Diagram of separable cryogenic vacuum system	59
3-2	Photos of the interior part of the cryostat	61
3-3	Photo of PCB octagon interposer	66
3-4	Filter motherboards	67
3-5	Schematic for the in-vacuum resonator	68
3-6	Diagram of the 2D MOT setup	70
3-7	Block diagram of the narrow-linewidth qubit laser	77
3-8	Diagram of diffraction-limited imaging system	81
3-9	Diagram of the Lincoln Standard trap	86
4-1	Motional heating rate of trapped ions versus trap temperature for a few different materials	98
5-1	Schematic of integrated DAC circuit	107
5-2	Illustration of DAC programming and example data	108
5-3	Electrode isolation switch design and concept	111
5-4	Voltage noise measured with the EIS open and closed	113
5-5	Layout of the ion trap chip with integrated DACs	115

5-6	Heating rates measured with the EIS open and closed	118
5-7	Measurement of trap temperature when using integrated voltage sources . .	120
6-1	SBS laser setup	126
6-2	SBS vibration isolation setup	128
6-3	Measuring and applying the temperature correction	129
6-4	Beat note measurement and feedforward simulation	131
6-5	Feedforward demonstration	133
6-6	Principle of operation of the clock	134
6-7	Initial spectroscopy and drift tracking	137
6-8	Clock operation	139
7-1	Integrated photonic structures	144
7-2	Ion trap with integrated photonic structures	146
7-3	Installation of integrated photonics ion trap into the cryogenic chamber . .	149
7-4	Beam profiles determined from laser-atom interaction with integrated pho- tonics trap	156
7-5	Protection from decoherence effect due to vibration when using integrated optics	159
7-6	Ion trap with multi-zone photonics	160
8-1	Example trajectories for ion transport	166
8-2	Schematic of the full control voltage circuit	169
8-3	Simulated ion trajectory during a simple transport operation	170
8-4	Predistortion of voltage waveforms to account for filtering effects	171
8-5	Comparison of outer-segmented and center-segmented surface electrode trap	172
8-6	Design of the MaxBeta trap	175
8-7	Loading and detection in the MaxBeta trap	177
8-8	Heating rate versus frequency on the MaxBeta trap	179
8-9	Transport verification experimental scheme	181
8-10	Spectroscopy during a split	182
9-1	Diagram of a trap with a junction	187

A-1	Schematic of the circuit for the high-voltage amplifier	192
A-2	Layout of the high-voltage amplifier	193
A-3	Step response of the voltage circuit	194

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

2.1	Landé g-factors for low-lying internal states of the ion	54
3.1	Dimensions of the Lincoln Standard trap	87
8.1	Dimensions of the MaxBeta trap	176

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction to Trapped Ion Quantum Information Processing

*“Rode here on the bus, now you’re one of us
It was magic hour”*

—Alvvays

The history of quantum mechanics and of modern computation are intertwined or—should we say?—entangled. The fundamentals of quantum theory emerged in the middle of the last century, at around the same time as the modern understanding of solid state physics developed, leading to the invention of the solid state transistor. A growing understanding of phenomena at the smallest of scales helped scientists and engineers at the time to make leaps and bounds in the creation of computers capable of performing increasingly-complex calculations. More computing power means more ability to process data, control experiments, simulate results and further our understanding of physics. These days, a physics student can use a processor with billions of transistors to integrate the Schrödinger equation and visualize the probability distribution of some experiment with a level of clarity that might have once been unimaginable. Viewed in this way, it’s only natural to think that the future will see even more complex methods of computation, and even richer understanding of quantum phenomena.

Richard Feynman famously put forth the concept of using a quantum system to model and simulate the behavior of other quantum systems [Fey82]. The idea garnered enough attention that mathematicians and physicists began to develop algorithms using an analog of

classical bits of information that were subject to the interference effects of the superposition states in quantum mechanics [BB14, Deu85, DJ92, Sho94, Gro96]. Actually implementing these proposed computations with physical systems took a few years; systems which have control down at the level of single quanta of information are usually delicate and sensitive to environmental perturbation. To date, we remain in an adolescent phase of the development of quantum computers which has been termed the "noisy intermediate scale quantum" regime [Pre18]. Considering that it will take orders of magnitude more qubits to perform calculations which might be useful (at least within our current understanding) and that it will take even more to perform these operations with error correction [CS96] to ensure high-fidelity results, the path from this point may seem unclear.

How can we ever hope to control such a complex and sensitive system at the level required for large scale computations? In this thesis we will attempt to answer this question by honing all the way down to the level single qubits interacting with control systems which have clear means for ramping up to many qubits. We will test the limit of behavior of our quantum systems as they are brought into close interaction with these noisy, classical control elements in a design which, if successful, may enable a path beyond the current scale of quantum computing.

1.1 A Brief History of Quantum Information with Trapped Ions

Trapped atomic ions are a strong contender for a system with the degree of control required for quantum computation. Ions can be cooled down to the ground state of their motion, and the spin of the ion's valence electron can be adjusted precisely. It may seem astonishing that we can perform routine experiments with single atoms, so how did the field get to this point?

The confinement of a charged particle using only a static configuration of electric fields is not allowed by Maxwell's equations. In the middle of the last century, techniques were developed to circumvent this rule by introducing other confining forces. Hans Dehmelt developed the Penning trap [Deh68] and Wolfgang Paul developed a trap with RF electric fields based on a quadrupole mass filter [Pau90]. With these techniques, it became possible to store only a handful of charged particles, though independent control of the ions remained

an outstanding problem due to the long wavelength of the microwave fields used to control their state.

With the invention of lasers, a number of new possibilities opened up. Laser cooling allowed the motion of ions to be controlled down to the level of μK and below [WI79, DBIW89]. With focused laser beams, the electronic state of individual ions in a trap could be controlled and read out, which enabled precision spectroscopy of states that were once inaccessible [WI81]. In David Wineland's group at NIST, the level of control over a single ion reached new heights with the demonstration of quantum logic gate using an ion's internal (electronic) and external (motional) state as two quantum bits [MMK⁺95]. From there, the demonstration of a quantum gate between two separate ions [SKHR⁺03] laid out all of the pieces that would be required to achieve quantum computation with even larger numbers of ions. As experimenters continue to refine the control techniques and eliminate or mitigate sources of noise, quantum gates with ions can be performed with both increasing speed [SBT⁺18] and increasing fidelity [GTL⁺16, BHL⁺16, SBK⁺21]. These quantum gates can be used to demonstrate quantum algorithms [MNM⁺16], and recently these efforts are approaching the point at which these experiments are able to perform useful calculations [WBD⁺19, NCP⁺20]. The precise level of control over the quantum states of trapped ions also allows them to be used as quantum sensors and as the basis for high-accuracy optical clocks [RKL⁺17, DL18, BCH⁺19].

The path towards a large system with many ions was rejuvenated with the advent of a planar trap [CBB⁺05]. With this architecture, modern fabrication techniques could be employed to design ion traps with complex structures at a scale which was once impossible [SCR⁺06, SHO⁺06, AUW⁺10]. At this point, the analogy to the growth of classical computers might seem clear, but the progress since then has not followed the same explosive trajectory as the transistor. While the path forward for the next decade may seem clear, there are many challenges which loom on the horizon and threaten to impede experimental progress.

1.2 Why Choose Ions as Qubits?

We often say that an ion is a natural choice for a qubit, but what exactly do we mean by this? Well, by "qubit" we simply mean a well-isolated quantum two-level system, often

treated as a spin-1/2 particle in a magnetic field. The fact that ions (or at least certain ions—and more precisely—the *valence electron* of certain ions) are quite literally spin-1/2 systems should explain why ions are "natural" qubits, but is it also true that ions are "good" qubits?

One way to answer the question is to consider a criteria for a physical quantum information processor put forth by DiVincenzo [DiV00]. In this criteria, a quantum computer should have:

1. A scalable physical system with well characterized qubits
2. The ability to initialize the state of the qubits
3. Long decoherence times, much longer than the gate time
4. A universal set of quantum gates
5. A qubit-specific measurement capability

In the past few decades of research in atomic physics, many aspects of these criteria have become routine. Certain ion qubit choices can have coherence times many orders of magnitude longer than gate times [WUZ⁺17], and focused lasers or microwave fields can be used to precisely control and initialize the electronic state. State-dependent fluorescence can then be used to measure electronic states with the highest fidelity of any qubit platform currently under investigation [MSW⁺08, HAB⁺14]. Ions also benefit from strong ion-ion coupling via the Coulomb interaction, which can be used to perform gates between ion qubits, thus rounding out the requirements for a universal set of quantum gates [HRB08].

Perhaps the only part of the criteria above we must make some effort to contrive is hidden in the first point, in that pesky little word "scalable." Going from a few ions in a small crystal to a chain of a few hundred comes with challenges [BCMS19], but moving even further beyond this, into groupings which can no longer fit in modern trap architectures seems at first to be insurmountable. In this thesis, we will consider one potential way forward to a larger picture. Rather than immediately charging ahead with an intractable number of ions, we instead consider what it will take to create a design which is naturally scalable, in the same way that an ion is naturally a qubit.

1.2.1 Why Choose an Optical Qubit?

There are several potential choices for an ion that we could use¹, so how do we settle on a species? First, we want to pick an ion with electronic states that can be easily manipulated, which generally requires an ion in which there is a single electron in the highest-occupied angular momentum orbital, such that we can approximate the behavior as that of a hydrogen atom. This limits the field to a few well-studied species [DeP21]. From here, we can further cull our choices by eliminating ion species which require highly ultraviolet ($< 380\text{ nm}$) radiation which may be difficult or expensive to produce². After this we are left with the alkaline earth ions, particularly calcium, strontium, and barium (and radium, which is usable [FWJ19], but requires special handling and can be a challenge due to finite half life). There is some freedom for preference at this point, and in this thesis, we use both strontium and calcium (see section 2.6), which have both been studied extensively.

Even after choosing an ion to use, there is still a choice to be made for the levels that will define the qubit. Ion qubits can be divided into a few categories [BCMS19]: (1) hyperfine qubits, (2) Zeeman qubits, (3) optical qubits and (4) fine structure qubits. There are advantages to each choice, but here we will focus on optical qubits, which are chosen for their high state detection fidelity and for their compatibility with integrated photonics.

1.3 Looking Forward: Trapped Ion Array Architectures

We’ve settled on a choice of ion, and we’ve got a few of them in our trap. We even have pretty good control over those ions, but now how are we going to get to the point that we have many hundreds? What about many millions?

Perhaps we can just make the standard linear trap longer and longer, such that it can accept a long chain of ions. Indeed this is the current approach of some efforts, but there are limits to how far this path can go. Though distant ions in a long chain share motional modes which can be used to send quantum information between pairs, as the length of the chain increases, these higher order modes bunch up in frequency space, leading to cross-talk, and the interaction strength between distant ions decreases, reducing gate times [MRR⁺14]. It

¹At least 118 choices, not counting the molecules, but trapping Seaborgium comes with its own set of challenges.

²Another reason for avoiding UV light is to ensure compatibility with integrated photonic structures, discussed further in Chapter 7.




Figure 1-1: (Left) An array based picture of a scalable ion trap quantum processor, called a "quantum CCD" architecture [Kie01]. (Right) Concept of an ion trap with control technology integrated beneath the surface of a trap. In the interior layers of the trap, there are dedicated areas for routing and focusing laser light, generating control voltages, and performing state detection.

thus becomes difficult to maintain high fidelity and high-speed connectivity. Transmitting quantum information by moving the ions themselves is also difficult since carefully-tuned swapping protocols are required to move ions down the chain [KRS⁺17]. One potential solution is to use photonic interconnects to entangle and teleport quantum information between several smaller ion chains [MK13, MRR⁺14], but remote entanglement currently takes long enough as to be a limitation in such an architecture [SNN⁺20].

Another idea to scale up the ion trap computer is to increase the density and connectivity of qubits by storing ions in a two-dimensional array [WMI⁺98, KMW02, LWF⁺17]. In this picture (Figure 1-1), the ions can be moved around between legs of the array in order to assemble small, manageable chains for performing quantum logic operations; afterwards, the qubits can be redistributed out to other locations. The array architecture is a framework which allows for a great deal of flexibility in the interactions that can be composed and executed. We note that, even outside of quantum computation, the array design and the concept of integrated control may have applications with portable clocks [DL18] or sensors [SDF⁺11].

A key challenge of the array approach is in getting the density of control required to manage all of the zones. With a surface trap, there is a potential to put control systems into the trap itself, in such a way that everything required to perform operations in a zone fit in an area less than or equal to the area of the relevant portion of the trap [RHS⁺20].

We call this concept a "unit cell," to denote a self sufficient building block that can then be combined with many other such cells. One can imagine some zones that might be well-suited for detection, with integrated detectors, and other zones which are designed to do multi-qubit interactions with integrated light delivery and optical modulators. Determining how best to lay out the cells in an array will require analysis of the tradeoffs of each control technology, and will likely end up as a compromise among device area, power, speed, and noise [BXN⁺17].

There has been great deal of prior work in the field of developing control technologies which can be integrated into an ion trap. A few particular results that pertain to work in this thesis will be discussed in detail in the relevant sections and below. For a thorough presentation of the work in the field, see the review articles [EGS⁺16] or [RHS⁺20].

1.4 Main Contributions of this Work

The main question we pose in the work presented here is: does the integration of control technology into an ion trap help establish the array approach as a viable path towards larger numbers of qubits? As we try to answer this question, we will consider whether our integrated systems perform at similar levels to what we achieve with our standard apparatus. We investigate how our delicate quantum systems behave in the presence of these added systems. In this process, we also look for ways in which integrated technologies pose key advantages which might be independent of the scalability they enable. Also along the lines of trying to improve the viability of the array architecture, we consider as well some techniques for developing the motion protocols that will be required for ensuring high connectivity in the array.

1.4.1 Integrated Voltage Control

Flexible control over the positions of ions in an array requires adjustable voltage sources to move ions around in the trap and to execute other reconfiguration operations, like splitting chains or reordering pairs. The time required to move ions can be a significant part of the overall operation time, so voltage sources must have fast update rates, on the order of the time required to perform quantum gates. However, since ions are sensitive to voltage noise, this must also be minimized in the design of the voltage source. The constraints of




Figure 1-2: Examples of prior work in the field of integrating control technologies into ion traps. (a) Decapsulated integrated circuit for a voltage source that was packaged onto an interposer near an ion trap and installed into a vacuum chamber [GFH⁺14]. (b) Ion trap fabricated in a commercial chip foundry process, demonstrating the compatibility with ion operations [MEB⁺14]. (c) High-stability laser system in the back of a van [LTBR11]. The laser sits in a volume much less than 1 m³ and can be operated while the van is in motion. (d) Reference cavity for a small high-stability laser which occupies a space less than 30 L for the whole system [DMM⁺18]. (e) Plan view of an ion trap with waveguides and diffractive gratings couplers for routing light on a chip and focusing onto ions trapped above the surface [MBM⁺16]. Light was coupled onto the chip by focusing laser light down onto diffraction gratings on the chip. (f) Side profile of light emitted from a diffractive grating coupler, used to couple light from waveguides into free space above the chip [MR17].

high speed and low noise are typically at odds, and often require the use of large filtering components, which may not be compatible with the increasing number of control voltages that will be required in larger ion arrays.

If the issue of voltage noise can be overcome, then one way to decrease the footprint of the control voltage system is to combine it with the trap itself. Before this work, the state-of-the-art for integrating voltage control into an ion trap system was the in-vacuum packaging of commercial voltage control hardware (Figure 1-2a) [GFH⁺14]. It had also been shown that ion trap could be fabricated in a CMOS foundry process (Figure 1-2b) [MEB⁺14], though this design did not include active electronics.

Here, we introduce the first ion trap chip with voltage sources embedded into the trap substrate with a voltage range and bandwidth compatible with doing rapid ion motion operations at our frequencies of interest [SPB⁺19]. Our design makes use of an analog

switch which allows for a high degree of voltage noise isolation between the source and the ion. This voltage source could be used to increase the number of independent controls on an ion trap beyond the level that is currently possible, enabling more densely-populated arrays.

1.4.2 Integrated Lasers for Controlling the Optical Qubit

Controlling the state of the optical qubit requires a laser source³ with an ultra-low frequency noise, equivalent to about 10^{-15} stability over one second. Typically this necessitates the use of an optical reference cavity, around 10 cm long, which must be kept isolated from its environment by several layers of acoustic and thermal shielding, often requiring a vacuum chamber. Noise on the laser which addresses the qubit can show up directly as infidelity in quantum gates, so it is not favorable to compromise the laser's stability to reduce the system size.

There has not been much effort to identify and develop a candidate for a qubit laser design which is compatible with monolithic integration. However, some work has been done to develop portable optical atomic clocks⁴ with small frequency sources [LTBR11, KGV⁺17, DL18]. Usually, these efforts proceed by miniaturizing existing components, so there is a limit to how small these lasers can be made, given by the size of the optical cavity and the chamber designed to hold it (Figure 1-2c-d).

Recently, work has taken place to develop a highly-frequency-stable laser based on an integrated resonator, which was demonstrated to have noise comparable to that of a standard qubit laser [LYOJ19]. In an experiment presented here, this integrated-photonic-compatible design is used to make a laser with which we implement a novel temperature compensation scheme to correct for temperature fluctuations down to the level of 200 nK. With this stabilized system, we were able to realize an optical atomic clock protocol, which is the first time this has been done with a laser that was not stabilized by a traditional frequency reference cavity [LSR⁺20]. Our architecture suggests a way toward a future qubit laser that may be realized at the chip scale.

³Here and through the rest of this thesis, the laser that addresses the qubit transition is sometimes called the "qubit laser."

⁴Optical atomic clocks, which use similar transitions to the one we choose to be our qubit, have similar stringent requirements on laser noise.

1.4.3 Integrated Photonics for Laser Routing and Focusing

In addition to the laser that addresses the state of the qubit, several lasers, over a spectrum of wavelengths (see section 2.6), are required to ionize and Doppler cool ions (see section 2.1 and 2.4) and to sideband cool to the ground state of motion (see section 2.5). To get high intensities at the ion, and to avoid cross-talk between neighboring ions in a chain, lasers must be focused down, ideally to the level of a few microns. Given the size of the vacuum chamber, it is challenging to realize small focused beam sizes, and these are usually limited to around $100\text{ }\mu\text{m}$ in diameter. As the number of lasers focused onto the trap increases in ever larger arrays, it will become a challenge to pack in all of the lenses required around the viewports on our chamber. Another issue with using external optics is that focusing optics, which are mounted some distance away from the ion, can pick up vibrations from the environment, which can translate to vibrations in the focused beam spot, thereby limiting the fidelity of qubit operations (see section 4.3.1).

Integrated photonics offer potential solutions to these problems [MSJ⁺21]. Prior to experiments here, the state of integration of laser routing and focusing had been recently advanced to the point that light addressing the optical qubit could be guided on and focused out of an ion trap chip (Figure 1-2e) [MBM⁺16]. Two outstanding challenges were to develop a platform which was compatible with wavelengths required for full control of the ion qubit and to come up with a way to couple light directly onto the chip from optical fiber.

In the course of this thesis work, an ion trap was developed, tested, and implemented which had integrated photonics structures for all required wavelengths [NSSA⁺20], which was particularly challenging due to the increased optical loss for blue wavelengths [SABK⁺18]. We also developed a packaging scheme for gluing a block of optical fibers to the edge of our chips which was able to maintain precise alignment even as we cooled the chips down to 4 K. The operation of this trap represents a critical step toward robust and portable trapped-ion-based clocks and quantum sensors, while demonstrating the promise of this technology for array architectures with many qubits.

1.4.4 Techniques for Improving Ion Motion Protocols

The ability to control the position of ions is a critical part of the design discussed in section 1.3. A large-computation may require many ion motion events, so any errors in

these operations can accumulate. Particularly, the total amount of motional excitation that ions acquire as they are transported must be limited, since this can reduce the fidelity of quantum operations between multiple ion qubits. One solution to the problem of additional excitation is to use additional ions which do not participate in quantum gates but instead help to cool back down to the ground state without affecting quantum information stored in the target ion⁵, but this takes process time, which can also accumulate as operations become more complex [BXN⁺17].

There has been significant effort to push the speed boundary while keeping motional excitation low [WZR⁺12, BGL⁺12], but these protocols often require precise timing to properly de-excite motion that is generated during the transport. Additionally, a split operation, which separates two ions into two separate wells, has not been demonstrated with an excitation less than one motional quanta, even though protocols exist to achieve this [PMGP⁺15]. Experiments with motion have been limited with how well control voltages can be matched to the prescriptions, and even slight deviations can lead to increased excitation.

Here, we introduce a new method for compensating for the effect of voltage noise filters in our system using a detailed numerical model of our circuit. We also present a trap design based on analysis conducted in [NH12] which should be able to perform split operations more rapidly than has been previously achieved. Together with this trap and the predistortion technique, it may be possible to realize a split with less motional excitation than has previously been demonstrated.

1.5 Structure of this Thesis

In the rest of this thesis, we first lay out the pieces required for trapping ions using our standard experimental apparatus, then move forward to consider in detail a few demonstrations of technologies we have explored, and then finally consider some trap designs and control techniques for realizing rapid motional operations. Here we outline this structure in detail:

Chapter 2 presents the theoretical concepts for trapping and controlling the optical qubit in alkaline earth ions. Many formulas are laid out which we will then reference throughout the rest of the work.

⁵This process of using additional ions of a different species to cool target ions is called sympathetic cooling [KKM⁺00].

Chapter 3 describes the experimental apparatus and discusses some relevant components in the lab. This chapter aims to be a comprehensive presentation of our system, which had previously only been described in bits and pieces.

Chapter 4 highlights the major technical challenges which typically lead to noise in our experiments. The concept of anomalous heating is presented, along with the techniques that we use to mitigate it.

Chapter 5 discusses experiments with an ion trap with integrated electronics. Several characterizations of noise are presented over the course of several design revisions. The work here overlaps strongly with [SPB⁺19].

Chapter 6 covers the design and implementation of a temperature-stabilized stimulated Brillouin scattering laser. This chapter overlaps with work reported in [LSR⁺20].

Chapter 7 discusses experiments with and characterizations of an ion trap with multi-wavelength integrated photonics. Here there is a strong overlap with [NSSA⁺20], though there is work presented at the end of this section which has not been published yet.

Chapter 8 describes the design of an ion trap created to perform ion motion operations with low-noise protocols. We lay out a general technique for accounting for filtering effects in our voltage sources, and present some preliminary results. The work in this segment is ongoing and will (hopefully) culminate in a publication soon.

Chapter 9 concludes with a few descriptions of some experiments and ideas that would make logical extensions of the work presented in the rest of the thesis.

1.6 Contributions of Co-workers and Published Papers

Many aspects of the work presented here were part of collaborative efforts between teams of researchers. In general, I have focused on the aspects of the various components for which I was most responsible⁶, but it is hard to avoid discussion of work performed by other team members. In particular, the work with the integrated electronics in Chapter 5 was enabled by the design work of Rich Panock. The work with the SBS laser described in Chapter 6 was done with equal contribution between William Loh and myself (though William is far and away the SBS guru). The integrated photonics work of Chapter 7 could

⁶Out of respect for my collaborators, and because I find it awkward to write this all in first person, I am sticking with "we" in the remainder of this work. In the cases where I was truly the only one to blame, then we will take this to mean you (the reader) and I.

not have happened without the integrated photonics team at Lincoln Laboratory, and most of the integrated-optics experiments with ions were performed by Bob Niffenegger, though we usually planned experiments and discussed results together. The ion motion work in Chapter 8 was done in close collaboration Luke Qi, a student whom I was mentoring; Luke handled the generation of the fancier waveforms and also developed the ODE integrator used to estimate the heat gained during motion operations.

Here is a list of the various publications which were produced during the course of the thesis. Several chapters of this thesis overlap strongly with the work presented in these articles.

- H. Zhang, M. Gutierrez, G.H. Low, R. Rines, **J. Stuart**, T. Wu, and I. Chuang, "Iterative precision measurement of branching ratios applied to 5P states in $^{88}\text{Sr}^+$ ", *New Journal of Physics*, **18**, 123021 (2016).
- C.D. Bruzewicz, R. McConnell, J.A. Sedlacek, **J. Stuart**, W. Loh, J. Sage, and J. Chiaverini, "High-fidelity, single-shot, quantum-logic assisted readout in a mixed-species ion chain", arXiv:1709.05102 (2017).
- J.A. Sedlacek, A. Greene, **J. Stuart**, R. McConnell, C.D. Bruzewicz, J.M. Sage, and J. Chiaverini, "Distance scaling of electric-field noise in a surface-electrode ion trap", *Physical Review A*, **97**, 020302 (2018).
- J.A. Sedlacek, **J. Stuart**, W. Loh, R. McConnell, C.D. Bruzewicz, J.M. Sage, and J. Chiaverini, "Method of determination of technical noise contributions to ion motional heating", *Journal of Applied Physics*, **124**, 214904 (2018).
- J.A. Sedlacek, **J. Stuart**, D.H. Slichter, C.D. Bruzewicz, R. McConnell, J.M. Sage, and J. Chiaverini, "Evidence for multiple mechanisms underlying surface electric-field noise in ion traps", *Physical Review A*, **98**, 063430 (2018).
- **J. Stuart**, R. Panock, C.D. Bruzewicz, J.A. Sedlacek, R. McConnell, I.L. Chuang, J.M. Sage, and J. Chiaverini, "Chip-integrated voltage sources for control of trapped ions", *Physical Review Applied*, **11**, 024010 (2019).
- C.D. Bruzewicz, R. McConnell, **J. Stuart**, J.M. Sage, and J. Chiaverini, "Dual-species, multi-qubit logic primitives for Ca^+/Sr^+ trapped-ion crystals", *npj Quantum*

Information, **5**, 1-10 (2019).

- R.J. Niffenegger, **J. Stuart**, C. Sorace-Agaskar, D. Kharas, S. Bramhavar, C.D. Bruzewicz, W. Loh, R.T. Maxson, R. McConnell, D. Reens, G.N. West, J.M. Sage, and J. Chiaverini, "Integrated multi-wavelength control of an ion qubit", *Nature*, **586**, 538-542 (2020).
- W. Loh*, **J. Stuart***, D. Reens, C.D. Bruzewicz, D. Braje, J. Chiaverini, P.W. Juodawlkis, J.M. Sage, and R. McConnell, "Operation of an optical atomic clock with a Brillouin laser subsystem", *Nature*, **588**, 244-249 (2020).

* - *equal contribution*

Chapter 2

Trapping and Controlling Alkaline Earth Ions

*“She is moving to describe the world
She has messages for everyone”*

—Talking Heads

So you want to precisely control the quantum state of the valence electron in an alkaline earth ion. The process of converting a hot vapor of neutral atoms to an ultracold crystal of just a few ions is an interesting journey from classical electrodynamics to fully-quantized laser-atom interactions, using many different theoretical tools which have been developed over the last several decades. In this chapter, we present a theoretical description for each of the steps that happen when loading ions into our trap, cooling down to the ground state of motion and then controlling the state of the optical qubit.

In section 2.1, we present our scheme for photoionizing neutral atoms. The theory of RF Paul traps is discussed in section 2.2, and the control of the confining potential is expanded in section 2.3. Section 2.4 covers the basics of Doppler cooling, and section 2.5 goes deeper into the laser control of the optical qubit, including sideband cooling and measuring the ion’s motional state. Finally, in section 2.6, we lay out relevant details of strontium and calcium, the two ion species covered in this thesis, including the wavelengths of the various lasers required to control the ion qubits.

2.1 Ionization

There are a few techniques for ionizing neutral atoms. In our experiments, we use two-stage photoionization (PI) which has a few advantages. For both strontium and calcium, the first stage of photoionization takes place on the $^1S_0 \leftrightarrow ^1P_1$ of the neutral atom. These transitions are a few tens of MHz broad, which does not require a lot of power to saturate, and they are also isotope-selective. Unlike laser ablation methods, which can potentially ionize other, undesired natural isotopes of our chosen metal, this method only produces the isotope for which we have tuned our laser to the appropriate mass shift. Photoionizing also allows us to produce ions precisely at the point where we focus our laser. This means that we can create ions at points where they will be instantly within the potential well of the trap confinement. Were it not for this, we might need to temporarily turn off our trapping fields to allow ions to enter into the space with the potential minimum before engaging the trap.

We use the dipole-allowed $^1S_0 \leftrightarrow ^1P_1$ transition in the neutral atom to perform an additional useful function for loading. By creating a two-dimensional magneto-optical trap with a laser field at this wavelength, we are able to pre-cool neutral atoms before they are trapped and then deliver them precisely to the trap center (see section 3.3).

For both ion species in this work, the second stage of photoionization is less restrictive. For calcium [LRH⁺04], we simply need to provide enough energy to get the excited electron into the continuum. For strontium [VCLB06], we are able to drive a second closed transition to a level which automatically ionizes, which is a much more power efficient process. To generate ions in our trap, we simply align both PI lasers to the desired location, switch them on, and then direct the pre-cooled neutral atoms towards the trap.

2.2 RF Paul Traps

It is not possible to confine a charged particle with static electric fields alone. This fact, a specific instance of Earnshaw’s theorem, comes from the satisfaction of Laplace’s equation $\nabla^2 V = 0$. From this equation, we see that a potential minimum (i.e. a point where the second derivative of the potential is negative) cannot simultaneously exist in all three dimensions; either the potential will be trivially zero or a saddle.

Two ways around the limitations of static electric fields are to either (1) use a combi-




Figure 2-1: (Left) Illustration of a Paul trap [CBB⁺05] (Right) Picture of blade trap [Chw09]

nation of static electric magnetic and electric fields in a Penning trap [Deh68] or (2) use a combination of static and oscillating fields in a Paul trap [Pau90]. Of these two possibilities, the Paul trap configuration is advantageous for quantum computation since ions in these traps are stable when confined to small oscillations about a point. In a Penning trap, ions are stable when undergoing cyclotron oscillations around a magnetic field axis. The constant motion of ions in Penning traps makes single-ion addressability a challenge.

The general form of the quadrupole potential in a Paul trap is [LBMW03]:

$$\Phi(x, y, z, t) = U \frac{1}{2} (\alpha x^2 + \beta y^2 + \gamma z^2) + V \cos(\Omega_{rf} t) (\alpha' x^2 + \beta' y^2 + \gamma' z^2) \quad (2.1)$$

The potential configuration we study here is based closely on the *linear* Paul trap, in which the oscillating part of the potential is defined by a four electrodes in quadrupole configuration. Here we assume that the electrode surfaces are all parallel to the $\hat{y}$ direction. As we will see, ions will tend to line up along the $\hat{y}$ axis, which we refer to this as the axial direction. We call the transverse direction the radial plane. In the linear Paul trap, we have $-(\alpha + \gamma) = \beta > 0$, $\alpha' = -\gamma'$, and $\beta' = 0$. The oscillating component of the field is a saddle point in the radial plane rotating along with the oscillating part of the potential at Ω_{rf} . The potential can be separated into an independent term for each of the spatial coordinates, which leads to three separate equations of motion, with the form:

$$\ddot{x} = -\frac{Z|e|\hbar}{m} \frac{\partial \Phi}{\partial x} = -\frac{Z|e|\hbar}{m} (U\alpha + V \cos(\Omega_{rf} t)\alpha')x \quad (2.2)$$

where m is the mass of the ion and $Z|e|$ is the total charge.

For the y (axial) direction, the equation simplifies to that of a 1D harmonic oscillator. Note that, without the confinement along the y direction, we would simply have a quadrupole mass filter, in which ions are able to slide freely along the central axis. The equations of motion in the radial directions can be simplified to the following form:

$$\frac{d^2x}{d\xi^2} + (a_x - 2q_x \cos(2\xi))x = 0 \quad (2.3)$$

by using the substitutions:

$$\xi = \frac{\Omega_{rf}t}{2}, a_x = \frac{4Z|e|U\alpha}{m\Omega_{rf}^2}, q_x = \frac{2Z|e|V\alpha'}{m\Omega_{rf}^2} \quad (2.4)$$

Equation 2.3 has the form of the Mathieu equation which has both stable and unstable solutions, depending on the values of a_x and q_x . The stability diagram has several regions in which either the x or z trajectory is independently stable [DB00]; however, we work exclusively in the lowest-order stability region in which:

$$0 < a_x \ll q_x^2 \ll 1 \quad (2.5)$$

The solutions of the Mathieu equation can be determined exactly from the Floquet theorem and in general consist of many oscillating terms [ŠB02, LBMW03]. To first order, the equation of motion in the x direction has the form:

$$x(t) \approx x_0 \left(1 + \frac{q_x}{2} \cos(\Omega_{rf}t) \right) \cos(\omega_x t + \phi_x) \quad (2.6)$$

where $\omega_x \approx \frac{\Omega_{rf}}{2} \sqrt{\frac{q_x^2}{2} + a_x}$ and x_0 is a term which depends on initial conditions. Equation 2.6 has a term oscillating at ω_x called the *secular motion* and a higher-order term with driven motion at Ω_{rf} which is called *micromotion*. Given the condition in equation 2.5, the secular frequency ω_x can be further approximated as $\omega_x \approx \frac{q_x \Omega_{rf}}{2\sqrt{2}} = \frac{qV\alpha'}{m\Omega_{rf}\sqrt{2}}$.

Another useful way to interpret the confinement in the radial directions is through the pseudopotential approximation[Deh68, Gho95], in which the fast terms oscillating at Ω_{rf} are averaged out [Bla12]. The effect of the oscillatory applied field can be described as an

effective pseudopotential:

$$\Phi_P(x, y, z) = \frac{q^2}{4m\Omega_{rf}^2} |\vec{E}^2(x, y, z)| \quad (2.7)$$

where the electric field $\vec{E}$ is that due only to the oscillatory potential. We will find the pseudopotential approximation to be a useful tool when calculating trap voltages (see section 2.3).

Excessive micromotion can be a source of noise in our experiments, and we usually tend to avoid it. Arranging potentials such that ions tend to stay near the null of the RF quadrupole reduces micromotion. In the linear trap, the RF null is a line along the central axis of the trap. In order to get chains of ions to line up along the axis, we require $\omega_y < \omega_x, \omega_z$. To increase the radial secular frequencies, we see that we can either increase the RF voltage or decrease the frequency; however, to make sure that our lowest-order approximation remains valid, we cannot lower the RF frequency indefinitely since we require $q_x \propto \Omega_{rf}^{-2} \ll 1$.

Ideally, the RF Paul trap is composed of hyperbolic surfaces which follow the equipotential planes of a quadrupole potential. In practice, this geometry is difficult to construct, but it turns out not to be strictly necessary. Two more common approaches are to make electrodes from rods arranged in a square or from rectangular strips arranged in an X shape in a so-called blade trap (see Figure 2-1). These kinds of three dimensional traps can be constructed using precisely machined electrode surfaces and spacers, but there is no clear path for traps like this to store and independently control the large number of ions that will eventually be required for useful computations.

2.2.1 Surface Electrode Traps

The same ponderomotive confinement of a Paul trap can also be obtained in a package with a more straightforward path to scalability by moving all of the trap electrodes into a planar geometry [CBB⁺05, SCR⁺06]. In this work, we use a so-called "five wire" design in which two RF electrodes sandwich a center electrode at RF ground, with two additional ground electrodes on the outside. The potential lines above a trap with this design are topologically similar to the potential within a standard four-rod trap. These types of traps are called planar traps or surface electrode traps (SETs).




Figure 2-2: (Left) Cross section of the pseudopotential above a five-wire style point trap [CBB⁺05] (Right) Isometric view of the pseudopotential minimum line above a surface electrode trap, also showing segmented electrodes which can be used to control the DC part of the ion's potential

With some approximations [Hou08], we can determine the potential above a planar trap analytically and apply the same pseudopotential approximation of equation 2.7 (see Figure 2-2 for some visualizations of the pseudopotential). The depth of the trap in the direction normal to the surface is reduced in a planar design, which is one of the major tradeoffs; these types of traps require more careful laser cooling and will not support trap lifetimes as long as could be achieved in a large 3D trap. Another disadvantage is that, as the shape of the electrodes move further away from the ideal hyperbolic shape, more instabilities are introduced [AHM⁺95], which may require *ad hoc* adjustment of trap parameters in order to trap ions¹. Surface traps have the potential to allow ions to be brought much closer to the electrode surface, which is nice for reducing overall size when scaling up, but reducing the ion-electrode distance has ramifications for noise (see section 4.2) and tends to make laser access more difficult when trying to avoid having lasers clip on the edges of the trap.

The advantages of surface electrode traps are numerous, particularly for the focus of this thesis. A planar design of metal electrodes can be created using photolithography [SCR⁺06, SHO⁺06, AUW⁺10, HLBH11, MKW⁺16, BMCS16], which allows for a high degree of control and reproducibility. Fine scale features, which may improve the control of potentials above the trap can be created with precision far greater than can be achieved

¹Fortunately, the stability diagram of the Mathieu equation is easy to explore by adjusting the value of (the Mathieu) q via the applied RF voltage amplitude. This is easy to adjust with a function generator (see section 3.2.3), so we will often adjust this knob in cases where ions are elusive.

with a macroscopic assembly. Using microfabricated traps also opens a door to including other technologies within the substrate of the trap, which will be the major focus of this thesis. Recently, many research efforts into scalable quantum computing have moved in the direction of surface traps, with several adopting a design by Sandia National Labs which attempts to incorporate the high degree of optical access afforded by larger 3D traps [LoEOoSI16]. In this thesis, all of the traps used were designed by our group. Most follow a standard geometry discussed in section 3.6 unless otherwise noted in the relevant sections.

2.3 Controlling the DC Potential

The total potential above a Paul trap can be written as:

$$\Phi(x, y, z) = \Phi_{RF}(x, y, z) + \Phi_C(x, y, z) \quad (2.8)$$

Where $\Phi_{RF}(x, y, z)$ is the potential due to the ponderomotive confinement of the RF electrode, with the pseudopotential approximation applied to remove the time dependence², and $\Phi_C(x, y, z)$ is the part of the potential due to the DC electric fields, here called the "control potential," which comes from the voltages applied to the various control electrodes³. At the desired position of the ion, we can perform a Taylor expansion of the control potential as [HS06]:

$$\Phi_C(x, y, z) = V_0 + E_x x + E_y y + E_z z + \alpha_{xx} x^2 + \alpha_{yy} y^2 + \alpha_{zz} z^2 + \alpha_{xz} xz + \beta_{yyyy} y^4 + \dots \quad (2.9)$$

Where we have chosen just a few of the lower-order terms in the expansion; in general, the expansion may require many higher-order terms⁴. At the trap position, we want for the E_i terms to go to zero, such that the ion experiences only symmetric, confining forces. It is possible to create a saddle point of $\Phi_C(x, y, z)$ at a position other than the minimum

²Note that $\Phi_{RF}(x, y, z)$ does not necessarily satisfy the Laplace equation, so it is possible for it to have a three dimensional minimum (a so-called *point* trap). For our linear Paul trap configuration, we expect harmonic confinement in the $\hat{x}$ and $\hat{z}$ directions.

³In the picture presented in section 2.2, this was treated as strictly a DC potential, but when we wish to do experiments which involve ion dynamics, these voltages and potentials are not constant. They still change on a time scale much slower than the period of the RF signal, but to call them "DC" voltages could lead to confusion. Here we are adopting the term "control" voltages/potentials/electrodes to refer to these quasi-static controls.

⁴The units of the coefficients in the expansion above may sometimes be inconsistent in the literature. Sometimes a charge factor is explicitly included to make Φ_C have units of energy. Later on, when we discuss the coefficient β_{yyyy} in more detail in section 8.2, we will use units of V/m^4 .

of $\Phi_{RF}(x, y, z)$, but this introduces additional, unwanted micromotion and will eventually prevent stable trapping if the minima are separated too far.

When the control potential is aligned to the RF null, the main effect of the RF will then be to provide harmonic confinement in the $\hat{x}$ and $\hat{z}$ directions. The total radial frequencies will be a combination of the frequency due to the pseudopotential and the frequency due to the applied control potential⁵. By introducing a nonzero α_{xz} term, the axes of the normal modes can be rotated away from the geometric axes of the trap; this is useful in order to ensure that all of the motional modes will have some projection along the propagation direction of the lasers used for cooling (discussed further in the next section). The introduction of the α_{xz} term is called "tilt."

The terms α_{yy} and E_y define the position and frequency of the axial confinement. By the symmetry of the RF electrode, the ponderomotive confinement has little effect in this direction. We are thus able to control the position of the ion dynamically by introducing a time-dependence to these terms. With the use of β_{yyyy} and other, higher-order terms, more complex dynamics can be obtained, like rotating chains of ions or splitting and joining multiple ions. These concepts will be discussed further in section 8.1.

The various terms in the expansion of $\Phi_C(x, y, z)$ can be controlled by the voltages applied to the control electrodes. With judicious choices of the geometry and separation of these control electrodes, we can ensure that the electric field projection the electrodes are sufficiently orthogonal that we are able to independently control the different coefficients. Generally, we have a set of target coefficients $\mathbf{C}_{target} = [E_x \ E_y \ E_z \ \alpha_{xx} \ \dots]$ and a set of control voltages $\mathbf{V} = [V_1 \ V_2 \ V_3 \ \dots]$ to adjust. Finding the ideal $\mathbf{V}$ can be framed as an optimization problem in which we try to minimize the distance $\|\mathbf{C} - \mathbf{C}_{target}\|$ where $\mathbf{C}$ is the realized set of coefficients given a particular voltage configuration. We typically have many more voltages to adjust than we have target coefficients, which leads to a problem that is underconstrained. To overcome this, we use a voltage optimization tool which uses a few other cost functions in order to find a unique solution in the otherwise degenerate space of acceptable voltage configurations. More details about the simulation of control potentials and the optimization of voltages are discussed in sections 3.6 and 8.1 and appendix B.

⁵By Earnshaw's theorem, one of α_{xx} or α_{zz} must be negative, but as long as the absolute value is lower than the square of the radial frequency from the pseudopotential, we will still get confinement along that axis.

2.4 Doppler Cooling

In this section, we describe the process by which we cool ions down below the trap depth. We take a few particulars of the atomic structure for granted here, namely the existence of two well-isolated levels, an excited state $|e\rangle$ and a ground state $|g\rangle$, separated by an energy difference commensurate with a photon in the optical regime. We assume that there is an electric dipole coupling between an external laser field and these two levels, with a coupling strength or rate given by Ω . We finally assume that there is some projection of all of the ion's modes of motion onto the direction of the cooling laser $\vec{k}$, which can be achieved using tilt as described above.

When we illuminate the ion with laser light which is near-resonant with the optical transition, the atom will absorb and then re-emit a photon. Both absorption and emission impart momentum to the atom, but the emission direction, and thus angle of the momentum kick, is random. Averaged over many scattering events, the atom will experience a net force given by [MvdS99]:

$$F = \hbar k \Gamma \rho_{ee} \quad (2.10)$$

Where the momentum of the photon is $\hbar k$, Γ is the decay rate from the excited state and ρ_{ee} is the population in the excited state. This population depends on the laser field detuning and the strength of the coupling; it is also susceptible to saturation effects, due to the finite rate of re-emission:

$$\rho_{ee} = \frac{1}{2} \frac{s_0}{1 + s_0 + (2\delta/\Gamma)^2} \quad s_0 \equiv \frac{2\Omega^2}{\Gamma^2} \quad (2.11)$$

where $\delta = \omega_L - \omega_{eg} - \vec{k} \cdot \vec{v}$ and $s_0 \equiv I/I_s$ is the saturation parameter. Here the detuning has a term to account for the Doppler shift of the atom as it moves in the laser field with velocity $\vec{v}$.

The velocity dependence of ρ_{ee} leads to a velocity-dependent damping force, given by:

$$F = F_0 - \beta v \quad (2.12)$$

where F_0 is the light force on an atom at rest and:

$$\beta = -\hbar k^2 \frac{4s_0(\delta/\Gamma)}{(1 + s_0 + (2\delta/\Gamma)^2)^2} \quad (2.13)$$

The maximum value of β is obtained when $\delta = -\Gamma/2$ and $s_0 = 2$, leading to:

$$\beta_{max} = \frac{\hbar k^2}{4} \quad (2.14)$$

Even though the average velocity of the atom is damped by the presence of the laser field, it is not possible for the atom to equilibrate to zero velocity, since there is still a random kick of momentum each time that the atom re-emits a photon. Equating the rate of rate of cooling to the rate of re-heating from random scattering leads to the Doppler limit:

$$k_B T_D = \frac{\hbar \Gamma}{2} \quad (2.15)$$

which occurs when $s_0 \ll 1$ and $\delta = -\frac{\Gamma}{2}$. Interestingly, the ultimate Doppler limit thus only depends on the decay rate Γ of the transition, as long as the laser intensity is sufficiently below saturation. For the atomic species in this work, we perform Doppler cooling on transitions with linewidths on the order of 20 MHz, for which $\Gamma \approx 1.2 \times 10^8$ Hz, leading to Doppler limits of around 450 μ K. For typical trap frequencies around 1 MHz, this corresponds to an average of about 10 motional quanta.

In the atomic species that we use, the excited state $P_{1/2}$ can decay to the ground state $S_{1/2}$ or to a metastable state $D_{3/2}$ with a smaller, but significant probability (see the level diagram in Figure 2-4). In steady state, the electron would end up getting stuck in this metastable level, so to keep the electron in the cycling transition that we use for Doppler cooling, we illuminate the ion with a second repumping laser resonant with the $P_{1/2} \leftrightarrow D_{3/2}$ transition. Since we do not usually care about maintaining population in the $D_{3/2}$ state, this repumping laser is typically left on at all times. While the polarization of the Doppler laser is not important, the polarization of the repumping beam must be at least partially perpendicular to the magnetic field (i.e. not strictly π light) in order to avoid trapping population in the stretched Zeeman sublevels. The dynamics of the Λ -type system can be complex [Lab08, CLO15]; one potentially unwanted phenomenon is coherent population trapping (CPT), in which electronic population ends up in a non-absorbing dark resonance

for particular values of laser detunings in the two relevant transitions. To avoid these resonances, which can limit the efficiency of Doppler cooling, we intentionally broaden the repumping laser to a linewidth of around 1 GHz. With this setup, we do not see the dip in Doppler spectroscopy which is associated with CPT.

2.5 Quantum Operations

We now move on to the laser which couples the two levels of our optical qubit. In our chosen ion species, this is a so-called electric quadrupole transition, sometimes called a "forbidden" transition (meaning that it is electric-dipole-forbidden). The transition is due to a high-order effect and is thus weak. As a result, the natural lifetime of an electronic state which can only decay by a forbidden transition can be quite long. These long-lived, or metastable, excited states make for good qubits. In this section, we lay out the physics for laser-atom interactions with these states.

2.5.1 Quadrupole Interactions

We focus now on two electronic energy levels, with an electric-quadrupole coupling. Without choosing any particular levels, we assume one of these levels is a stable ground state, $|g\rangle$ and the other is an metastable excited state, $|e\rangle$. The Hamiltonian for this simple two-level system is:

$$H_e = \hbar\omega_e|e\rangle\langle e| + \hbar\omega_g|g\rangle\langle g| \quad (2.16)$$

To simplify the Hamiltonian, we introduce the Pauli operator $\hat{\sigma}_z = |e\rangle\langle e| - |g\rangle\langle g|$. We define $\omega_{eg} = \omega_e - \omega_g$ and re-scale the total energy to get:

$$H_e = \frac{\hbar\omega_{eg}}{2}\hat{\sigma}_z \quad (2.17)$$

In the full Hamiltonian for the total energy of the ion, we must also account for the potential of the harmonic oscillator of the ion trap. Here we will focus on just a single motional mode with frequency ω_t . The total Hamiltonian becomes:

$$H_0 = \frac{\hbar\omega_{eg}}{2}\hat{\sigma}_z + \hbar\omega_t\left(a^\dagger a + \frac{1}{2}\right) \quad (2.18)$$

We introduce now a monochromatic laser field given by:

$$\vec{E}(\vec{r}, t) = \vec{E}_0 \left(e^{i(\vec{k} \cdot \vec{r} - \omega_L t + \phi)} + e^{-i(\vec{k} \cdot \vec{r} - \omega_L t + \phi)} \right) \quad (2.19)$$

The interaction Hamiltonian with the applied field is [Roo00]:

$$H_{int} = \frac{\hbar}{2} \Omega (\sigma^+ + \sigma^-) \left(e^{i(\vec{k} \cdot \vec{r} - \omega_L t + \phi)} + e^{-i(\vec{k} \cdot \vec{r} - \omega_L t + \phi)} \right) \quad (2.20)$$

where $\sigma^+ = |e\rangle\langle g|$ and $\sigma^- = |g\rangle\langle e|$. The strength of the interaction comes from the Rabi frequency Ω , which is [Lab08]:

$$\Omega = \left| \frac{eE_0}{2\hbar} \langle g, m | (\vec{\epsilon} \cdot \vec{r}) (\vec{k} \cdot \vec{r}) | e, m' \rangle \right| \quad (2.21)$$

where $\vec{\epsilon}$ is the polarization vector of the applied electric field. The evaluation of this coupling gives rise to selection rules which will be described further below.

We now make a few assumptions. We take the strength of the laser field to be roughly constant over the length scale of the atom, such that it can be treated as a plane wave, and we apply the rotating wave approximation, leaving out terms which oscillate at $\omega_{eg} + \omega_L$. We also define the Lamb-Dicke parameter:

$$\eta = k \sqrt{\frac{\hbar}{2m\omega_t}} \quad (2.22)$$

so that we can rewrite $\vec{k} \cdot \vec{r}$ as $\eta(a + a^\dagger)$. Here we assume that the propagation direction of $\vec{k}$ points along the oscillation axis. If there were instead an angle θ between $\vec{k}$ and $\vec{r}$, then we would reduce η by a factor of $\cos \theta$.

The resulting interaction Hamiltonian is:

$$H_{int} = \frac{\hbar}{2} \Omega \left(\sigma^+ e^{i\eta(a+a^\dagger)} e^{-i\omega_L t} + \sigma^- e^{-i\eta(a+a^\dagger)} e^{i\omega_L t} \right) \quad (2.23)$$

This Hamiltonian couples states of the form $\langle S, n' |$ to states of the form $|D, n\rangle$, where we have now expanded the state vector to contain the motional quantum number n . The term in the H_{int} with η acts only on the motional part of the state, so we introduce a

motional-transition-specific coupling strength as [WMI+98]:

$$\Omega_{n',n} = \Omega |\langle n' | e^{i\eta(a+a^\dagger)} | n \rangle| \quad (2.24)$$

$$\Omega_{n',n} = \Omega e^{-\eta^2/2} \eta^{|n'-n|} \sqrt{\frac{n_{<}!}{n_{>}!}} L_{n_{<}}^{|n-n'|}(\eta^2) \quad (2.25)$$

where $n_{<}$ is the lower of the two motional states n and n' and $n_{>}$ is the greater of the two states. The function L_n^α is the generalized Laguerre polynomial.

Finally, we have:

$$H_{int} = \frac{\hbar}{2} \sum_{n',n=0} \Omega_{n',n} \sigma^+ e^{-i\omega_L t} |n'\rangle \langle n| + \Omega_{n,n'} \sigma^- e^{i\omega_L t} |n\rangle \langle n'| \quad (2.26)$$

and in the frame rotating along with H_0 , which we will call the "quantum computing" frame, we have [Lab08]:

$$H_{QC} = \frac{\hbar}{2} \sum_{n',n=0} \Omega_{n',n} \sigma^+ e^{-i\delta t} |n'\rangle \langle n| + \Omega_{n,n'} \sigma^- e^{i\delta t} |n\rangle \langle n'| \quad (2.27)$$

where $\delta \equiv \omega_L - \omega_{eg} - (m - n)\omega_\nu$.

2.5.2 Selection Rules

From the evaluation of equation 2.21, we obtain the allowed transitions between the Zeeman sublevels of the excited and ground states. For a quadrupole transition, the quantum number for the projection of angular momentum m_j can change by $|\Delta m| = 0, 1, 2$. The particular jumps that are allowed depend on the direction of the applied magnetic field $\vec{B}$, the laser polarization $\vec{\epsilon}$ and the laser propagation direction $\vec{k}$. We define the angles ϕ , the angle between $\vec{k}$ and $\vec{B}$, and θ , the angle between $\vec{\epsilon}$ and $\vec{B}$.

Evaluating the matrix element precisely requires Wigner 3- j symbols [Jam98], but the geometry dependent part can be captured in a factor which depends only on ϕ and θ

[Roo00, SSK12]:

$$g^{(\Delta m=0)} = \frac{1}{2} |\cos(\theta) \sin(2\phi)| \quad (2.28)$$

$$g^{(\Delta m=\pm 1)} = \frac{1}{\sqrt{6}} |\cos(\theta) \cos(2\phi) + i \sin(\theta) \cos(\phi)| \quad (2.29)$$

$$g^{(\Delta m=\pm 2)} = \frac{1}{\sqrt{6}} \left| \frac{1}{2} \cos(\theta) \sin(2\phi) + i \sin(\theta) \sin(\phi) \right| \quad (2.30)$$

Some useful points to note from these terms is that for $\phi = \theta = 90^\circ$ (all three vectors mutually orthogonal), the only non-zero terms are $g^{(\Delta m=\pm 2)}$. When $\phi = 0$ and $\theta = 90^\circ$, the only non-zero terms are $g^{(\Delta m=\pm 1)}$. As long as the propagation direction is orthogonal to the magnetic field, then $\sin(2\phi) = 0$ and the $\Delta m = 0$ transitions are not driven. Due to our chosen trap and magnetic field geometry, this is usually the case in our experiments.

2.5.3 Sideband Transitions

The formulae given above can be simplified in the cases where the average motional occupation is low. In the so-called Lamb-Dicke regime, defined by the condition $\eta^2(2n+1) \ll 1$, many terms in the expansion in equation 2.25 can be neglected.

In our experiments, we typically work with atom-laser interactions in which only one motional quantum is added or subtracted. To lowest order, we have the relevant Rabi coupling rates:

$$\Omega_{n,n} = \Omega(1 - \eta^2 n) \quad (2.31)$$

$$\Omega_{n-1,n} = \eta \Omega \sqrt{n} \quad (2.32)$$

$$\Omega_{n+1,n} = \eta \Omega \sqrt{n+1} \quad (2.33)$$

There are some interesting things to note about these frequencies. The frequencies $\Omega_{n-1,n}$ and $\Omega_{n+1,n}$ are called the frequencies of the red and blue motional sidebands, respectively; we choose this terminology since we usually observe these transitions as sidebands on the carrier, separated by ω_t when performing Rabi spectroscopy on the ion. For large values of n , the frequencies of the first motional sidebands approach the same value, and for very small values of n , we see that $\Omega_{n-1,n} \rightarrow 0$. Both sideband frequencies are suppressed from the carrier frequency $\Omega_{n,n}$ by the Lamb-Dicke parameter η which is typically about 0.05 - 0.1 in our experiments.

2.5.4 Optical Pumping

When performing quantum gates with our chosen qubit, we must select two particular Zeeman sublevels to use as the $|0\rangle$ and $|1\rangle$ states. For this, we traditionally choose the stretched states $|1\rangle = |S_{1/2}, m_j = -1/2\rangle$ and $|0\rangle = |D_{5/2}, m_j = -5/2\rangle$ ⁶. During Doppler cooling, the electron will eventually end up in a mixed state of $|S_{1/2}, m_j = -1/2\rangle$ and $|S_{1/2}, m_j = +1/2\rangle$. In order to prepare the electronic state in the $|1\rangle$ state, we can use a technique called optical pumping. To do this, we apply a pulse with our quadrupole laser tuned to the carrier frequency of the $|S_{1/2}, m_j = +1/2\rangle \leftrightarrow |D_{5/2}, m_j = -3/2\rangle$ transition. The pulse time is chosen to be equal to half the Rabi period⁷ (i.e. a π pulse) in order to maximally move population into $|D_{5/2}, m_j = -3/2\rangle$. From here, we apply another laser⁸, resonant with the $D_{5/2} \leftrightarrow P_{3/2}$ transition. Once up in the $P_{3/2}$ state, there are dipole-allowed transitions back down to $D_{3/2}$, $D_{5/2}$ or $S_{1/2}$. With the Doppler repumping laser also on to clear out the $D_{3/2}$ level, after a brief period of time, the electron can be brought back down into the $S_{1/2}$ level. Transitions to both of the sublevels of $S_{1/2}$ are allowed, so it is possible to end up back in the unwanted state. But after enough successive rounds of this procedure, the $|1\rangle$ state can be prepared with $> 99\%$ probability.

2.5.5 Sideband Cooling

In addition to preparing the electronic state, the quadrupole laser can also be used to prepare the motional state. The concept is similar: we first apply a laser pulse resonant with the $|1, n\rangle \leftrightarrow |0, n-1\rangle$ transition, where the pulse time is adjusted to maximally transfer population into the excited state. The quench laser is then used to clear out the state back down to the $S_{1/2}$ level (see Figure 2-3). In the Lamb-Dicke regime, most of these transitions do not change the motional state, aside from the ones we intentionally drive. Thus, in each cycle of this procedure, a motional quanta is usually subtracted. This process is known as resolved sideband cooling.

Over time, the unwanted state $|S_{1/2}, m_j = +1/2\rangle$ builds up, so we must occasionally interrupt sideband cooling to re-initialize with optical pumping. During each cycle of side-

⁶Note that driving this transition requires a particular choice of laser polarization and magnetic field direction; the laser propagation, polarization vector and magnetic field axis must all be mutually orthogonal.

⁷The Rabi frequency for this "optical pumping" transition is different from the Rabi frequency of our qubit, and must be tuned separately.

⁸We call this laser the quench laser; so-called since it is used to quench the population stored in the metastable state.




Figure 2-3: (a) The concept of sideband cooling. Upper levels are the metastable excited levels and lower levels are the ground state. The motional part of the state is shown along the horizontal axis. Solid arrows represent driven sideband transitions and squiggle lines represent spontaneous decay. [Roo00] (b) The rate of sideband cooling can be enhanced by speeding up the decay via a transition with dipole-allowed spontaneous decay. These $|P\rangle$ states are shown at the top of the diagram.

band cooling, the Rabi frequency on the red sideband decreases as $\bar{n}$ decreases and thus the optimal π time increases. In our experiment, we can programmatically step up the pulse time in each pass of a loop. We start with a π time that is suited for an ion cooled to the Doppler limit, and then we increase to about $5\times$ this value at the end.

2.5.6 State Detection

In order to detect the state of our qubit, we use a method known as electron shelving [NYD89]. With this method, we use fluorescence from the Doppler cooling transition, re-emitted by the ion, to detect when the ion is in the $|1\rangle$ or *bright* state. When the ion is in the $|0\rangle$ or *dark* state, the ion will not respond to the Doppler light, and thus it will not emit radiation at this wavelength. We are able to collect photons scattered from the ion using a high-numerical aperture imaging system (see section 3.5). Since the ion is a very small and monochromatic source of light, we can make use of wavelength-selective filters and a high degree of spatial filtering to isolate light from the ion and eliminate scatter from the Doppler laser or other sources. At the end of an experiment, we can turn on the Doppler laser and then count photons using a single-photon-sensitive detector. If the number of photon counts observed is above a threshold set by background light levels, then we have determined that we were in the $|1\rangle$ state. Otherwise, we must have been in the $|0\rangle$ state.

This kind of detection can be done with very high fidelity in relatively short amounts of time (short with respect to the coherence time of the qubit) [MSW⁺08, HAB⁺14]. Typically, readout is limited by the amount of light that can be collected from the ion, though this can

potentially be improved with the use of in-vacuum or fully-integrated detectors [TVM⁺21]. State determination can also be accelerated with the use of photon arrival times [BSWL10].

2.5.7 Probing the Motional State

To monitor the performance of sideband cooling, and to measure the effects of heating (see section 4.1 and 4.2), it is useful to determine the average motional occupation, $\bar{n}$. This can be determined quite readily by comparing the probability of transitions on the red sideband to the probability of transitions on the blue sideband. For a given pulse time t , the probability to shelve the electron in the $|0\rangle$ state is:

$$P_{RSB}(t) = \sum_{n=1}^{\infty} P_n \sin^2(\Omega_{n-1,n}t/2) \quad (2.34)$$

$$P_{BSB}(t) = \sum_{n=0}^{\infty} P_n \sin^2(\Omega_{n+1,n}t/2) \quad (2.35)$$

where P_{RSB} (P_{BSB}) is the probability to drive the red (blue) sideband. We assume a thermal distribution for $P_n = \frac{\bar{n}^n}{(\bar{n}+1)^{n+1}}$. We can rewrite P_{RSB} as [TKK⁺00, Lab08]:

$$P_{RSB}(t) = \sum_{n=0}^{\infty} P_{n+1} \sin^2(\Omega_{n,n+1}t/2) \quad (2.36)$$

$$= \frac{\bar{n}}{1+\bar{n}} \sum_{n=0}^{\infty} P_n \sin^2(\Omega_{n,n+1}t/2) \quad (2.37)$$

$$= \frac{\bar{n}}{1+\bar{n}} P_{BSB}(t) \quad (2.38)$$

using the fact that $\Omega_{n,n+1} = \Omega_{n+1,n}$. Solving for $\bar{n}$, we have:

$$\bar{n} = \frac{P_{RSB}(t)/P_{BSB}(t)}{1 - P_{RSB}(t)/P_{BSB}(t)} \quad (2.39)$$

This expression for $\bar{n}$ has a few interesting properties, most notably that it does not depend on the base Rabi frequency Ω and that it is valid for any probe time t . When measuring the motion state, we need only take two points, one on resonance of each of the two motional sidebands for a particular pulse time. To make the measurement a little less prone to experimental error, we typically scan across each sideband with Rabi spectroscopy and then extract the value of $P_{RSB}(t)$ and $P_{BSB}(t)$ on resonance from a fit.




Figure 2-4: (Left) Relevant levels, transitions and linewidths for strontium [DeP21] (Right) Levels for calcium

Table 2.1: Landé g-factors for low-lying internal states of the ion

	$S_{1/2}$	$P_{1/2}$	$P_{3/2}$	$D_{3/2}$	$D_{5/2}$
g_J	2	$2/3$	$4/3$	$4/5$	$6/5$

2.6 The Particulars of Strontium and Calcium

In our experiments, we use both strontium and calcium, particularly the ions $^{40}\text{Ca}^+$ and $^{88}\text{Sr}^+$. Both of these ions have well-characterized quadrupole transitions [SSJ17] which make for good optical qubits. For both species, we choose the most naturally-abundant isotope, which both have zero nuclear spin, and thus no hyperfine splitting. An applied magnetic field is required to split the degeneracy of the different total angular momentum projection sublevels. The energy shift of the various levels is given by $\Delta E = m_j g_J \mu_B B$, where m_j is the quantum number for the projection of the total angular momentum, B is the applied field strength, μ_B is the Bohr magneton, and g_J is the Landé g-factor. These g-factors depend only the angular momentum and spin of the electron, so they are the same for both $^{40}\text{Ca}^+$ and $^{88}\text{Sr}^+$ and are given in Table 2.1 [Lab08]. It is interesting to consider how the different quadrupole transitions discussed in section 2.5.2 are shifted in frequency by the applied field. Of particular interest is the shift of the "qubit" transition $|S_{1/2}, m_j = -1/2\rangle \leftrightarrow |D_{5/2}, m_j = -5/2\rangle$ which is $-2\frac{\mu_B B}{h}$. The "optical pumping" transition $|S_{1/2}, m_j = +1/2\rangle \leftrightarrow |D_{5/2}, m_j = -3/2\rangle$ is shifted by $-\frac{14}{5}\frac{\mu_B B}{h}$. By measuring these two transitions frequencies, and taking a difference, we're able to get a quantity which depends only on B as $\Delta f = \frac{4}{5}\frac{\mu_B B}{h}$, which is a handy way to measure the magnetic field, which then allows us to determine the exact frequency of our laser.

The wavelengths of the various lasers required for each ion are different but similar⁹ (see Figure 2-4). Notably, calcium requires a fairly ultraviolet set of lasers for Doppler cooling (397 nm) and for second-stage photoionization (< 390 nm), which can be a potential issue for laser-induced charging (see section 4.1.3). However, calcium has the advantage that all of its required wavelengths fall well below the bandgap of Si, which makes them easier to image on a CCD when aligning beams to our trap. Another important difference between the two ions is the mass. Lower voltages are required to get similar trap frequencies with calcium, and thus higher motional frequencies can be achieved. Of course, the tradeoff is that calcium is more sensitive to electric field noise heating.

In our lab, we have some systems with calcium and some systems with strontium (and one with both, to take advantage of dual-species operations, like sympathetic cooling and quantum-logic readout). Experiments with ion motion are typically done with calcium systems, and everything else happens in strontium systems. In this thesis, chapters 5 and 8 are done with calcium and chapters 6 and 7 are done with strontium.

⁹Throughout this thesis, we try to use generic terminology when possible (e.g. "Doppler" laser to refer to either 397 nm or 422 nm or "qubit" laser to refer to either 729 nm or 674 nm).

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Experimental Apparatus: Lincoln Standard One

*“He looks around, around.
He sees angels in the architecture.”*

—Paul Simon

The experimental apparatus that we use has been developed over the course of a few years to allow for a high-degree of flexibility and a fast cycle time. This planning has resulted in a system which is capable of performing quite different experiments, with just a few modifications. The quick turn around for installing new traps has been invaluable when debugging or doing rapid development. Here we will describe the elements of our system which are common to the various incarnations throughout the rest of this thesis.

In section 3.1, we discuss the construction and operation of our cryogenic vacuum chamber. The full chain of the voltage sources that we use for controlling the trapping potential is described in section 3.2, both for the RF voltage and for the quasi-static control voltages. In section 3.3, we present the 2D magneto-optical trap which is used as a pre-cooled source of neutral atoms. The various lasers required for controlling the ions are introduced in section 3.4, including the setup for the highly-stabilized laser used to address the optical transition we use as our qubit. The imaging system that we use for collecting, observing and counting photons emitted from our ions is described in section 3.5. In section 3.6, we discuss the geometry of the "standard" trap design that we use in many of the experiments presented here. Lastly, in section 3.7, the timing control system that we use is discussed

briefly.

3.1 Cryogenic Vacuum Chamber

In order to hang on to the ions that we trap for finite periods of time, our ion trap must be operated in an ultra-high vacuum (UHV) environment, since collisions with background gas can impart kinetic energy to trapped ions, potentially exciting them out of the trap if they are not sufficiently recooled. There are a few different methods for achieving UHV; in many cases, the vacuum chamber must be pumped at an elevated temperature (i.e. "baked") for a period of a few days in order to pump away water and other adsorbates on the walls of the chamber, which would otherwise slowly outgas over time, increasing the pressure in the chamber. An alternative pathway to UHV is to instead cool the chamber down, all the way to the boiling point of helium at 4.2 K or lower, at which temperatures most substances in the chamber will freeze out. When going this route, only part of the chamber will be cooled down (the cryostat and anything else attached to the low-temperature segments), and the cold surface of the cryostat will act as a pump to reduce the pressure in the rest of the chamber. One of the main benefits of cryogenic operation is a greatly reduced cycle time; UHV can be achieved with a cryostat on the timescale of hours, rather than days, as would be required with a baked room temperature system. Another important benefit to low temperature operation is a reduction in electric field noise originating from the trap material (discussed further in section 4.2). The price of running a cryogenic system is that the system must be constantly refrigerated, which comes at a high power cost and usually requires special attention to ensure systems run without interruption. Cryocoolers also tend to generate acoustic noise and mechanical vibration during their expansion and compression cycle, which can cause decoherence issues (see section 4.3.1).

In our vacuum system¹, we use a vibration-isolated cryostat, based on a Gifford-McMahon (GM) cryocooler. To isolate the vibration of the cryocooler, the cryostat is composed of an inner part and an outer shell, separated by a flexible rubber bellows (see Figure 3-1 and [PHK⁺18]). The inner part of the assembly is a commercial GM cryocooler from Sumitomo. When operating without a significant thermal load, the lowest part of the inner cryocooler, here called the cold finger, will typically reach temperatures slightly below 4 K. The space

¹Some aspects of the system have been described previously in more detail in [BMCS16, BMS⁺19].




Figure 3-1: Schematic of the vacuum chamber, lifting mechanism, inner radiation shields and trap mounting point. Drawing is not to scale, but is generally topologically correct. Components are discussed further in the text, but there are a few key points to note. The upper part of cryostat, with the valve motor, is either attached to the gantry (and thus the ceiling) during normal operation or is suspended by the chain hoist (also attached to the ceiling). The middle part of the vacuum chamber, including everything up to the break point and all of the interior components, is either attached to the optical table via the posts at the bottom of the diagram or is attached to the upper part of the cryostat by the removable plates. The break point is a 10" conflat seal. With the bolts removed, the entire assembly can be lifted by the chain hoist so that we can access the radiation shields and cold head, in order to exchange traps.

between the inner and outer parts, referred to as the thermal exchange space, is filled with high-purity helium gas. When the temperature of the cold finger decreases below the condensation temperature of helium, the gas will convert to a liquid at the bottom of the exchange space, effectively transferring thermal energy from the cold finger to the bottom of the outer part of the cryostat assembly, here called the cold head. There are high-power cartridge resistors inserted into the cold head which allow us to heat it.

Onto the cold head, we attach an interposer "pedestal" which has various connections for attachments that will be discussed in more detail below. The pedestal also has a thermal shield which fully surrounds the ion trap. This thermal shield is sunk to the cold head, so it will reach approximately the same temperature. Keeping a cold surface around all sides of the trap is useful for effective cryopumping. We find that the addition of this shield increases the trap lifetime to many hours, which we attribute to a reduced background pressure at the ion's location.

Due to the limited thermal capacity of helium, it is difficult to achieve very high cooling power at ~ 4 K on the cold head. In our system, the cooling power is about 1 W. At this level, the thermal load from blackbody radiation from the room temperature walls of our vacuum chamber will prevent the cold head from reaching its base temperature. Thus we use an intermediate radiation shield assembly which completely surrounds the main body of the cryostat. The radiation shield mounts to the top of the cryostat at a point that is able to provide much more cooling power, approximately 50 W. In this configuration, we find that the radiation shield routinely reaches < 60 K. The radiation shield assembly is made of a few large, solid pieces of $1/4$ " thick copper, which is electropolished and coated with nickel so that it will reflect the infrared blackbody radiation of the room-temperature walls of the chamber. There are two main parts of the radiation shield; the bottom part (shown in Figure 3-2) can be easily removed for better access to the interior of the cryostat. There are a few holes in the radiation shield assembly to allow for wiring to get from the outside of the chamber to the cold stage. At the bottom of the lower part of the radiation shield, there are small glass windows to allow lasers to be sent in to our ions and a large window for letting ion fluorescence out of the chamber during state detection (see Figure 3-2).

The outer walls of the vacuum chamber are composed of various custom pieces which are connected together with conflat seals. The outer vacuum chamber is rigidly attached to an optical table via some 1.5 " diameter support posts. Thus the majority of the vacuum




Figure 3-2: (Left) Photo of the innermost part of the chamber. The gold-colored component, the pedestal, attaches to the cold head of the cryostat. At the bottom of the pedestal, the PCB octagon (Figure 3-3), which holds the ion trap, is attached. Various cables and wires can be seen, which route signals from feedthroughs at the top of the chamber down to the trap. Wires are wrapped around the cryostat to sink as much heat as possible away from the trap location. (Right) Photo of the cold head with the lower part of the 50 K radiation shield installed. The shield is coated with electropolished nickel, to reflect blackbody radiation emanating from the room temperature walls of the vacuum chamber. Glass windows can be seen at the bottom of the shield, where laser beams can be routed in towards the trap. The windows appear green-tinted since they have been anti-reflection (AR) coated at our wavelengths of interest, which are typically in the blue/NUV and red/NIR regimes. In this photo, the large 4 K and 50 K imaging windows, which are in place during normal operation, have been removed for better viewing of the trap.

chamber (including the part that the ion trap attaches to, when in the default, lowered position) is well-registered to the optical table, which simplifies laser alignment and focusing (see section 3.4.3). One of the conflat seals in the outer chamber is designed to be easily accessed so that the seal can be broken when venting the chamber and exchanging traps. We have a chain hoist above the system which can be used to raise half of the vacuum chamber after the seal is broken, so that we have easy access to the cold head at the bottom where the trap attaches (Figure 3-1).

3.1.1 Feedthroughs and Wiring

At the top of our vacuum chamber, we have a cross piece with eight 2.75" conflat flanges. The ports are used for a few different feedthroughs. On two of these ports, we have 4.5"-to-2.75" conical reducers which are used to attach dual DB25 4.5" flanges. One of these flanges is used to send in 50 total analog voltages which are used to supply the quasistatic control voltages (section 3.2). Inside the vacuum chamber, the wiring for the control voltages is done with quad-twist phosphor-bronze wire, chosen for its favorable ratio of electrical conductivity to thermal conductivity [Pob07, WM02]. The other dual DB25 feedthrough is used to send in wires for temperature sensors and heaters, which will be discussed in the next section. Sensor wires are quad-twist phosphor-bronze, and heaters are enamel-coated 30AWG copper wire. One of the 2.75" flanges is dedicated to a DB9 feedthrough which goes to an array of four cartridge resistor heaters which can be used to heat the radiation shield assembly. The final wire feedthrough on the upper part of the assembly is a 3-pin cable, from which two of the wires supply the positive and negative lead for the current that drives the magnetic field coils which are used to define the quantization axis.

In the middle of the chamber, there is one final wire feedthrough, which carries just one single-ended signal which is used to supply the RF voltage. Near the bottom of the chamber, there is a feedthrough with a combination Bayard-Alpert/Pirani gauge, which is able to read out the chamber pressure from atmosphere all the way down to 4×10^{-10} Torr. Opposite the gauge, the final feedthrough is used as a pumping port, with an all-metal valve that is used to close off the chamber. It's important to pull a moderate vacuum in the chamber before turning on the cryocooler, since the atmosphere in the chamber would otherwise freeze to the cryocooler and trap assembly. The vacuum must also be held when warming up the cryostat to avoid pressurizing the chamber. When starting and stopping the cryocooler, we attach a mechanical roughing pump to the pumping port and leave the angle valve open. When the cryocooler is in normal operation, the vacuum can be maintained by the cryocooler itself, so this valve can be shut and the mechanical pumps removed.

3.1.2 Thermometry

When operating the cryostat, it's important to have a few temperature sensors to monitor the function of the cryocooler and confirm that things are fully cold before proceeding

with ion operations. At least, a reading of the temperature of the outer radiation shield and of the cold head are the most useful. In some cases, it might also be advantageous to get a more precise measurement of the trap chip temperature, in which case an extra sensor might be added to the system. To sense temperature, we use specially constructed diodes, supplied by Lakeshore Cryotronics, which have temperature-dependent IV curves that have been calibrated by the manufacturer. By supplying a constant current to these devices, the voltage drop across the diode then becomes a proxy for the temperature. The relationship between voltage and temperature is nonlinear; a special temperature controller (Lakeshore 336) uses the calibration data to interpolate the temperature given an arbitrary voltage. The temperature measurement is done in a 4-wire configuration to overcome the relatively high electrical resistance of the phosphor-bronze wires used to route signals to the temperature sensor.

While the diodes are nice sensors which are sensitive over our full range of interest, there is one major drawback to this choice. The presence of an AC voltage signal across the diode, due, for instance, to capacitive coupling from the RF wiring, will be half-rectified by the diode. After subsequent low-pass filtering in the voltage measurement electronics, the rectified AC voltage signal will appear as a DC offset in the measurement. In some portions of the voltage-to-temperature curve, the deduced temperature is quite sensitive to the exact voltage, so even a small offset can lead to a significant temperature error. We have found that the reported temperature can vary by up to a few K when the RF voltage is turned off and on². An alternative temperature sensor is a resistor which is specially fabricated to have a significant temperature-dependent resistance at the temperatures of interest. Lakeshore offers such a resistor under the trade name Cernox. When feasible, we have recently switched to using this type of temperature sensor instead of using diodes.

We also use some high-power resistors as heaters in order to achieve closed-loop control at a few of the stages. The ability to heat the stages is especially useful when we are warming up the chamber, at which time the only other source of heating is blackbody radiation from the room temperature walls of the chamber. With the heater resistors, we can apply between 50 and 100 W of Joule heating to the cryostat. The Lakeshore 336 controller is able to manage the PID for the closed-loop control; it can also slowly ramp

²The apparent change in temperature in these cases is immediate, which rules out a true thermal coupling and points instead to an electrical coupling.

the setpoint which is useful when warming up, to ensure that the stages do not heat up too quickly.

In some cases, it is useful to operate the cryostat slightly above the base temperature. We have observed that there is a non-negligible vibration coupling (measured via qubit decoherence, as described in section 4.3.1) between the cold finger and the cold head when we allow helium to fully condense in the exchange space. If we intentionally heat the cold head to just above the condensation temperature, say 6 K, then the amount of vibration coupling is decreased and the effectiveness of the cryopumping is not impacted noticeably [PHK⁺18].

3.2 Voltage Sources

To define the DC part of the ion's potential (section 2.3), we use precision, low-noise voltage sources. The exact source varies depending on the experiment. Here we describe the system in the most basic configuration, in which voltages are generated by a separate system outside of the chamber and routed in along the lines described above.

3.2.1 DC Sources and Amplifiers

Recently, external voltages have come from a few different sources. Older systems use 32-channel National Instruments DAC systems (PXI-6723) which have 13-bit DACs with 800 kS/s (kilo-sample per second) maximum update rate and ± 10 V full-scale output range³. LabVIEW experiments are written for applying voltages based on precalculated voltage sets. In the LabVIEW framework, the voltages to get a particular trap frequency (the so-called "nominal" voltages) are stored on disk, along with the voltage change required to get a 1 μ m shift along each axis (called "shim" voltages). In the control panel, the nominal voltage set is specified, and the shift along each axis is updated and applied. The Virtual Instrument (VI) will then multiply each shim voltage set by the prescribed shift, sum the results, and then add the shimming voltages to the nominal voltage set. Shim controls can be used for fine-tuning the position of the ion and for compensating residual micromotion. The compensation process will be described in section 4.1.3.

³Note that the 800 kS/s update rate with this device is only possible when updating on channel at a time. The batch update rate for updating all 32 channels is reduced to 45 kS/s.

Newer systems use 32-channel Sinara DAC systems (either the Zotino or the Fastino) which have 16-bit DACs with either 31 kS/sec or 2 MS/sec update maximum simultaneous update rate and ± 10 V full-scale output range. Sinara DACs are controlled with ARTIQ experiments, which will be described in section 3.7. In the ARTIQ framework, voltages are calculated live whenever they are needed. Voltage optimization software runs as a standalone server that can be queried by specifying a trap position, frequency and tilt. Optionally, a shim can be specified as an electric field in units of V/cm at the trap position. Unlike the LabVIEW scheme, in this case the optimization is done with the shim condition in place, so only one voltage set needs to be generated at a time. There is a separate library of functions which is able to map a voltage set onto the hardware and communicate with the DAC modules to update their voltages. The software for generating voltages and translating the voltages for the DAC hardware will be described at length in appendix B.

For some ion operations, like splitting a chain or trapping a single ion at a high axial frequency, it is advantageous to have a larger range of voltages, in order to achieve higher electric fields at the ion's position. For this purpose, we use two different designs of multi-channel $4\times$ voltage amplifiers. An earlier design, which is the main one used, unless stated otherwise, had a heavily-limited bandwidth in order to avoid voltage noise on the output. A newer design, explained extensively in appendix A, has a higher bandwidth and slew rate, to be used in experiments involving ion motion (see chapter 8).

3.2.2 DC Filtering

Ions are sensitive detectors of electric field noise (see section 4.1), so special care must be taken to filter the voltages applied to the trap. Fortunately, a trapped ion is a fairly narrow-band sensor, so there are not too many restrictions on where all in the spectrum the noise must be low. For our typical ion species and trap geometries, we would like to limit the noise at about 800 kHz and above. The target voltage noise in this range is $1\text{ nV}/\sqrt{\text{Hz}}$ or less⁴. In all of the experiments described here, at least some amount of low pass filtering is needed to achieve this level of noise.

In addition to suppressing noise from the voltage sources themselves, it is also useful to filter noise that may come from other environmental sources, such as capacitive coupling

⁴The reason for this value, and the expected amount of heating from this noise, will be explained further in section 4.1.1




Figure 3-3: (Left) Photo of the PCB octagon interposer, with important components labeled. The surface electrode trap is bonded to a CPGA and mounted to a zero-insertion-force (ZIF) socket. The PCB electrically connects the pins of the ZIF socket to cable headers via some low-pass filters, composed of a single resistor and capacitor. Also of note are some higher bandwidth, general use coaxial connections and a coaxial connection for the RF input signal, both of which are routed to dedicated pins of the ZIF/CPGA. A temperature sensor mounted to the PCB reads out the temperature, though, due to the high thermal resistance of the FR4 material, the temperature is not always a useful proxy for the temperature of the trap and the rest of the PCB. (Right) Block diagram of the control voltage generation, filtering and routing. The voltage amplifier is explained in further detail in appendix A. The external filters are described in the text and shown in Figure 3-4 (higher bandwidth version only).

along the long wires going into the cryostat. Since this noise can enter in at any point along the DC lines, filters are most effective when placed as close as possible to the trap electrodes. In the Lincoln system, we use a low pass filter, composed of a single resistor and capacitor on an interposer PCB board (Figure 3-3) to which the trap is attached at the bottom of the cold head. The value of these components is different for different experiments, but the standard is usually $100\text{ k}\Omega$ and 1 nF . The capacitor must be NP0/C0G-type since this material has finite capacitance at cryogenic temperature. Unless stated otherwise, these are the values that will be used, and thus the cutoff of this final low-pass is about 1.6 kHz .

On systems which do not require high-bandwidth manipulation of control voltages, an external multi-pole low-pass filter is used with has several order of magnitude of suppression at ion frequencies of interest. This filter array, provided by Kiwa electronics, has a nominal




Figure 3-4: (Left) Photo of filter motherboard assembly, showing seven daughterboards plugged into a single module. Each of the daughterboards has filters for four DAC channels. The input side to the motherboard is designed to be compatible with the Sinara DAC cards (either with or without the high-voltage amplifier in appendix A). The output of the motherboard is designed to plug into the DB25 feedthrough on the vacuum chamber. (Right) Schematic of the combined filter assembly; the part of the schematic on the daughter card is highlighted in blue. The assembly is designed to have many degrees of flexibility, in cases where the input comes from a floating source or in which an additional RF signal is to be added to the control electrode via a bias tee (though both of these applications are outside the scope of this thesis). Unlabeled components are either not populated or replaced with 0Ω resistors, depending on whether the connection is needed to complete the circuit.

cutoff frequency of 50 kHz. In the typical configuration, voltages originate from the sources, and then are routed via BNC cables through the amplifiers and into the Kiwa filter array. The feedthrough into the chamber for the control voltages is a dual DB25 pin connector header, so we have a custom box which combines 50 independent BNC voltage signals into two DB25 cables, which are then routed into the chamber.

A high-bandwidth alternative to this routing is to use a custom filter array board, instead of the Kiwa filter array (see Figure 3-4). In this arrangement, the HD68 cables which come out of either the Fastino or Zotino DACs are routed (either directly or through the high-bandwidth, high-voltage amplifier) to the filter board, which is composed of several interchangeable daughter cards. The input to the filter mother board is an HD68 connection, and the output is a single DB25 cable. Each of the daughter cards has components for filters on four channels. The filter components are intended to be swapped out to achieve different filter types, as necessary; in this work, the component values were borrowed from a design that was recently reported [KLS⁺20]. In this case, the filter assembly acts as a three-pole low-pass filter with a cutoff frequency around 120 kHz. When these filters are used, seven of the 32 channels coming out of the Sinara DACs are not connected to wires inside the




Figure 3-5: (Left) Block diagram for the RF network. The primary RF source is the origin of the signal applied to trap. A few other components in the network are used for either reducing or probing sensitivity to noise in the RF system; these will be discussed further in section 4.1.2. The orange block represents the chamber; the RF signal is directed into the chamber via a coaxial feedthrough. (Right) The schematic for the resonator inside the vacuum chamber [Gre16]. The 100 pF and 18 pF capacitors and the 27 nH inductor are part of an impedance matching network, in order to match the $50\ \Omega$ source resistance to the very low resistance when the resonator is driven at resonance. The 430 nH inductor forms half of the resonator along with the trap capacitance C_{trap} .

cryostat.

3.2.3 RF Source and In-Vacuum Resonator

For the typical radial confinement strength required in our experiments, we must apply an RF signal of about 70 V peak amplitude at 50 MHz to the RF electrode of the ion trap. This level of signal is difficult to apply directly, for instance from an amplified function generator, due to the high power that would be required and the fairly high impedance of the RF electrode by itself. To boost the voltage applied to the electrode without requiring a large input power, a resonant circuit is used, by combining the capacitance of the electrode with an inductor. In our system, we use a resonator board which has a 430 nH inductor preceded by a network of capacitors which are used to match the $50\ \Omega$ impedance of the input source to the fairly low impedance of the LC circuit when it is driven at resonance (see Figure 3-5). The value of the inductor is chosen to place the resonant frequency at a value which gives a nice compromise between trap depth and stability, for our typical ion species.

The entire resonator board assembly is attached to the bottom of the cold head, in close proximity to the trap. Keeping the resonator close to the trap reduces the amount of wire that is needed to connect the inductor to the trap, and thus reduces the resistance in the LC circuit, which improves the quality of the resonance. The input signal to the resonator board comes from a single-conductor feedthrough on the vacuum chamber. On the air side of this

flange, a small network of BNC cables and connectors, connected in parallel with the BNC input to the flange, is added to increase the capacitance of the matching network slightly and improve the impedance matching into the resonator. This amount of capacitance is fine-tuned while using a network analyzer to make an S11 measurement looking into the resonator from outside of the chamber. The quality factor (Q_L) of the resonance is determined from the 3 dB width of the resonance, and the depth of the resonance feature shows how well the impedance of the system is matched to $50\ \Omega$. When the resonance peak is maximally deep, the reflection of the input signal is minimized and thus the impedance is well-matched. While looking at the S11 measurement, the parallel capacitance is modified, by adding or removing BNC components, until the depth of the resonance is optimal.

When the resonator is well-coupled, the voltage across C_{trap} is given by [Gre16]:

$$V_{trap} = 2\sqrt{Q_L L \omega_0 P_{in}} \quad (3.1)$$

Where Q_L is the loaded Q⁵, L is the value of the main inductor in the resonator, and $\omega_0 = 1/\sqrt{LC_{trap}}$ is the resonant frequency. P_{in} is the RF power applied at the input to the resonator, and V_{trap} is the amplitude of the RF signal on the RF electrode. Typically the value of Q_L in our system is 80–150.

The input to the chamber comes from an RF amplifier. An RF power of about +20 dBm is required to reach the typical voltage amplitude required (between 50 V and 100 V amplitude). For this, a moderate-power, general purpose RF amplifier is sufficient. In our system, we use the Mini-Circuits ZHL-1-2W amplifier. Between the amplifier and the chamber, we add a directional coupler in order to monitor the power reflected back from the chamber. When this reflected power is minimal, then we know that we are driving on resonance. It's useful to have a live view of this resonance, since the trap capacitance will usually be modified slightly as the RF power is increased, due to small thermal effects from Joule heating. The amplifier is driven by a function generator, which is chosen to have a low phase noise and high stability. The SRS SG382 or SG384 are nice choices for this, and are also both easily programmed and monitored remotely. In some cases, we add additional components to the RF network, such as an RF attenuator or bandpass filter, to either reduce phase noise in the system or to probe the ion's susceptibility to this effect. We also occasionally

⁵The loaded Q is the quality factor of the resonance measured when the resonator is connected to a source. This is the value we measure when performing an S11 measurement, so it is convenient to use.




Figure 3-6: (left) Schematic diagram of the laser and magnet configuration in the 2D MOT, described further in [TGLW09, BMCS16]. Stacks of permanent magnets create magnetic field gradients. This creates position-dependent shifts in the $^1S_0 \leftrightarrow ^1P_1$ transitions which can then create position-dependent forces by using polarization selectivity. Along each arm of the MOT, $\sim 1''$ diameter laser beams detuned by 60 MHz from the $^1S_0 \leftrightarrow ^1P_1$ are sent through the chamber. One of the beams is σ^+ polarized and the other is σ^- polarized. Each beam is retroreflected back through the chamber, after passing twice through a quarter-wave plate, which reverses the handedness of the polarization. This configuration of beams creates a cooling force directed radially inward in the plane depicted in the figure. When the MOT beams are on, a cloud of cooled atoms, with kinetic energy around a few milliKelvin, is formed in the plane with the MOT lasers. Attached to the side of the MOT chamber, there is a source of neutral atoms (not shown) which can be heated up to $\sim 400^\circ\text{C}$ to increase the vapor pressure to the point that there are enough atoms in the MOT. (right) Photograph of the interior of the 2D MOT chamber with the lasers on. A bright cloud of calcium atoms can be seen in the center. The same concept can be applied to the strontium MOT.

use a secondary source of RF in order to probe technical noise sensitivities. All of these extra components will be discussed in more detail in section 4.1.2.

3.3 2D MOT

To introduce neutral atoms to the ion trap region, we use a pre-cooled source of atoms which are prepared in an ancillary chamber. This extra chamber, which is baked and kept under constant vacuum via an ion pump which is always running, contains an oven which is filled with metallic chunks of either strontium or calcium. The oven, itself essentially just a tube with a nozzle on the end, is wrapped with a heater made out of a meander of tungsten

wire. By running current through the wire, the oven can be heated to around 350° C for strontium or around 420° C for calcium in order to increase the vapor pressure of the metal in the chamber. The vapor in the chamber can then be cooled in two directions using a two dimensional magneto-optical trap [TGLW09]. It is possible to set up the requisite field for this configuration using only permanent magnets. Four stacks of magnets are arranged in a circle around the outside of the chamber, as in Figure 3-6. This setup creates a magnetic field gradient which decreases radially inward. In the standard MOT arrangement, circularly polarized beams red-detuned from the neutral $^1S_0 \leftrightarrow ^1P_1$ transition are pointed in from the sides of the chamber, through viewports which are anti-reflection coated at the required wavelength. The helicity of the polarization can be flipped by retro-reflecting the light after it passes through a quarter-wave plate. These beams are all red-detuned from the zero-field transition. Away from the center of the chamber, the magnetic field causes one of the three possible $^1S_0 \leftrightarrow ^1P_1$ transitions to be brought into resonance with the detuned laser via Zeeman splitting. The atoms will absorb photons from one of the beams, depending on which polarization is allowed, and the photon will impart some momentum to the atoms in the interaction. By careful choice of magnetic field direction and laser polarization, a trap can be setup in which the atoms are always pushed back towards the center of the trap when they move off the middle⁶. In steady state, with the MOT beams left on, the atoms will form a cylindrical cloud at the center of the chamber with a reduced average kinetic energy in the radial direction (see Figure 3-6, right).

Due to the large size of the MOT beams, we require a significant amount of power, around 15 mW, in each of the beams to effectively cool the neutral atoms⁷. To obtain lots of power close to the chamber, we use an injection locking setup [SBTL15] in which a high power (~60 mW) laser diode close to the $^1S_0 \leftrightarrow ^1P_1$ is seeded by a primary laser. For the injected diodes, we use high-power blue diodes from Nichia, which have been frequency-selected to naturally lase close to the target wavelength. The injection can be seeded with as little as 100 μ W of input power, though we may sometimes operate with a few hundred μ W for better long-term stability. We need most of our laser power to be red-detuned for the MOT, but we also need a fraction tuned to resonance for photoionization. To

⁶In practice, there are just four possible configurations for the circularly polarized beams. It is often easiest to just test all of the configurations, by adjusting wave plates, until the MOT fluorescence is observed.

⁷Also, with less than 5 mW of power in the MOT beams, the MOT becomes difficult to see with our chosen oven temperatures, and the status of the laser becomes difficult to ascertain.

do this, we lock our primary laser to a reference cavity (see section 3.4.1) such that it is already detuned. The output of the injection laser can thus be sent directly into the MOT chamber. We sample part of this light and use a low-frequency AOM to shift the frequency onto resonance, which becomes the source of the photoionization beam. Since the frequency of the PI beam is locked to the frequency of the MOT beams, visibility of the pre-cooled cloud is thus a useful check to confirm that the laser is both injection locked and that the frequency of the primary laser is correct.

To get neutral atoms from the 2D MOT chamber into the chamber with the ion trap, we use another red-detuned "push" beam derived from the same laser that supplies the MOT beams [SKC12]. The push beam is aligned down the axis of the pre-cooled cloud, pointing towards the trap. While the MOT beams are normally left on, the push beam is only turned on when ions are being loaded into the system. When an ion absorbs a photon from the push beam, the photon imparts some momentum in the direction of the trap. Unlike the MOT beam interaction, this push operation is not sensitive to the polarization. In the middle of the MOT cloud, there is not a clearly-defined quantization axis, so any polarization will have some chance of being absorbed by an atom in the cloud. Whenever the push beam is on, a stream of pre-cooled neutral atoms is flowing towards the trap. Ionization beams, aligned to the trap site, can then be switched on in order to convert some fraction of the incoming neutral atoms into ions, which will then be confined in the trap zone by the electric fields coming from the trap electrodes [BMCS16].

In some ion trap experiments, the source of neutral atoms is only heated when loading is happening. In our system, the oven in the MOT chamber is left on throughout the day. The advantage of this scheme is that re-loading ions can be quite quick, which is a valuable feature, especially when testing out motion operations that might be prone to kicking ions out of the trap or when operating with traps that have noise sources that might tend to heat ions enough that they have some chance to escape the confining potential. One disadvantage to leaving the oven on is that the chamber with the MOT has an elevated pressure, sometimes in the 10^{-8} Torr range. To keep the higher-pressure chamber isolated from the experiment chamber, we use a tube with a narrow hole down the middle, called a differential pumping tube. The constriction limits the allowed solid angle for atoms to pass through without colliding with and sticking to the walls of the chamber. When operating, we find the lifetime of ions in the trap—a relevant proxy for background pressure in the

experiment chamber—is not affected by the oven. Since the chamber with MOT is too far away from the cryostat to be effectively cryopumped, care must be taken to maintain the vacuum in this chamber through other means. This is why an ion pump is kept on the MOT chamber. To prevent the vacuum from being disrupted when the experiment chamber is not cold, a gate valve is used to connect the MOT chamber to the experiment chamber. The gate valve is only left open when experiments are running.

There are a few other advantages to the 2D MOT loading scheme that are worth mentioning. In some oven loading schemes, a neutral source of atoms is placed above the ion trap, and when the source of neutrals is heated, atoms flow towards the trap and end up deposited on the surface of the ion trap. In some large 3D traps, this small coating of metal is acceptable, but in planar traps, even a thin layer of metal can lead to shorts between electrodes⁸ With the 2D MOT, the flux of neutral atoms incident on the surface of the trap is greatly reduced since the pointing of the push beam allows a great deal of directional control. The neutrals above the trap also have a higher chance of having a kinetic energy below the depth of the trap, since they have already been cooled in two directions. Since the cooling in the MOT and the pushing operations are both based on resonant atomic transitions, the 2D MOT also provides isotope selectivity, since the mass shift of undesired isotopes will cause the atomic transitions to be shifted off of resonance with the laser fields. In our systems, we essentially never observe events where the wrong isotope is loaded into the trap, which is not always the case in other ion loading schemes.

3.4 Laser Sources and Alignment

Six wavelengths of laser light are required to ionize, cool and then manipulate the quantum state of an optical qubit in an alkaline earth ion. Four of these wavelengths address optical dipole transitions, which are relatively broad features, with linewidths around 1 MHz. One wavelength pumps into the continuum⁹ so there is no stringent linewidth requirement on this laser. The final laser address the optical qubit transition; the linewidth of this laser can often be the limitation to the fidelity with which the transition can be driven. It will be

⁸Some more advanced planar traps have holes in the trap so that the source of neutrals can be placed behind the trap, to limit this unintentional deposition. It is also possible to limit this effect with the use of baffles between the oven and the trap.

⁹For strontium, there is a closed transition to an auto-ionizing electronic level, but this transition is also very broad, at the level of 100's of GHz.

treated separately from the others, and discussed further in section 3.4.2 and section 4.3.1.

The lasers that pump into the continuum are a 405 nm laser for strontium and a 375 nm laser for calcium. These lasers do not require any additional frequency stabilization beyond the feedback from the reflecting facets of a Fabry-Perót laser diode. For these two wavelengths, we use Toptica iBeam smart systems, which contain embedded temperature controllers, current sources, and collimation optics. The wavelengths of these lasers can be adjusted via temperature tuning to get close the relevant atomic transition. The lasers are shared between multiple systems, so some beam distribution, consisting of wave plates and polarizing beam splitters, is required to send laser light to the different experiments. Ideally, given our typical beam sizes, each system requires about 1 mW of power at these wavelengths to be sent into the chamber. Fortunately, there are blue and near-UV diodes available which can supply many 10's of mW at the needed wavelengths.

The lasers which address dipole transitions all require some additional frequency stabilization techniques. In most cases, the linewidth available with a commercial external-cavity diode laser (ECDL), which is around 100 kHz or less, is sufficient. For some of the more infrared lasers, distributed feedback (DFB) lasers are available which do not require an external grating to adjust the frequency. Since DFB laser diodes are monolithic devices, they are less prone to drift, and we find that their frequency is more constant than those of the ECDLs over long time spans. For dipole transitions, and given our typical beam sizes at the ion, we usually want powers in the range of $1 - 100 \mu\text{W}$ in these lasers. In addition to being distributed between multiple experiments, these lasers also require frequency and amplitude modulation, which comes at the cost of laser power. Given all sources of power loss and splitting, we end up needing between 20 and 40 mW of power at the laser source.

All of our ECDL and DFB lasers are in the Toptica DL100 or DL100 Pro form factor, and come with controllers that provide temperature controllers, current sources, and high-voltage sources, in the case of piezo-tuned ECDLs. The diodes are frequency-selected by Toptica to be close to the relevant atomic transitions. To ensure good long-term stability, the temperature controllers for the lasers are left on all the time, but to preserve the lifetime of the diodes, the current is only switched on when experiments are running. Most lasers are kept in a separate room from the experiments, to isolate the lasers from the acoustic noise of the cryocoolers; we use $\sim 40\text{ m}$ of optical fiber to route each beam from the laser room into the experiment room.

3.4.1 Atomic References

To monitor the frequency of each of the lasers, we use a commercial wavelength meter (Angstrom WS6 and WS7). By tuning the temperature and current in the lasers, we can bring the frequency close to the atomic transition frequency. Imperfect temperature and current control and other sources of mechanical drift cause the frequency of the the lasers to change over time. For most of the lasers, a frequency drift of more than a few MHz is unacceptable, as this can lead to less-effective cooling or initialization of the electronic state. One way to stabilize the frequency of each laser is to lock the frequency to an atomic line, for instance using a vapor cell of a neutral reference atom. For the 422 nm Doppler cooling laser of strontium, there is a well-known atomic feature in rubidium [MMB98], which is offset from the required frequency by only about 400 MHz which can easily be spanned by an acousto-optic modulator. We use a saturated absorption spectroscopy technique to see the Doppler-free feature [Sag06] by dithering the probe laser and demodulating the transmission signal. This kind of atomic reference lock is ideal in that it is free from any calibration error in frequency measurement electronics. The rubidium line can be uniquely identified and thus the frequency of the 422 nm laser can be known exactly.

Not every wavelength has a useful atomic line nearby, and even if they did, the overhead of having vapor cells and spectroscopic locks for each laser would be cumbersome. The atomic stability of one locked laser can be translated to other lasers with the use of a Fabry-Perót (FP) reference cavity. Using a PDH locking scheme, the length of an FP cavity can be locked to the wavelength of a laser which is separately locked to an atomic line. The mirrors of the FP cavity are coated to have high-reflectivity at the atomic reference wavelength and at any other wavelengths that also need to be stabilized. The length of the cavity is chosen such that the free spectral range is small enough to be spanned by an acousto-optic modulator. Once the length of the cavity has been locked, then any other laser that is locked to the cavity (using a similar PDH setup) will benefit from the same stability as the reference laser.

This transfer cavity scheme is useful for achieving long-term stability of the various wavelengths, but none of the transfer-locked laser frequencies are known precisely. The laser frequencies must be initially determined using other methods, i.e. by looking for ions and tuning up the fluorescence signal. Once the frequencies have been optimized, a pickoff

of the laser can be brought into resonance with one of the fringes of the transfer cavity using an acousto-optic modulator. Each day that the laser and transfer lock are used, the fringe of the transfer cavity can be tracked by monitoring the voltage on the piezo required to bring the length of the cavity into resonance. This ensures that the cavity is always locked at the same exact length. The other lasers that are locked to the transfer cavity are monitored with the wavelength meter. As long as the resolution of the wavelength meter is higher than the free spectral range of the transfer cavity, then the reading can be used to identify the correct fringe.

The 422 nm laser for the strontium Doppler transition doubles as an atomic reference for a transfer cavity which also locks the 397 nm and 423 nm lasers for calcium. For the infrared wavelengths, we use an ancillary 852 nm laser which is locked to a feature in a cesium vapor cell. The infrared transfer cavity is used to lock a 922 nm laser (which is frequency-doubled to 461 nm for strontium), a 1033 nm laser for strontium, and an 854 nm laser for calcium.

Two additional infrared lasers, the 866 nm laser for calcium and the 1092 nm laser for strontium, are not locked to a reference cavity. Instead, these two lasers are intentionally frequency-broadened to 1 GHz, which limits the effect of frequency drift. There are other advantages to broadening, as was discussed in section 2.4.

3.4.2 ULE Cavity System

The lasers which address the optical qubit, the 674 nm laser for strontium and the 729 nm laser for calcium, have special requirements. Since the atomic transition in this case is a quadrupole transition, it is challenging to saturate, and thus we are virtually always limited by the amount of laser power that can be delivered. At a minimum we would like to have at least 1 mW of power at the chamber. Given the losses in the frequency shifters and power distribution between several experiments, we end up requiring more than 300 mW of laser power at the source; to achieve this, we use a cascade of injection-locked lasers and tapered amplifiers (see Figure 3-7).

For both qubit beams, the primary laser is an ECDL, which will have a stabilized linewidth of around 100 kHz. Since the atomic transition is on the order of 1 Hz wide, additional stabilization is required. The main trick for reducing the laser linewidth is to lock the frequency of the laser to an high-finesse optical cavity composed of highly-reflective mirrors optically-contacted to an ultra-low expansion (ULE) glass spacer [YCIB99, JLL⁺11].




Figure 3-7: Block diagram of the lasers and locks that compose the narrow-linewidth qubit laser generation. The setup here closely follows a system presented in [ANK⁺15]. A primary laser is locked to a ULE cavity, and the transmitted light is used to seed an injection lock. Boxed arrows represent optical isolators which can be used to prevent etaloning from unintentional feedback and also used to direct control signals into injected lasers. An AOM between injection lock stages can be used to account for a frequency difference between the atomic resonance and the nearest fringe of the ULE cavity. Following the injected lasers, tapered amplifiers (TAs) can be used to get the power high enough to overcome modulation and control losses on the way to the experimental chamber. We use fiber noise cancellation (FNC) to mitigate vibrational noise picked up in the optical fiber between our laser room and the room with the experiments [MJYH19]. The diagram above is representative of the distribution for the 674 nm laser for strontium. For calcium, only one injection stage and one TA is required. A cavity fringe is close enough to the calcium optical transition that a shifting AOM is not necessary.

The ULE cavity is isolated from the laboratory environment in a vacuum chamber, providing acoustic and thermal isolation, which is further bolstered by two temperature-controlled radiation shields in the vacuum chamber. The temperature of the cavity is chosen to be at a point where the length of the glass in the spacer is first-order insensitive to temperature changes. To minimize vibration coupling, the spacer has a conical shape and is held up by supporting at just three points around the middle. The entire vacuum chamber is further isolated by a vibration isolation platform. The mirrors of the cavity are specially coated for high-reflectance at the wavelength of the qubit laser. Both of our ULE cavities (for 674 nm and 729 nm) have a free spectral range of approximately 2 GHz and a finesse of $\sim 10^5$, resulting in a cavity linewidth of around 10 kHz. These cavities also exhibit a linear

drift, of a few Hz per minute, which requires tracking via atomic spectroscopy using the locked laser.

When we lock the primary laser to the ULE cavity, the ultra-high mechanical stability of the spacer is converted to a stable frequency. With high-speed feedback on a PDH error signal, the laser linewidth can be controlled to a level much smaller than the cavity linewidth. Without two ULE-stabilized lasers at the same wavelength, it is difficult to precisely determine the exact linewidth of the final system. From some calibration measurements using a different narrow-linewidth laser, we can bound the linewidth to at most 20 Hz. The noise of the laser is only controlled up to the bandwidth of the servo electronics. Close to the bandwidth, the phase of the control signal begins to shift, and the servo feedback causes a bump to appear in the spectrum. Given the bandwidth of our electronics, this servo bump can appear between 100 kHz and 1 MHz. Since this is close to typical trap frequencies, the frequency content in the servo bumps can lead to situations in which we unintentionally drive motional sidebands while trying to drive the carrier, or vice versa. The result of this off-resonant driving is noise in spectroscopy, due to a finite probability for the electronic state to change, even at frequencies where no atomic transition should occur.

A natural way to bandpass filter the frequency of the laser is to use the cavity itself. In the beam that is transmitted by the cavity, the servo bumps are attenuated since they are not resonant. This transmitted beam is not immediately useful since it has quite low power. To avoid having too much intracavity power, we send only about $700\text{ }\mu\text{W}$ into the cavity. Ideally, the transmitted power should be equal to the incident power in steady state, but we find that the measured power is much less than this, presumably due to loss in the mirrors and in the viewports on the vacuum chamber. We get about $10\text{ }\mu\text{W}$ out of the transmitted beam. In order to step this power up, we use an injection lock. This is composed of a second temperature-controlled diode, which is AR coated on the output facet, such that there is no cavity-enforced lasing mode of the diode by itself. The temperature of the diode is adjusted such that the center of the laser gain profile is brought near the frequency of the primary laser. Then, by using the side port of an optical isolator, the transmitted beam of the ULE cavity can be directed into the injection diode. With careful mode-matching between these two lasers, the mode of the injection diode will synchronize to the mode of the primary laser. This scheme effectively boosts the power of the transmitted signal from tens of μW up to tens of mW. For the 674 nm system, we find that a second injection stage

helps to achieve long-term operation. The first injection is operated at a lower current, where the injection is easier to maintain, and then the second injection can be seeded with much more input power.

Even this level of power is not enough for distributing to all of our systems, so we follow the injection laser system by a tapered amplifier. The 674 nm system, which is distributed between more systems, consists of two tapered amplifiers, with some additional pick offs between the amplifiers. At the end of the full system, we have a few hundred mW of highly-stable laser power.

To transfer the light from our laser room over to the room with our experiments, we use ~ 40 m optical fiber. The fiber runs through the ceiling of our lab and is prone to temperature fluctuation and acoustic interference, both of which influence the effective optical path length. To mitigate these effects, we use fiber noise cancellation (FNC), which is a scheme based on comparing a laser signal that has passed through the length of the fiber with a local reference signal. The beat note between these signals is broadened by the effect of perturbations in the fiber, so by using a frequency modulator to stabilize the beat note to a stable RF oscillator, the frequency fluctuations in the fiber can be suppressed. The setup requires the use of a single-pass AOM, which transmits about 60% of the input optical power at most. We also typically lose about 50% of the power coupling into the fiber, so the total power that reaches the experiment table is about 30% of the power from the laser room.

3.4.3 Frequency Modulation and Focusing

For a few of the lasers, namely the detection laser (either 422 nm or 397 nm) and the qubit laser (either the 674 nm or the 729 nm), rapid frequency control of the laser frequency is required at the experiment, for performing spectroscopy, tuning the lasers to optimal cooling settings, or for qubit rotations. For the quench lasers (either the 1033 nm or the 854 nm), we require the ability to pulse the amplitude of the laser, in order to clear out the metastable levels and reset the electronic state on demand. Both frequency and amplitude control can be obtained using AOMs in double-pass configurations, which can shift the frequency of laser without affecting the beam pointing [DHL⁺05]. In most configurations, laser light reaches the experiment's optical table via an optical fiber. This light is then routed in free space through a double-pass, through some optional attenuators, and then through a final

half-wave plate for adjusting the polarization; after this, light is then coupled back into optical fiber to simplify the final routing into the chamber.

The AOMs in the double-pass configurations are driven by adjustable direct digital synthesizers (DDSes). The DDS frequency is written on to the laser frequency, so these need to be at least as precise as the laser frequencies themselves. For most lasers, this is not a stringent constraint, but at least for the qubit laser, care must be taken to ensure that no additional frequency noise is introduced by the RF source. In our experiments, the timing and pulsing of the DDSes is controlled by a programmable real-time experiment control system, which will be discussed in section 3.7.

Optical access into the chamber is provided by seven 2.75" viewports around the bottom of the chamber. Each of the viewports is anti-reflection coated over the range of wavelengths we care about. The windows are all single pieces of BK7 glass which are fixed onto the side of the chamber using indium seals, which form a vacuum seal. The advantage to using indium seals is that this avoids stress in the glass which is common in preformed viewports and can lead to strain-induced birefringence. Inside the vacuum chamber, there are additional windows on the radiation shields which have the same dielectric coating. The windows on the shields in the chamber are attached with indium but do not form leak-tight seals. They are intended to screen the infrared blackbody radiation from the walls of the chamber and also to block the path of higher-temperature gases which originate further away from the cryostat. The distance from the center of the trap to the outside surface of the viewport is about 6". The smallest beam waist that can be achieved on the surface of the trap is limited by clipping at the edge of the ion trap and the PCB interposer, which prevents us from making the beam arbitrarily large outside of the chamber. Typically, we use lenses with focal lengths between 175 mm and 200 mm with 1 mm or 2 mm input beam diameter in order to obtain a beam diameter of about $50 - 100 \mu\text{m}$ at the ion's position. These final focusing lenses are placed on compact 3-axis positioning stages with about 1 cm of travel in each direction. For small adjustments, translating the lens translates the focused beam by the same amount, and thus with micrometer stages, the beams can be positioned at the ion's location to within less than $10 \mu\text{m}$. Initial alignment is done with the help of a camera focused on the trap, which will be discussed in the next section.

Recently, the micrometers on a few of the stages have been exchanged for servo motors¹⁰

¹⁰We use Thorlabs Z625B, driven by a custom H-bridge circuit and controlled by an FPGA running custom




Figure 3-8: Diagram of the imaging system. Distances are not to scale and only intended to show the order of the components. The main element of the imager is the primary objective, which is a large, compound lens designed to give diffraction-limited performance with a high numerical aperture (0.43). The lens is designed to produce an image plane with $10\times$ magnification, indicated in the diagram as the "intermediate focal plane." Placing an iris at this plane allows for spatial filtering of the image that is produced, which we use to pass only light coming from near the ion(s). Depending on the placement and focal length of the relay lens, additional magnification can be achieved. The EMCCD and the PMT are both at equal distances along the optical axis, and the path is chosen by the state of the flipper mirror.

which can be used to programmatically control the position of the beams on the trap. This control is useful for ion transport experiments where the trap zone might be different between experiments. Being able to move the position of the Doppler cooling beams is also useful when ions are loaded at a different position from where experiments occur.

3.5 Imaging System

In order to detect ions in the trap, and later to determine their electronic state, we need to collect photons scattered by the ion when it decays along its detection transition. The two primary challenges in imaging the ion are maximizing the amount of light that is collected and engineering a high resolving power so that ions in a chain can be distinguished. These two requirements can be satisfied with an imaging system with a high numerical aperture,

firmware.

which means that the first lens must either have a large diameter or be close to the ion, since the numerical aperture is related to the ratio of the diameter to the focal length. It is possible to place the first lens of the imaging system inside the vacuum chamber, but this can present many technical challenges. Tilt in an in-vacuum lens is difficult to adjust without large and potentially-vacuum-incompatible stages, and the aberration from this tilt may not be possible (or at least overly cumbersome) to correct with external optical elements. In our system (shown in Figure 3-8), the first optic in the imaging system is placed outside of the vacuum, but the chamber is engineered to get the ion as close as possible to the bottom of the chamber, with large windows at the bottom to maximize the clear aperture. The distance between the ion and the bottom of the lowermost window is approximately 62 mm.

The main lens in the imaging system is a compound objective which is designed to achieve diffraction-limited focusing given the available working distance. The lens is mounted right below the chamber, as close as possible to the ion. The numerical aperture of the lens is about 0.43. Newer systems have a lens designed to have diffraction limited focusing at both 397 nm and 422 nm, the detection wavelengths for both calcium and strontium, with minimal focal length shift between the two wavelengths. Both styles of objective have similar numerical aperture. These lens systems are designed for 10 $\times$ magnification; on the opposite end of the lens from the ion, a image of the ion will be formed at a distance of around 400 mm from the back of the lens, at an intermediate focal plane. The main lens is placed vertically directly underneath the chamber; to keep us from having to build the rest of our imaging system one meter below the table, we insert a 45 $^\circ$ "folding" mirror to move the intermediate focal plane to a position further back on the optical table, where there is room to set up the detectors.

A second lens, called a relay lens, can be used to obtain additional magnification. One can use the thin lens equation to determine the additional magnification, given the focal length of the relay lens and the distance to the intermediate focal plane. It is useful to take advantage of the intermediate focal plane to apply some spatial filtering. An iris at the intermediate focal plane can block out light that is scattered from the trap surface that enters the imaging system via other pathways. We use a so-called "zero aperture" iris which is able to close completely, and with this, we are able to close down the aperture such that only scatter from directly beneath the ion will affect detection. Since we detect ion

fluorescence at a particular wavelength, it is also helpful to include narrow bandpass filter to completely block scatter at irrelevant wavelength. This filter is included just before the iris in the optical path, simply for convenience; it does not matter where in the imaging system the filter is inserted. A quick-change filter carriage allows the filter to be removed easily, which is helpful to focus the imaging system at other wavelengths, when other lasers are being aligned to the trap surface.

Ideally, the entire imaging system should be translated together, to preserve the distance between the main objective and the relay lens and to keep all elements centered on the optical axis, but in practice it is difficult to fix everything to the same stage. Instead, we allow translation only for the main objective. For small adjustments, moving the main objective will move the image by a corresponding amount (depending on the system's magnification). Lateral translations of the objective allow us to focus on different parts of the trap, which is necessary for trapping at different zones on the trap. Often, it is also beneficial to translate to the edges of the chip when performing initial beam alignment, since the extra scatter at the edges of the trap makes for a prominent alignment marker. Vertical translations of the objective allow the focus to be adjusted, which is required when focusing on different wavelengths during alignment. Small vertical adjustments are also required when traps are exchanged, due to imperfect repeatability, and also when the temperature of the trap is changed by large amounts, due to thermal expansion and contraction. If the front surface of the objective is not completely normal to the optical axis, the image formed by the system will be affected by coma, which causes the point source of the ion's fluorescence to appear spread out into an oblong shape. To overcome the various problems that tilt can be unintentionally introduced into the system, for instance by imperfect machining tolerances, we mount the lens on a four axis tip-tilt and translation stage. The tilt of the lens is adjusted to minimize coma once ions have been imaged. After this initial adjustment, we find that additional tuning is rarely required.

To detect the image of the trap during alignment, or to watch ions during loading or other calibrations, we use an electron-multiplied charge-coupled device (EMCCD) camera. Each pixel of the camera can be sensitive to light at the level of single photons, which is important to overcome the limited collection efficiency of the entire system. The camera is moved into position such that the magnified image of the trap is in focus at the surface of the EMCCD; fine adjustments of the position and focus of the image can then be performed

using the translation stages of the lens, as stated previously. The camera we use is sensitive to all of our wavelengths of interest¹¹, so it is a valuable tool for the initial alignment to the trap.

The camera is useful for initial tune up, but it is difficult to use it to read out the state of the ion quickly¹². Experiment time scales are usually limited by the amount of time required to detect the electronic state of the ion, so additional detection time increases the total duty cycle of experiments. It is possible to read out from the camera quickly by pulling counts from only a subset of the pixels in the CCD, but feeding the detected counts into the experiment controller also requires a direct interface to the camera and a programmable interface for adjusting the area of the CCD that is read out. To obtain a more easily interpreted measure of the ion’s fluorescence, we use a photo-multiplier tube (PMT) which is also sensitive to light at the single-photon level. The PMT generates a small voltage pulse when a photon is detected, which can then be converted to a TTL pulse with analog detection electronics. The TTL signal can then be counted by our experiment timing controller. To select between the camera and the PMT as a detector, we use a motorized flipper mirror which can direct the focused image of the ion onto either instrument. In the typical tune up procedure, we start with the image on the camera, and then use this while loading to watch for ions to be trapped. After trapping, we can optimize the pointing of the various beams to maximize the fluorescence signal from the ion. We then close the iris to the point that only the ion is visible, while using the translation stages on the lens to center the intermediate image of the ion on the hole in the iris. Finally, we then flip the image over to the PMT, after which point we can monitor the status of the ion in terms of photon counts.

3.6 Surface Electrode Trap Summary

In section 2.2.1, we introduced a miniature style of ion trap which is defined in a plane. These surface traps are compatible with lithographic fabrication, in which layers of metal

¹¹For the 1033 nm and 1092 nm laser, the camera sensitivity is at the edge of its range, so imaging these wavelength requires increased laser power and is usually only possible if there is a feature which scatters a lot of light, like a dust particle on the surface.

¹²Difficult, but by no means impossible. With a fast readout interface from the EMCCD camera and adaptive readout techniques, state detection can be performed with a camera in a similar time to what we require with our PMT. Indeed, this camera-based readout is potentially useful in ion arrays, when the states of multiple ions, spread over some area of the trap may need to be read out simultaneously. See [Bur10] for more discussion about camera-based readout.

and dielectrics are deposited with precise thicknesses and then patterned into shapes with μm accuracy using photoresist and material-selective etching. Traps fabricated in this way can be combined with other features embedded into the substrate¹³, but there are a few major features which will be common among different traps. It will be helpful to review these components of the design and introduce some terminology that will be useful in future discussion.

The trap axis, the line on which chains of ions form and along which ions may be transported, is defined by the RF electrode. In linear segments of the trap, in which there are no bends or other features, the RF electrode consists of two strips on either side of the trap axis. It is important to apply the same voltage to both sides of this electrode, so the electrode is usually fabricated in a contiguous plane which forms a shape like a tuning fork. There will not be a minimum in the ponderomotive potential when the two ends of the RF electrode join together, so this area is kept far away from regions where ions will be trapped. In practice, when the joint is about $500\ \mu\text{m}$ or further, it will have a negligible effect on the potential seen by the ion. The geometry of the RF electrode, in particular, the width of the two strips, and the separation between them, determine the shape of the pseudopotential in the radial plane (i.e. the planes normal to the trap axis) [Hou08]. For a given ion height, there is an optimal choice of RF electrode width in order to maximize the trap depth, which is defined as the highest kinetic energy the ion can have while maintaining a bounded trajectory [Wes08, NH12]¹⁴.

The electrodes that are used to define the DC part of the potential are here called the "control" electrodes. In the simplest possible configuration, in which we simply want to define a quadratic potential at a fixed position, we might only need three control electrodes, one in the center with a relatively negative voltage, and two on either side, called "endcap" electrodes, with relatively positive voltages. The next constraint we require is to be able to rotate the principal axes of the ion's motion in the trap, in order to ensure that the Doppler cooling beam, which comes in parallel to the plane of the trap, will have a nonzero projection on each of the three modes. This rotation, or tilt, is achieved by manipulating the DC components of the radial potential, and requires having control electrodes on either side

¹³See, for instance, this thesis.

¹⁴Choosing the ion height may not be a simple task, since this is constrained by the geometry of other parts of the apparatus, like achievable beam diameters. Some considerations, like minimizing applied voltages, or increasing trap frequencies, motivate us to bring ions closer to the surface, but an unfavorable scaling of noise from the trap surface (see section 4.2) tends to tip the scale in the other direction.




Figure 3-9: Drawing of the Lincoln standard trap. The electrode pattern is mirrored about the vertical and horizontal axis. The RF electrodes (thicker horizontal segments), define the pseudopotential, and set most of the trapping parameters, including the trap height of $50\text{ }\mu\text{m}$. The 18 outer segmented electrodes, including the four larger endcap electrodes are used to define the quasistatic or "control" part of the potential. Two additional control electrodes, along the middle of the trap, provide some additional degrees of freedom with a large (and opposite) projection along the $\hat{x}$ direction, which can be useful for tilt generation. Ions are typically trapped in the middle of the narrower segmented electrodes, at intervals along the y axis given by multiples of the electrode pitch, $120\text{ }\mu\text{m}$; these positions are sometimes referred to as "trap zones."

of the trap axis¹⁵. Finally, in order to achieve even finer control of the trap potential (e.g. beyond manipulation of the quadratic terms only) and to be able to generate minima at other locations along the trap axis, we wish to divide the control electrodes into "segments" with a fixed length along the axial direction. Given the distance between the segmented electrodes and the ion, there is an optimal choice for the pitch, which maximizes the quadratic field that is generated by one control electrode while minimizing the effect on the field generated by adjacent electrodes [RLB⁺06, NH12].

3.6.1 "Lincoln Standard" Geometry

Most of the traps used in this thesis all have a common set of dimensions, so for convenience, those will be presented up front. Unless stated otherwise, the relevant dimensions are given

¹⁵We may sometimes say that electrodes which participate strongly in the axis rotation are "tilt" electrodes, but in general, most of the control electrodes will play some part in the potential generation.

Table 3.1: Dimensions of the Lincoln Standard trap

Segmented electrode width	116 μm
RF electrode width	122 μm
Space between RF electrodes	38 μm
Center rail width	13 μm
Gaps between electrodes	4 μm (sometimes 6 μm)
Pseudopotential minimum height	50 μm

in Table 3.1 and the geometry is shown roughly in Figure 3-9.

The standard trap geometry mostly follows along with the optimized choices stated above, given the chosen ion height of 50 μm . One quirk to call out is the split center "rail" electrodes which run parallel to the RF electrodes. Having these electrode free is not strictly required to generate our typical potentials, but the extra control is helpful to keep voltages within range, especially when generating tilt.

Unless stated otherwise, the top metal layers in the traps discussed in this thesis are all niobium. Though niobium is superconducting at the temperatures at which we usually operate, we do not find a significant difference in ion behavior using this material over other materials (see Figure 4-1). We simply stick with niobium traps as they have been well-characterized in the various experiments in our lab.

3.6.2 Voltage Simulation

We use a home built software package which optimizes the voltages that are applied to trap electrodes given some requirements on the desired potential. A detailed discussion of the optimizer itself (ITVG) is covered in appendix B. The input to the optimizer comes from electrostatic simulations performed using a finite element differential equation solver (COMSOL). With COMSOL, we are able to take into account many odd features of trap geometry, including effects from electrode thickness, exposed dielectric layers and other nearby conducting surfaces. However, solving for the potential with COMSOL directly is cumbersome, and can take up to tens of minutes, depending on the level of discretization. Instead of re-solving each time, we do a series of simulations in which we apply 1 V to each of the trap electrodes, and in each case, we solve for the potential in a range of interest around the ion and save the result to files (called "grid" files). By the linearity

of the Laplace equation, we are able to sum together the grid files, after multiplying each by the voltage applied to the corresponding electrode, to obtain the potential for a given voltage configuration. This makes finding a potential much faster, as it is reduced to simple (however large) matrix multiplication and addition.

ITVG allows us to introduce several constraints, including the position and trap frequency of an ion. Some electrodes can be fixed during the optimization, which is useful in cases where there are lots of electrodes and few required conditions. Indeed, in most cases in which there is just a single ion or a few-ion chain, the voltage optimization problem is underconstrained. During the optimization, the RF electrode is treated separately: we use the same electrostatic simulation as with the DC electrodes, but to this field we apply the pseudopotential approximation in order to obtain the time-averaged effective potential. This RF pseudopotential is taken into account when determining how to adjust voltages in order to realize the desired tilt.

3.7 The Brain: ARTIQ

Experiment control typically consists of changing voltages, changing frequencies applied to modulators, pulsing the amplitude on those modulators, and then counting photons to determine the result of an experiment. For some operations, it is sufficient for these parameters to be set once, but many of the tasks must happen with a fixed timing relationship. For instance, lasers used for detection must be switched on while photons are counted and pulses with the qubit laser must be done with accurate phase accounting, to avoid decoherence. In our experiments, we use a timing control system (ARTIQ) [BWJ⁺18] which is commercially developed by M-Labs in close collaboration with NIST and Oxford University. The system consists of various pieces of hardware, which are connected to an FPGA controller running a specialized firmware. The controller is able to execute programs, which might include loops, decision branching and data storage. A dedicated timing pipeline, in combination with the deterministic behavior of the FPGA, ensures that the sequence of actions all happens as specified by the user. The programs that are executed by the FPGA are interpreted from Python code, in an architecture that allows precisely-timed code to interact seamlessly with higher-level scripts which are run on the host computer. For a full discussion of the capabilities of ARTIQ, we defer to the online documentation and tutorials.

Chapter 4

Technical Challenges

*“I know it’s hard but
You’ve gotta deal with it.”*

—Sleigh Bells

In the course of analyzing the interaction of our ions with the technology we integrate, we will be on the lookout for the ways in which performance is limited by these additions. Isolated charged atoms are highly sensitive detectors in many ways, which means that they are also great sensors of noise! We now consider the various ways in which our experiments can be sensitive to non-ideal operation, how we have learned to mitigate those effects, and how we are able to use these mechanisms in order to assess the performance of our devices.

In section 4.1, we discuss technical noise, which arises due to imperfections in our control systems, and discuss a few sources and their mitigation in detail. In section 4.2, we introduce and then discuss briefly the concept of anomalous heating, which is due to noise of an unknown origin and attributed to some properties of the trap surface. Finally, in section 4.3, we discuss some sources of decoherence of prepared qubit states, in particular laser frequency noise and magnetic field noise.

4.1 Technical Electric Field Noise

In contrast to the dipole and quadrupole interactions discussed in the previous chapters, the ion, being a charged particle, can also have monopole interactions with electric fields, which can be quite strong, even for relatively small fields. Fortunately, when confined in

a trap, the ion has only narrow bands of resonance frequencies, near the normal modes of its motion, where it is most sensitive to noise in the electric field. We consider a spectral density of electric field fluctuations:

$$S_E(\omega) = 2 \int_{-\infty}^{\infty} d\tau e^{i\omega\tau} \langle \epsilon(t) \epsilon(t + \tau) \rangle \quad (4.1)$$

Which is measured in units of $(\text{V/cm})^2 \text{Hz}^{-1}$. The function $\epsilon(t)$ is a fluctuating, uniform electric field. The rate of change of the average thermal occupation of the harmonic oscillator is [WMI+98, TKK+00]:

$$\dot{\bar{n}}_i = \frac{q^2}{4m\hbar\omega_i} \left(S_E(\omega_i) + \frac{\omega_i^2}{2\Omega_{rf}^2} S_E(\Omega_{rf} \pm \omega_i) \right) \quad (4.2)$$

Where ω_i is the frequency in each of the modes of motion; $\dot{\bar{n}}_i$ is usually called the heating rate. The term on the right of equation 4.2 comes from cross coupling between the RF field at Ω_{rf} and the noise field; it will only contribute to heating of the modes which see a significant projection of the RF field (i.e. the radial modes).

It may not be immediately obvious how heating of the ion's external state within the trap can lead to noise which is relevant for quantum behavior of the ion's internal (electronic) state, since these two states are not normally coupled. On the extreme end, excessive motional heating can overwhelm our ability to laser cool ions below the trap depth, and we will see ion loss. The heat also shows up when we turn on the laser-atom interaction, as discussed in section 2.5. When $\bar{n}$ is nonzero, and thermally distributed, the probability to drive carrier transitions (see section 2.5.1) turns into a sum over an ensemble of frequencies [Roo00]:

$$\rho(t) = \sum_{n=0}^{\infty} p_n \sin^2(\Omega_{n,n}t) \quad (4.3)$$

for each n , weighted by:

$$p_n = \frac{1}{\bar{n} + 1} \left(\frac{\bar{n}}{\bar{n} + 1} \right)^n \quad (4.4)$$

Thus, when $\bar{n}$ is large, the fidelity of single-qubit rotations will be reduced¹. Multi-

¹Since the carrier Rabi frequency can be determined from either first principles or independent fitting, it is possible to fit equation 4.3 to a decohering Rabi oscillation with $\bar{n}$ as the only free parameter. This technique is a useful way to measure the motional state when the value is too high to determine with the sideband technique discussed in section 2.5.7 [SVB+14].

qubit gates, which use the shared motion of chains of ions in order to exchange quantum information between them [HRB08], can be performed when $\bar{n}$ is nonzero [KBZ⁺09], but if the value of $\bar{n}$ changes during the gate operation, due to heating, the fidelity can still be reduced [SM00].

Fortunately, many noise processes exhibit an inverse dependence with frequency. Since mode frequencies are typically in the range of ~ 1 MHz, the list of major contributors is limited to a select few sources, which will be discussed in this in the next section. We first consider noise which is unintentionally applied by the apparatus.

4.1.1 Voltage Noise on DC Sources

Voltage noise on trap electrodes is a direct pathway for electric field noise to enter into our system. Given a voltage spectral density of $S_V(\omega)$, in units of V^2/Hz , the electric field noise spectral density is:

$$S_{E_j}(\omega_i) = \frac{S_{V_j}(\omega_i)}{D_{i,j}^2}, \quad (4.5)$$

for each of j trap electrodes, where ω_i is the frequency of the mode i and $D_{i,j}$ is a geometric factor (an "effective distance") determined by the distance from the ion to the trap and the shape of the control electrode [LYS07, SSL⁺18]. The value of $D_{i,j}$ can be obtained from the same electrostatic simulation which is used when calculating trap voltages (see section 3.6.2). In order to determine the contribution from this mechanism to the total heating rate, the electric field contributions from each electrode must be added together, in quadrature², and then plugged into equation 4.2 to find the heating rate of the mode of interest.

The exact level of voltage noise which leads to noticeable heating depends on the trap geometry, and on the contribution from other sources, which will be discussed later. In our experiments, as a rule of thumb, we generally want to keep voltage noise at frequencies of interest at or below $1 \text{ nV}/\sqrt{\text{Hz}}$ ³, and a reasonable effort is made to stay an order of magnitude below this value when possible.

There are several different mechanisms by which voltage noise can end up on trap

²Noise terms are added in quadrature if we assume that the noise on the different electrodes is uncorrelated, as would be the case with voltage noise from different, independent sources on each electrode in the trap. If the noise on electrodes is actually correlated, then those electrodes must be grouped together to calculate a total $D_{i,j}$ for the group which can then be used with equation 4.5.

³With $1 \text{ nV}/\sqrt{\text{Hz}}$ of voltage noise on each electrode in the standard trap geometry (see section 3.6.1), a calcium ion would experience ~ 10 quanta/s of heating in an axial mode at 1.3 MHz. This heating rate is close to the level of anomalous heating (see section 4.2) that we observe, so by keeping voltage noise below this point, we can be fairly confident that we are not limited by this mechanism.

electrodes. Even in the absence of control voltage sources, there can still be voltage noise on trap electrodes due to electromagnetic interference (EMI) or Johnson noise [BKRB15]. The voltage noise spectral density from a resistance R at temperature T is:

$$S_V^{(Joh)}(\omega) = \sqrt{4k_B T R} \quad (4.6)$$

where the units are $V/\sqrt{\text{Hz}}$. Trap electrodes themselves typically have very small resistances, which will be insignificant for Johnson noise, but sometimes, the addition of higher-resistance wires (like our phosphor-bronze wires) or resistors in filter configurations, can bring the contribution of equation 4.6 up to more relevant levels. The level of EMI is much harder to predict, since it can depend on conditions in the laboratory environment. The vacuum chamber, when linked to the ground of the control electronics, makes a good Faraday cage to prevent pickup inside the chamber⁴. Outside of the chamber, signals bound for feedthroughs are sent in along shielded cables whenever possible. Perhaps the most pernicious kind of interference is due to noise on the lab ground, which can show up when currents from different instruments flow to ground along the same non-zero-impedance path. To find these ground loops, we make sure that all instruments are connected back to the same power circuit (a so-called "star" configuration); devices which draw a lot of current, like pumps and compressors, are relegated to a different circuit.

Typically, the control voltage sources themselves are the primary contributor to heating via this pathway. It is for this reason that we take such care when choosing low-pass filter components as described in section 3.2.2. Usually the requirements of increasing the voltage range or bandwidth push the noise up, and so in most cases, the level of performance of our voltage sources will be set by the amount of heating that we can accept.

4.1.2 Phase Noise on the RF Source

Voltage noise on the RF electrode can also contribute to the heating rate via the same coupling discussed in the previous subsection. Unlike the DC electrodes, it is not possible to simply low-pass the RF signal, since this will also attenuate the RF signal itself. However, there is some filtering naturally from the RF resonator discussed in section 3.2.3. For an

⁴Not counting, of course, pickup from the RF input signal. In practice, this will show up to some degree on all wires in the chamber, but at least for this noise coupling mechanism, the extra bit of RF voltage is not a problem for the ion since Ω_{rf} is much greater than any of the motional mode frequencies ω_i .

RLC resonator with resonant frequency Ω , the voltage noise at a frequency $\Omega \pm \omega$ is:

$$S_{V_{RF}}(\Omega \pm \omega) = \frac{QL\Omega}{1 + 4Q^2 \frac{\omega^2}{\Omega^2}} S_{P_{RF}}(\Omega \pm \omega) \quad (4.7)$$

where Q is the quality factor of the resonator and L is the value of the inductance. Here $S_{P_{RF}}(\Omega \pm \omega)$ is the phase noise on the RF source (with units of W/Hz), when it is set to Ω and measured at $\Omega \pm \omega$. We typically use RF sources with high frequency and amplitude stability, and these usually also have low phase noise. In combination with the filtering from the resonator, direct coupling of voltage noise (via the mechanism of equations 4.2 and 4.5) is often not an issue for this electrode.

There is another mechanism by which ions can have enhanced sensitivity to noise in the RF, which is especially prevalent in instances with asymmetry in the geometry of the RF electrode, as is the case with junctions [BOV⁺09, BOV⁺11]. When the ion sits at a point in the potential where there is a gradient in the RF pseudopotential, it is possible for the ion to experience a force resonant with the secular frequency due to noise at $\Omega_{rf} \pm \omega_i$ beating with the carrier tone at Ω_{rf} . The heating due to this mechanism is given by:

$$\dot{n}_i = \frac{q^4}{16m^3\hbar\omega_i\Omega_{rf}^4} \left[\frac{\partial}{\partial i} E_0^2 \right]^2 \frac{S_{V_{RF}}(\Omega_{rf} \pm \omega_i)}{V_0^2} \quad (4.8)$$

where V_0 is the amplitude of the signal applied to the RF electrode. The term $\frac{\partial}{\partial i} E_0^2$ is proportional to the gradient of the pseudopotential, taken along the axis i . In this work, we will typically focus on the axial mode in which $i = y$. The term can come from an electric field projection along the trap axis due to the RF electrode, which can come variety of sources, including features in the trap design (see for instance [SGS⁺18]) or due to fabrication imperfections. It may not be possible to reduce this term to zero, but in practice it can be mitigated.

One technique to mitigate heating due to this effect is to intentionally apply a signal on the RF source at $\Omega_{rf} \pm \omega_y$, which can be done with a power combiner [SSL⁺18]. In figure 3-5, a secondary RF source is shown which can be controlled independently of the main RF drive for this purpose and added together in the total applied signal. The applied noise is increased to the point that the heating from the gradient begins to compete with Doppler cooling, and then the average motional occupation in the mode of interest will

increase. When the noise is high enough, the high motional state will result in oscillation that is visible on the EMCCD camera. Since the oscillation frequency is much higher than the frame rate of the camera, the ion appears to spread out along the mode of interest; this degree of spreading can be used as a metric for how much the ion has been heated (and thus for how strong the pseudopotential gradient is at this location). Adjustments can then be made to the applied voltages, in order to shift the ion's position around slightly. At each position, the degree of spread is re-measured, and we move around, adjusting and minimizing excitation, one axis at a time, until a local minimum has been found⁵. To increase sensitivity, the power of the noise tone can be progressively increased; in most cases, we stop our optimization when the injected power (measured before the resonator) is within 10 dB of the carrier power.

This procedure minimizes the gradient term in equation 4.8, but the noise spectral density $S_{V_{RF}}$ can also be decreased. In our most commonly-used RF sources, the SRS SG382 and the SG384, we have found that there is white noise (independent of the carrier amplitude), which contributes to the heating due to the gradient mechanism more than the phase noise on the carrier signal. Thus, if we simply add a 10 dB attenuator to the output of the function generator (shown in figure 3-5) and increase the carrier tone amplitude to compensate, we can reduce the white noise contribution. Another way to reduce noise on either side of the carrier tone on the RF is to use a bandpass filter, centered on the RF drive frequency. Since this frequency changes for different traps, it is useful to have a tunable bandpass filter, but it is difficult to make a bandpass filter with a narrow (i.e. < 500 kHz) passband which is also tunable over our range of interest. As a compromise, we use two tunable bandpass filters (K&L 5BT-30/76-5-N/N) which each have a 2.3 MHz bandwidth and then adjust the passbands to put the RF carrier frequency on the edge of either passband, such that a signal passing through both of the filters will end up effectively bandpassed at the RF frequency. With this method, we are able to attenuate noise at $\Omega_{rf} \pm \omega_y$ by 20 dB or more. Finally, the issue of excess phase noise can also be avoided by using a more stable RF source; we found that the performance of the HP8640B signal generator was good enough that it did not require the same level of phase noise mitigation

⁵As stated previously, in this work, we focus primarily on the axial mode in which $i = y$, so we typically only go to the trouble of minimizing the gradient in this mode. By the symmetry of our standard trap configuration, usually only radial shifts make a difference, and we find that in most cases a suitable minimum can be found by exclusively adjusting the z position.

as our standard SRS sources.

4.1.3 Stray DC Electric Fields

All of the mechanisms discussed so far have led directly to motional heating. Some of these noise sources are worse depending on the ion's position. For the pseudopotential gradient heating discussed in the last section in particular, the degree of technical noise heating can depend strongly on the ion's exact position. When the ion is held at a position away from the RF null line, it will also experience excess micromotion, which can introduce additional sidebands in spectroscopy and reduce the efficiency of laser cooling [BMB⁺98]. The geometry of the RF electrode (and the corresponding shape of the pseudopotential) sets the position of a minimum in the RF potential, and we do our best to make sure that the DC potential arising from the applied voltages is co-aligned with the RF minimum. Several factors can hinder our ability to place the ion exactly on the null line, including errors in our electrostatic simulations or fabrication imperfections⁶. Unknown electric fields at the ions position can also effect the potential seen by the ion.

Background electric fields can come from charge deposited on the surface of the trap. Many exposed surfaces on the trap have some conductive path to mollify this added charge, but in the surface electrode traps studied here, there are gaps between electrodes with exposed dielectric material. In some cases, there may also be thin layers of non-conductive contaminants on top of metal electrodes which can also accumulate charge (leftover photoresist from trap fabrication, for instance). Finally, more distant surfaces, like the outer edge of the chip carrier or nearby windows, can also lead to stray electric fields. The arrangement of charges does not necessarily have any particular shape or overall sign, since they can be the result of many sources, so it is not usually possible to predict the degree or direction of the ion's shift. One commonly-attributed source of charged particles is from photo-liberated electrons which are freed from the metal electrodes by the ultraviolet lasers used during photoionization [HBHB10, WLL⁺11].

In order to get ions into the exact potential configuration that we intend, we must apply so-called "shim voltages" which are designed to change only one aspect of the potential (e.g. to only add an electric field in the z direction) while keeping all other terms constant.

⁶Both of these sources can be modeled as effective electric field shifts on top of the "ideal," intended field configuration. We can compensate for these effects in the same way that we do for external fields.

Adjusting the linear component of the electric field can be done while minimizing the pseudopotential gradient⁷ (as described above), minimizing the effect of parametric excitation [ITU11, TMA⁺12], or while minimizing some other hallmark of micromotion (e.g. a micromotion sideband in spectroscopy). Higher order terms in the potential will require other techniques for determining the difference between the measured and expected configuration. One other example is the second derivative along the axial direction, which determines the axial frequency of oscillation. This can be adjusted while monitoring the trap frequency via sidebands on the optical transition.

4.2 Anomalous Heating

So far, we have discussed noise and heating which originate from known sources. In these cases, noise can be independently measured (without an ion), and then the expected heating can be calculated using the formulae given above. However, in all experiments with trapped ions, a level of noise is observed which exceeds these expectations and which cannot be immediately attributed to any known sources. The so-called "anomalous" heating that remains has been the subject of investigation by many groups over the last two decades [TKK⁺00, BKRB15]. It is sometimes frustratingly-difficult to draw conclusions from comparisons between totally different experiments, and there is a high degree of variability in the results reported in the recent literature⁸. Even so, some progress has been made on discovering techniques for decreasing the effect of anomalous heating, and on placing constraints on the underlying mechanisms that might be to blame.

Perhaps one of the most clear features of all of the anomalous heating data collected so far is a strong inverse dependence on distance between the trap electrodes and the ion [LHC11, BKW18, SGS⁺18, AMUH19]. There are many benefits to decreasing this distance, including a decrease in the size of the overall trap geometry (which is fixed mostly by the chosen ion height) and a reduction in the required voltages to achieve stable trapping, but the high cost of increased anomalous heating has kept surface electrode trap distances at 30 μm or more [ISK⁺21].

In lieu of a full understanding of the underlying functional dependence of this anomalous

⁷Minimizing the gradient of the pseudopotential is not necessarily the same as aligning the ion to the RF null, though in our usual trap geometries, we find that the outcome is roughly similar.

⁸In all likelihood, there are probably many different mechanisms all at work simultaneously, contributing at about the same level. Different experiments see the effects coming in at slightly different levels.

heating, we use the following placeholder:

$$\dot{n}(\omega, T, d) = \frac{q^2}{4m\hbar\omega} S_{EAH}(\omega, T, d) \quad (4.9)$$

where S_{EAH} is the electric field spectral density attributed to anomalous heating with an unknown frequency, temperature and distance dependence.

4.2.1 Mitigation Techniques

In many experiments, the anomalous heating is observed to decrease with decreasing temperature [LGL⁺08, BSC15]. (Indeed, this is one of the main reasons for using a cryogenic apparatus.) In our experiments, we observe that the heating follows the temperature as a power law, with an exponent of about 1.5, down to some phenomenological "turn on" temperature T_P , according to:

$$\dot{n}(T) = \dot{n}_0 \left[1 + \left(\frac{T}{T_P} \right)^\beta \right] \quad (4.10)$$

Thus by operating our chips at cryogenic temperatures, we are able to achieve levels of anomalous-heating-limited performance over two orders of magnitude lower than would be expected at room temperature. We have performed extensive investigations of temperature dependence on our traps using one of our experiments which is designed to be able to vary the temperature of the trap continuously from 295 K–4 K. We find that, for a variety of different materials, we measure very similar temperature dependence [CS14] (see Figure 4-1).

There are several methods to prepare the surface of a trap in order to remove surface contaminants, and a few of these have been studied in detail in ion trap applications. One technique is to use an ion mill, in which high-energy ions⁹ are used to bombard the surface of the trap. Ion milling has been shown to reduce the heating rate by up to two orders of magnitude [HCW⁺12, DGB⁺14], and some experiments have been carried out in carefully-controlled conditions such that traps are not allowed to be recontaminated after milling. Ion milling has multiple microscopic effects, including removing contaminants, removing trap material itself, and possibly altering the crystalline structure at the surface of the trap. Milling experiments have typically been paired with other surface analysis tools, though to

⁹Not to be confused with the low-energy ions we store in our traps.




Figure 4-1: Measured heating rate versus temperature for a few different materials, all exhibiting roughly similar behavior. All data points are heating rates taken with strontium atoms with a trap frequency near 1.3 MHz. The trap temperature is varied using a special experiment that is able to locally heat the trap above the temperature of the cold head, described further in [CS14]. All traps are fabricated with different materials as called out in the legend but have the standard geometry described in section 3.6.1, with a $50\text{ }\mu\text{m}$ ion-electrode separation. For experiments presented here, these heating rates will be taken as a lower limit when evaluating the performance of trap that have other materials or mechanisms that increase the amount of heating observed.

date it is not well understood precisely what aspect of milling most affects ion heating.

In our experiments, it is not feasible to install an ion mill onto our chamber, due to the difficulty of getting the milling beam to pass through the radiation shields, so instead we use an ancillary chamber in order to mill traps before installing them into the apparatus. Care is taken to minimize the amount of time during which the trap is exposed to laboratory air. With this arrangement, we find that we are able to realize reductions in heating rates which are similar to those observed in systems that have the ion mill attached directly to

the main experiment chamber [SSS⁺18].

While ion milling has been the most studied cleaning technique for ion trap preparation, there have also been promising results using pulsed laser cleaning [AGH⁺11] and *in situ* plasma cleaning [MBCS15]. These two methods, while encouraging, require additional components to be added to the system, either a coil for generating the plasma or a high power laser, so they were not implemented in the traps studied in this thesis.

4.2.2 Potential Sources

A few different microscopic models have been proposed to explain anomalous heating [SNRWS11, SNKW⁺13, HCW⁺13, KHR⁺16, KSNH⁺17, NBUM⁺19]. These models all predict certain behavior of electric field noise as a function of frequency, temperature, and distance, and they yield an overall order of magnitude expectation for the amount of heating [BKRB15]. In experiments, we can vary these parameters and measure the resulting heating rate in order to compare with theory [HMK⁺17, SSS⁺18, BUML⁺21].

One example of this was a measurement of the heating rate as a function of distance, which has recently been done in a few experiments that allow the ion height to be adjusted without exchanging the trap [BKW18, SGS⁺18]. A model for the distance scaling was proposed based on the combined effect of uncorrelated patches, each producing electric field noise, which predicts a distance scaling which goes as d^{-4} when the size of the patches is much smaller than the ion's height [LHC11]. Two of the recent experiments agree well with this model, but the third found a scaling which was consistent with a smaller exponent, and thus a much larger average patch size [AMUH19].

Temperature dependence of anomalous heating has been measured at two extremes, both down to cryogenic temperature [LGL⁺08, BSC15] and up to a few hundred degrees Celsius [NBUM⁺19, BUML⁺21]. The high-temperature behavior has not been studied as extensively, but in one experiment, it has been found that the increase in heating with increasing temperature levels out at temperatures above 400 K and then eventually increases rapidly above 600 K. Neither of these temperature regimes is fully consistent with a microscopic model.

Frequency dependence is more difficult to measure than the other parameters, since the range of available trap frequencies is usually set by the trap geometry to be less than an order of magnitude. Given the already significant spread in heating measurements, it can

be difficult to extract an exponent for a fit to data as a function of frequency. Even so, a $1/f$ electric field noise behavior is often observed and reported [BKRB15], corresponding to an ion heating scaling of $1/f^2$, by equation 4.9.

As stated previously, over a wide range of materials that have been tested, the heating rate dependences have all been fairly similar. One outstanding challenge is to determine if there is some material or treatment technique which can result in a more well-behaved single-component crystal structure which can break this paradigm [BCSH21]. There has been some recent progress in uncovering a difference in material behavior after ion milling [SSS+18], but this work is still in early stages.

4.3 Qubit Decoherence

The natural linewidth of the optical qubit transition that we use is ~ 0.4 Hz (~ 0.15 Hz) for strontium (calcium), governed by the lifetime of the excited state, $T_1 = 390$ ms ($T_1 = 1200$ ms). This linewidth sets a maximum time over which the phase of a superposition $|\psi\rangle = \alpha|0\rangle + \beta e^{i\phi}|1\rangle$ can remain coherent, $T_2|_{max} = 2T_1$. In a real-world system, there are other mechanisms, aside from the natural linewidth, which can reduce the effective T_2 . The processes we describe here introduce an extra, unknown phase to the superposition, and will thus lead to decoherence [WMI+98].

Decoherence can be quantified using Ramsey spectroscopy [LGRS04]. In this procedure, a coherent superposition $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ is prepared by first initialing to the $|1\rangle$ state and then applying a single $\pi/2$ pulse. A subsequent $\pi/2$ pulse with the same phase as the initial pulse will then take $|\psi\rangle$ to $|0\rangle$. By sweeping the phase of the second pulse, the final state can be swept from $|0\rangle$ to $|1\rangle$. Measuring a full contrast in the ability to prepare these two extremal states verifies the phase coherence of the initial superposition. If we introduce a wait time, between the two $\pi/2$ pulses, then as the coherence of the original state decays to that of a fully mixed state, due to the influence of phase fluctuations in the laser or in the atomic levels, the contrast of the measurement decreases to the point that we only ever measure $|0\rangle$ or $|1\rangle$ with equal probability, regardless of the phase of the second pulse. We measure signal contrast as a function of wait time between pulses and fit the data to a decaying exponential; the resulting decay time is our coherence time.

4.3.1 Laser Frequency Noise

No real laser is a perfect single-frequency source; practically, lasers always exhibit some amount of frequency fluctuation, however small, about a center frequency. When the laser is used to address the electronic state, the phase of the electric field is written onto the quantum state. If this phase were always fixed in time, then this added phase could be ignored in the final accounting. However, if the the laser phase is changing over time, then a differential phase is added to $|\psi\rangle$, which cannot be simply removed. A random phase fluctuation, due to the imperfectly stabilized frequency, leads to a different phase in each operation, and thus leads to decoherence. This decoherence time is related to the inverse of the laser’s linewidth [BXN⁺17].

We thus desire a way to stabilize the frequency of the laser, ideally to at least the level of the natural linewidth of our chosen qubit transition. A good deal of this stabilization comes from our use of an high finesse cavity (see section 3.4.2). As described previously, locking our laser to the cavity allows us to achieve a laser linewidth of less than 20 Hz. Once we achieve this level of stability, we must then fight the tendency of the laboratory environment to re-broaden this line. One source of laser frequency noise is in the fibers that we use to route laser light from our injection stages and amplifiers into the lab with the ion experiments. Acoustic perturbations in the fibers can modulate the optical index of the fiber, introducing effective path length fluctuations that can appear in the laser as phase noise. To mitigate this effect, we use a technique called fiber noise cancellation (FNC) [MJYH19] in order to observe and remove the added phase noise, at a cost of power lost due to diffracting through an AOM. We employ FNC on most fibers longer than 5 m on the path to the chamber. The chamber itself vibrates, as discussed in the previous section, which causes the ion to shake back and forth in the applied laser field, thus experiencing a Doppler shifted version of this beam. When possible, we attempt to prevent this vibration by ensuring that the cryocooler is well-decoupled from the rest of the vacuum chamber¹⁰. It is also possible to mitigate the vibration effects using a technique similar to FNC: a reference beam which is reflected off the vibrating cold head can be used to make an interferometric

¹⁰Common issues with the decoupling are (1) other gases in the exchange space which can freeze and introduce mechanical coupling and therefore must be pumped out, (2) liquid helium at the bottom of the exchange space which can introduce some mechanical coupling unless we heat the cryocooler above helium’s boiling point, and (3) external mechanical coupling, e.g. due to hanging cables that touch both the vibrating upper part and the stable lower part.

measurement of the motion of the cold head. This kind of vibration cancellation has been employed in lattice clock experiments [Nic06], but extending this to our experiment, in which there are other modulators and switching optics in the beam path, was not trivial. In practice, we have used a scheme in which a separate reference beam, sent along a different but similar path to the qubit beam, is used to create the interferometer signal which serves as the error signal for suppressing the noise in the qubit beam. We find that we achieve similar coherence times when using this vibration cancellation scheme to what we observe when briefly turning off the cryocooler. Vibration cancellation also comes at a cost of laser power, and thus we only employ the technique when high coherence times are necessary, as in the experiment described in section 6.3. Later on, in section 7.4.2, we will also see that photonic integration has the potential to passively suppress this sensitivity to vibration.

4.3.2 Magnetic Field Noise

Instability in the qubit transition itself can also lead to decoherence. In our chosen optical qubit, based on transitions between Zeeman sublevels, the energy spacing depends on the applied magnetic field. Thus fluctuations in the magnetic field lead to fluctuations in the qubit frequency. In all of our systems, the magnetic field comes from a Helmholtz coil configuration, which is driven by a stable current source. It can be difficult to achieve a stable magnetic field in this way, since resistance fluctuations, voltage pickup, or other noise in the current source can lead to a time-varying field. A passive solution to this problem is to use superconducting Helmholtz coils, which can maintain a persistent current to prevent changes in magnetic field when cooled below the superconducting transition. In our experiment, we use rings of niobium, thermally sunk to the 4 K stage, above and below the trap [BMS⁺19]. We leave the applied magnetic field on as the system is cooled down, but once the temperature of the rings has settled below 9.2 K, the external field can be switched off while the rings maintain the field [WLG⁺10].

In our ion, there are multiple allowed quadrupole transitions between the $S_{1/2}$ and $D_{5/2}$ orbitals, which go between different Zeeman sublevels. The Zeeman shift of these different levels is given by the Landé g-factors and leads to a slightly different frequency and magnetic field sensitivity for the different optical transitions (see section 2.6). By comparing a coherence measurement on two of these transitions, we can extract the part of the decoherence that is due to the magnetic field fluctuation, and the part which is due

to all other sources. This technique allows us to assess how beneficial the superconducting rings are. We find that the rings can improve the coherence that would be attained with an external field only, yet we still observe a magnetic-field-limited coherence in some systems, even with the rings¹¹.

¹¹The measured coherence is different for the various experiments in our lab, depending on the level of mitigation employed. We also observe different laser-limited coherence times for strontium and calcium, suggesting a measurable difference in laser performance for our stabilized 674 nm and 729 nm lasers. In our most stabilized systems, we see coherence times of around 5–10 ms.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

An Ion Trap with Integrated Voltage Sources

“Doing nothing’s sacrilegious; triple overtime ambitious.”

—Charly Bliss

In our first take on integration of classical technology into our ion trap system, we consider the integration of control voltages. There are a few advantages to moving voltage sources inside the vacuum chamber. The most obvious benefit is one of scaling; having voltage sources inside the chamber allows us to send in a few lines of high-speed serial data instead of requiring a different wire feedthrough for each of the voltages that we might want to control. Another benefit to voltage integration is the potential for increased bandwidth. Reducing the cabling between the output of the voltage source and the electrode results in a reduced capacitive load and a limited amount of parasitic low pass filtering, allowing us to approach the limit set by the electrode capacitance itself.

Some prior work has been done to test the feasibility of putting commercial voltage control electronics into a vacuum system. With a few modifications, these components can be made vacuum compatible, and some implementations have been demonstrated [GFH⁺14, ALKH13, ALS⁺16]. It has also been shown that commercial integrated circuit fabrication processes can be used to make functional ion traps [MEB⁺14], though at the start of this work no active electronics had been included in a trap substrate. In the experiments described here, we advance the integration of control voltages to the microscopic level

by presenting a design for a voltage source, with performance similar to modern external sources, that is incorporated into the substrate beneath the electrodes of an ion trap. Adding active electronics in such close proximity to the trap presents a challenge since ions are sensitive to both the static and fluctuating electric fields generated within the integrated circuit. We carefully characterize the noise and speed of these chip-integrated devices with an eye toward future quantum information applications.

In section 5.1, we present the design of our high-voltage CMOS voltage source, along with some performance characterizations, particularly at low temperatures. In section 5.2, we introduce the concept of an analog switch which is added to the design of the voltage source to act as a means for controlling the voltage noise that is exposed to the ion. With a finalized design, some experiments were performed with ions, which are presented and analyzed in section 5.3. In section 5.4, we show some measurements of the trap temperature during the operation of the voltage source and discuss some potential ways to mitigate the effects of this. Finally, in section 5.5, we consider the ultimate limits to the bandwidth of our voltage sources, set by their internal resistance and the control electrode capacitance.

5.1 Integrated DAC Design and Testing

The first challenge that we face in implementing commercially-fabricated devices in our system is cryogenic operation. Traditionally, integrated circuits can be well characterized and simulated before they are ever produced, using tools and data supplied by the manufacturer, but in our case, the specifications provided by the CMOS foundry do not characterize behavior at temperatures below 220 K. Thus we sacrifice some of the reproducibility usually afforded by CMOS and must ensure basic operation at our chosen temperature. We first test the DACs on the bench by connecting the trap chip to an interposing circuit board with several measurement test points. After verifying room temperature operation, we are then able to immerse our test rig in either liquid nitrogen or liquid helium to quickly characterize the behavior without having to install the chips into our vacuum chamber.

5.1.1 DAC Description

Our voltage source is based on an R-2R resistor ladder digital-to-analog converter (DAC) [HH15] with 12-bits of resolution and an output range of approximately ± 8 V (see Fig. 5-1).




Figure 5-1: (a) Schematic illustrating the stages of each integrated circuit block. The R-2R DAC converts a 12-bit digital code into an analog voltage from 0 V to 1.8 V. A SPI bus (not shown) programs the state of the switches in the R-2R. A unity-gain buffer (BUFAMP) follows the DAC and isolates the resistor ladder from the rest of the circuit. The high-voltage amplifier (HVAMP) circuit has a gain of approximately $9\times$ and also centers the output range at 0 V. The values of the components in the circuit are $R_1 = 70\text{ k}\Omega$, $R_2 = 2\text{ k}\Omega$, $R_3 = 8\text{ k}\Omega$ and $C = 200\text{ fF}$. (b) Layer stack-up for the CMHV7SF 180 nm node (not to scale). The total thickness of the lithographic layers is about $15\text{ }\mu\text{m}$. The lowest metal layer (M1) is made of copper, unlike the other (aluminum) metal layers, and has a higher current capacity. M1 is used for routing power around the edge of the chip. Active devices (e.g. DACs and amplifiers) are patterned in the silicon layers at the bottom of the diagram.

The DAC accepts a 12-bit code word that it translates to an analog voltage on its output (see example data in Fig. 5-2). We program the DACs using an integrated serial peripheral interface (SPI) bus, controlled by a pulse train generator (see Fig. 5-2a). The analog output of the R-2R DAC is amplified to our range of interest using a combination of two operational amplifiers. These amplifiers are made with a standard design which is composed of both n-channel and p-channel field effect transistors (FETs)¹.

The devices were fabricated as part of a multi-project wafer through MOSIS in the Global Foundries CMHV7SF 180 nm node. We chose this process for its compatibility with high voltages (up to $\pm 10\text{ V}$) which are necessary to achieve ion trapping frequencies around 1 MHz for Ca^+ in typical ion traps. The process includes 7 metal layers (Fig. 5-1b); the top-most metal layer (AM), consisting of $4\text{ }\mu\text{m}$ -thick aluminum, is used to define the serial communication, power, and test point connections (and, in the full design, the trap electrodes themselves). The remaining metal layers route digital and analog signals between the CMOS

¹See, for instance, chapter 3 of [HH15] for a primer on FETs and an introduction to the related terminology.




Figure 5-2: (a) Pulse timing diagram for programming the DAC. (b) Voltage versus code for a few test points on one of the devices. Test points are at the ends of the colored segments of Fig. 5-1a. The output of the high-voltage section uses the scale on the right side of the graph.

transistors defined in the silicon substrate and the serial inputs and control electrodes on the top metal. Since each high-voltage transistor is comparatively large ($\sim 800 \mu\text{m}^2$, about 2% of the total DAC block size), we try to use as few as possible. Hence, the first two stages of the circuit, depicted in Fig. 5-1a, are formed using only smaller, low-voltage CMOS devices. The number of transistors in the high-voltage amplifier (HVAMP) is also minimized, which reduces input impedance and limits the output range to $\pm 8 \text{ V}$ instead of the full-scale $\pm 10 \text{ V}$ set by the power supplies. Due to the low input impedance, a low-voltage buffer amplifier (BUFAMP) was added between the R-2R DAC and the high-voltage amplifier stage so that current in the DAC does not leak into the HVAMP.

5.1.2 Testing Results

To test the cryogenic operation of these devices, we use a custom interposer PCB with input lines for power and serial connections and output lines for voltages at various test points. The interposer can be attached to the bottom of a dipstick apparatus and then submerged in either liquid nitrogen or liquid helium (or just operated on the bench at room temperature). Signals are sent into and out of the dewar via a parallel cable assembly and a custom break out panel on the room temperature end.

A linear power supply (Agilent E3631A) supplies the necessary $+1.8 \text{ V}$ for the low-voltage CMOS section and the $\pm 10 \text{ V}$ for the high-voltage section. A precision DMM (Keithley 2400) is used to monitor the voltage coming out of the DAC. In order to program the DACs, we use a pulse train generator with 50 ohm output impedance (Agilent B1110A). Three serial data lines are required. A clock and data line are used to load a shift register that stores the digital code word that will be sent to the R-2R switches. On each rising

edge of the clock line, the digital value present on the data line is entered into the shift register. When the entire code word has been stored, a third input line, the load signal, is used to send a single pulse which triggers the shift register to pass the values through to the 2-2R switches. A python script was used to programmatically sweep through each code, load the pulse sequence onto the pulse generator and then trigger the data transfer. The python script can also coordinate voltage measurements using GPIB commands to the DMM.

The first version of the DAC that was tested had test points immediately after the switchable resistor ladder, after the low-voltage buffer amplifier and after the final high-voltage amplifier. We have particular interest in how the behavior of the DAC at these points changes after cooling the chip down to cryogenic temperature. We find that the behavior of the R-2R section by itself is relatively unchanged. After the buffer amplifier, some distortion appears at the upper and lower voltage limits (see left and right end of Fig. 5-2b).

5.1.3 Design Revisions

The first round of testing informed a few design revisions that were implemented in a second fabrication run. One of the changes, adding a switch to the output, was an important enough feature that it will be described in full in section 5.2. A more subtle change was increasing the size of the high-voltage FETs, which had all been drawn with the smallest possible dimensions, in order to save space. We expected that the increased transistor size would improve the linearity, and this was borne out in code-to-voltage measurements taken with the second round of chips.

The rather high power draw of the circuit (especially when running all 16 DACs in the full trap configuration) led us to try out a design which powered off some current sources in the op amps. This partial power down (PPD) was operated by a single low-voltage control signal. However, since the PPD also affected high-voltage parts of the circuit, the control signal needed to be boosted to either -8 V or $+8\text{ V}$ with a high-voltage comparator in order to fully shut off or turn on the FETs in the PPD circuit. In the second design, the PPD circuit had a parity error in which the low-voltage amplifiers would be on when the high-voltage was off, and vice versa, instead of both parts of the circuit being turned on and off together. This design error rendered several of the designs unusable, but it was possible to

fix this error without adding active components, simply by bypassing an inverter between the low- and high-voltage PPD sections. To implement the fix in our integrated circuit, we sent our chips to a rework specialist who used a focused ion beam (FIB) apparatus in order to precisely mill away material from the die and cut through a metal trace to open up a connection on the circuit. The FIB can also be used to deposit metal with similar precision, and this was done to complete the inverter bypass. The chips which had the FIB operation worked as expected². In a later revision of the design (mentioned again in section 5.4.2), the parity error was fixed, and we found that these versions of our chips worked with no further modification.

5.2 Analog Switches and Voltage Noise

After characterizing basic DC operation and linearity, we moved on to looking at voltage noise. Initial tests of voltage noise, which will be discussed below, indicated a level of technical noise which would lead to ion heating far above the typical level of anomalous heating. In the second iteration of the DAC design, we included an analog switch which allows the output of the DAC to be disconnected, attenuating voltage noise in the process.

5.2.1 Analog Switch Design and Operation

The analog switch is shown schematically in Fig. 5-3a. An additional analog line is needed to control the state of the switch. The input to this circuit consists of a high-voltage comparator and inverter, which are used to convert a low-voltage input signal into a full-range $\pm 8\text{ V}$ signal; the switch control input can be controlled using an additional output from our timing controller. For reasons that will become clear soon, and to avoid ambiguity, we have chosen to refer to this switch as an electrode isolation switch or EIS.

The EIS approximates the function of a mechanical relay, in which no current flows and the input voltage signal is perfectly attenuated when the EIS is open (as in Fig. 5-3b-c). More precisely, the EIS is a voltage-variable resistance which is relatively small when the EIS is in the closed state and is very large when the EIS is in the open state. Since there is some capacitance associated with the trap electrode, we can think of the variable resistance

²Even after the FIB, the circuit was only fully functional when cooled down to 77 K or lower. We presume this is due to the increased current capacity in the low-voltage part of the circuit when the chip is operated at low temperature. A substantial amount of current is required to drive the comparator that shuts off the high-voltage section.




Figure 5-3: (a) Schematic of the analog switch design. A complementary pair of n-channel and p-channel FETs are either both allowing or blocking current. The use of the pair allows the switch to operate over the full voltage range. A buffer and inverter, made with high-voltage transistors, boost the low-voltage control signal up to the range needed to fully turn the FETs off and on. (b) Notional diagram for the coupling from the DAC to the output when the EIS is closed. (c) Diagram for the coupling when the EIS is open. In this case, the voltage on the output discharges with some characteristic time τ , and the voltage noise is attenuated.

of the EIS instead as a variable low pass filter, with cutoff frequencies either far above or below the nominal trap frequency, depending on the EIS state. The EIS can be closed during trap voltage updates, and then while measurements are performed, the EIS can be opened to freeze the output voltages in place and greatly reduce the voltage noise.

The EIS thus turns out to be an advantageous component of the DAC design. Since there are limited low-frequency passive filtering options in an integrated platform (due to limited space to create arbitrarily large capacitors and inductors), a large resistance is a useful tool for making a low frequency cutoff with a relatively small trap capacitance. But unlike the case with a fixed frequency filter, the switchable nature of the EIS allows us to compromise between voltage noise and output bandwidth, depending on the situation. In more complex designs, the EIS could also be used for switching between multiple integrated voltage sources, which has been shown to be useful in experiments with fast ion transport [ALKH13, ALS⁺16].

5.2.2 Switch Characterization

The isolation of the EIS can be characterized by measuring the resistance of the device when it is nominally on and off. The setup for measuring the resistance is simple: apply a constant voltage and measure the current that flows through the switch. However, since the resistance when the switch is open is quite high, the current must be measured with

a picoammeter. With this measurement scheme, the resistance when the EIS is closed is approximately $3.3\text{k}\Omega$ ³. The value of this resistance is not fundamental, and can be decreased by increasing the size and current capacity of the FETs. In practice, we find that this value of resistance does not limit the bandwidth enough to preclude ion operations. When the EIS is open, the resistance is above the limit of what can be measured by the picoammeter, which places it near the $\text{T}\Omega$ range. The amount of isolation from the EIS can also be measured in the suppression of voltage noise, which will be discussed in the next section.

There is an unusual detail in the performance of our transistors at low temperatures which affects the operation of the EIS. At temperatures near 4 K, we find that the FET essentially switches off when the voltage difference between the drain and source is small⁴, regardless of the voltage at the gate. In practice, this means that when the voltage at the input of the EIS is close to the voltage at the output, then the EIS behaves as if it's open, no matter state of the control, which hints at a potential for a kind of "automatic" operation. In this mode, when the voltage on the DAC changes, then the switch closes due to the voltage difference between the DAC and the electrode and then, when the voltage has settled to the desired value, the switch opens again and reduces the transmitted noise.

However, the low-temperature behavior of the FETs is not well understood. We found many chips in our second fabrication run that exhibited behavior consistent with the above description. However, the third iteration of chips, which were produced by a different foundry, did not function in the same way. There was some indication of the "automatic" shutoff, but only at very low temperatures. It's possible that a more thorough analysis could lead to a better understand and eventually result in a design that intentionally switches itself off when voltages are not updated, but this is beyond the scope of the present work.

In the data that follows, we will report the voltage noise suppression of the EIS at 77 K, at which temperature the FETs no longer exhibit the unusual shutoff behavior. These higher temperature results are more representative of the behavior when working with ions, since we find that the temperature of the chip is elevated when it is installed in the chamber, as will be described shortly (see section 5.4).

³The exact resistance depends on the voltage across the switch and the temperature of the circuit. For the situations discussed here, in which the EIS is operated in the full 16-electrode chip in the vacuum chamber (at $\sim 60\text{ K}$), the quoted value of resistance is representative.

⁴Stated more precisely: when the value of V_{DS} is less than $\sim 0.3\text{ V}$, $I_{DS} \rightarrow 0\text{ A}$.




Figure 5-4: (a) Simplified schematic of the measurement configuration. The capacitance of the electrode, $C_{ELECTRODE}$, set by the trap geometry and the thickness of the internal layers, is ~ 1 pF. The parallel capacitance $C_{||}$ is the sum of any external capacitance added for filtering and any capacitance from signal wires or instruments. The parallel resistance $R_{||}$ is only present when the electrode noise is being measured by the spectrum analyzer. (b) Measured noise power density (square of voltage noise spectral density) on one of the trap electrodes. The voltage noise spectrum is measured using an HP89410A FFT analyzer while the trap is immersed in liquid nitrogen ($T = 77$ K).

5.2.3 Voltage Noise Measurements

The voltage noise measured in the test rig comes from a variety of sources, including pickup from the environment or power line harmonics⁵. Special care must be taken to isolate the voltage noise of the DAC. These measurements were performed using the same dewar setup, which, when properly grounded, serves as a Faraday cage for suppressing environmental noise. The measured signal is routed out of the dewar along a coaxial cable for increased shielding. Coaxial cable is also used for routing the signal outside of the dewar. The noise is measured using an HP89410A FFT spectrum analyzer, which can be configured to have a large input impedance. Voltage noise measurements are conducted using a battery to power the circuit; without the battery operation, the 60 Hz line frequency and its harmonics would appear in the 60 Hz–10 kHz range. The pulse generator which programs the DAC is disconnected from the circuit after the voltage is programmed. To avoid ground loops, the ground of the DAC and the ground of the dewar are both connected to the ground

⁵In one "Aha!" moment, it was discovered that a source of noise was the pickup of acoustic interference (i.e. music) from this researcher's laptop when collecting data. The exact coupling is unknown, but it could possibly be a residual piezoelectric effect of the ceramic capacitors soldered to the interposer PCB.

connection of the spectrum analyzer.

Any stray capacitance in the test fixture affects the bandwidth of the measurement (see Fig. 5-4a), and must also be accounted for when analyzing noise data. In the measurements discussed here, the coaxial cabling used to route the signal to the spectrum analyzer contributes about 320 pF, and the spectrum analyzer itself quotes a parallel input capacitance of 80 pF. This total capacitance, in combination with the series resistance of the EIS, creates a single pole at 121 kHz when the EIS is closed. The low pass filtering of this pair can be numerically undone by dividing by the transfer function of the single-pole filter, which gives an estimate for what the noise would look like before the EIS.

The voltage noise with the EIS open and closed is presented in Fig. 5-4b. After undoing the effect of the stray capacitive filter (dashed green line), the noise measured with the EIS closed is well-described by a $1/f$ noise model, which is consistent with the behavior of the circuit in simulation. At relevant ion frequencies, say around 1 MHz, the unfiltered noise out of the DAC is about $\sim 1 \mu\text{V}/\sqrt{\text{Hz}}$. For reference, a start-of-the-art integrated circuit DAC with a similar update speed may have a voltage noise of between 8–30 nV/ $\sqrt{\text{Hz}}$ ⁶. As discussed in section 4.1.1, the heating from this level of noise can potentially be quite high, well above 10^6 quanta/sec. As can be seen in the filtered data of Fig. 5-4b, any parallel capacitance added onto the output of the DAC will serve to reduce the voltage noise with the EIS closed, at the cost of output bandwidth. The noise measurement with the EIS open is limited by the noise floor of the spectrum analyzer (dashed blue [lower] line). The shape of the EIS-open curve is due to the Johnson noise of the 1 M Ω input resistance of the analyzer (130 nV/ $\sqrt{\text{Hz}}$) rolling off to the noise floor (8 nV/ $\sqrt{\text{Hz}}$) due to a pole created by the analyzer’s input resistance ($R_{||} = 1 \text{ M}\Omega$) and parallel input capacitance ($C_{||} = 400 \text{ pF}$).

5.3 Experiments with Ions

To use the integrated DAC design with ions, the circuit described in the previous sections must be copied for each of the control electrodes in our chosen geometry (see Fig. 5-5). We use a trap design similar to the one described in section 3.6.1 but with two fewer electrodes on each side of the trap, for a total of 16 control electrodes. In order to program all of these DACs, the shift registers at the inputs of each of the DACs are chained together. Thus, by

⁶See for instance the Analog Devices AD5790 or the Texas Instruments TI-DAC8881.




Figure 5-5: (a) Micrograph of the top metal layer of the chip. The small squares around the edge of the chip are wire bond pads, defined in the top metal layer (AM in Fig. 5-1b), for connecting power (four redundant groups of 4) and digital control (one group of 5). The two largest parallel electrodes in the center of the trap are the RF electrodes, which provide ponderomotive radial confinement and are controlled by an external source. The other 14 electrodes at the top and bottom of the chip plus 2 additional electrodes in the middle of the trap are the control electrodes, which provide axial confinement and use the integrated voltage sources. The cutaway reveals the layout of the circuit in the internal layers. The different colors in the drawing represent different signal planes. (b) Layout of one complete DAC block shown in the schematic in (d). There are 16 such blocks on the chip. The integrated devices fit beneath the trap electrodes and take up 15% of the total device area. The squares on the left side of the diagram are the high-voltage FETs; each is approximately $27\mu\text{m} \times 29\mu\text{m}$ (c) Output voltage as a function of the 12-bit digital code sent to the DAC, for each of the 16 electrodes, measured while the trap is installed in the cryostat. There is some variability in the voltage output between electrodes on the order of 1 V, and one of the electrodes has significant nonlinearity at 3 V and above. Both of these effects are calibrated *in situ* by using wire bonds to monitor the electrode voltages. (d) Photograph of the chip after attaching and bonding to CPGA interposer.

writing a $16 \times 12 = 192$ length bit string to the input to the chip, the code assigned to each DAC can be specified. There is a single load signal for the entire device, so in the current design, the electrode voltages can only be updated all at once.

In the full DAC design, the second metal layer from the top (MT in Fig. 5-1b) forms a ground plane to isolate the ion from digital and power supply noise in the internal layers of the DAC, to isolate the DAC from pickup due to the large RF voltage on the top of the chip, and to shield the ion from laser-induced effects of carrier generation in the silicon

substrate [MEB⁺14].

To verify basic operation, some *ex situ* tests were performed with the full 16 DAC device, similar to those described above. Early results indicated a significant discrepancy between different channels (as can be seen in Fig. 5-5c, which was not consistent over several chips that were tested. Much of this deviating behavior occurred at low temperatures, where there is less characterization data for the behavior of the transistors. It's possible that some more design iterations could yield a circuit which was more robust and less variable as-fabricated. In lieu of further revision, we chose to proceed with repeated *in situ* calibration of the DACs when doing ion operations. To facilitate this, each electrode on the trap was wirebonded to the pins of the CPGA carrier that are normally used to send in control voltages (Fig. 5-5d). By doing this, each electrode is connected to the parallel capacitance of the low pass filter on the octagon, 1 nF. Outside of the vacuum chamber, the voltages from each electrode are sent into a switch matrix which can route each signal one-at-a-time to a precision voltmeter (Keithley 2400). A python script is used to loop through each of the electrodes and adjust the code on each DAC until the measured voltages matches the expectation. We found that operation of the DACs in the vacuum chamber was similar to what we observed in the dewar, and the function of the DACs were not impeded by the high-voltage (up to $\sim 80 V_p$ amplitude) signal applied to the RF electrode.

A more scalable alternative to adding a monitor wirebond to each of the electrodes on the trap is to add an on-chip output multiplexer to the design of the full chip. Such a multiplexer can be used to route all of the voltages to a single diagnostic pad, obviating the need for the extra bonds. The single monitor pad can then be read out using an analog-to-digital converter, either outside of the vacuum chamber, or somewhere in the chamber in a place with a higher temperature, to ensure more predictable behavior. Every so often, the DACs can be recalibrated automatically, to keep track of any time-varying behavior of the circuit.

5.3.1 Rudimentary Motion and Frequency Control

To test basic control with the integrated DACs, we use our voltage calculation software to determine the voltages required for a certain potential, and then program the DACs (including a closed-loop calibration cycle) to match those voltages. The range of available frequencies is limited by the effective $\pm 8 V$ output range of the DACs, but we are still able

to demonstrate trapping of calcium ions at axial frequencies from 800 kHz up to 1.6 MHz. In the experiments performed here, a discrepancy is observed between the expected frequency and the frequency measured via optical spectroscopy. In all cases, the measured frequency is higher; in fact, in order to access frequencies below ~ 1 MHz, a nominally anti-confining set of voltages must be applied. This behavior, which we attribute to the effect of laser-induced charging, is not totally unprecedented, since it is sometimes observed even in standard traps near the area where photoionization lasers are incident on the trap (see section 4.1.3). It is, however, unusual that the effect is so strong on this trap. One reason for this could be the relatively large surface area of exposed dielectric on the top surface of the trap (brown ring around the outside of Fig. 5-5a); in future designs, we have extended the electrodes in the top metal layer out to the edges of the chip, such that there are no large areas of oxide exposed to the ion. It is also possible that residue from the lithographic processing steps could leave small layers of dielectric material directly beneath the ion. Further surface treatment of the trap before installation could alleviate this problem.

We performed rudimentary ion transport over a distance of $80\text{ }\mu\text{m}$ at 2 mm/s by writing a continuous stream of voltages to trap electrodes at a SPI clock speed of 200 kHz. Ion motion is verified by watching ion fluorescence on an electron-multiplying CCD camera while the ion is repeatedly shuttled back and forth by varying the DAC voltages with the EIS closed. Ions were shuttled over hundreds of cycles without loss. When DAC voltages are fixed, trap lifetimes are similar to those we observe in standard planar traps controlled by external voltages sources (>30 seconds without cooling lasers, and several hours with Doppler cooling). We verify that trap voltages do not drift when the EIS is open by monitoring the trapped ion's position and trapping frequency and find that we can leave the EIS open for several minutes with no noticeable effect to the potential at the ion's location. In this configuration, the relevant capacitance for modeling the discharge time constant is the 1 nF of the octagon filter; if the capacitance is decreased, as discussed in section 5.5, then the time constant will decrease and the discharge time of the EIS may need to be reevaluated.

5.3.2 Heating Rates

In Fig. 5-6, we present the average motional state of the ion as a function of delay after resolved-sideband cooling, from which we determine the heating rate. With the EIS closed,




Figure 5-6: Axial heating rates of a single $^{40}\text{Ca}^+$ ion with a trapping frequency of 1.5 MHz. The heating rate measured with the EIS open [closed] is 120 ± 30 quanta/s [1090 ± 20 quanta/s], which corresponds to a voltage noise of 6.0 ± 0.8 nV/ $\sqrt{\text{Hz}}$ [18.2 ± 0.2 nV/ $\sqrt{\text{Hz}}$]. The ion’s motional state with the EIS closed is determined by fitting the amplitude of the Rabi oscillation on the $S_{1/2}$ to $D_{5/2}$ transition [NRL⁺00], while the state with the EIS open is determined using the sideband amplitude ratio technique [BKRB15]. These two methods work best in different ranges of motional quanta. To facilitate plotting both results on the same scale, the linear fit for the EIS-open data is extended, and the experimental data are presented on a smaller scale in an inset.

we can compare the heating rate and voltage noise with the bench measurement in Fig. 5-4b, but we must account for additional filtering seen by the ion. The wire bonds that we use to monitor DAC voltages put the trap electrodes in parallel with extra capacitance from the filter on the PCB octagon ($C_{||} = 1$ nF). The filter capacitor alone makes a pole at 48 kHz when the switch is closed ($R_{\text{EIS}} = 3.3$ k Ω) and reduces the expected voltage noise from $0.98 \mu\text{V}/\sqrt{\text{Hz}}$ to 32 nV/ $\sqrt{\text{Hz}}$ at 1.5 MHz. This is in agreement with the voltage noise obtained from the heating rate in Fig. 5-6, 1090 ± 20 quanta/s $\rightarrow 18.2 \pm 0.2$ nV/ $\sqrt{\text{Hz}}$, within a factor of two. With the EIS open, the voltage noise is below the noise floor of our spectrum analyzer (Fig. 5-4b), so we expect a heating rate of 180 quanta/s or lower. The measured

heating rate with the EIS open, 120 ± 30 quanta/s, is below this limit, but the noise is not reduced by as much as expected given the high isolation of the EIS at lower frequencies.


5.4 Trap Temperature

One hypothesis for the limitation to the heating rate when the EIS was left open was an elevated level of anomalous heating, due to an increased trap temperature (see section 4.2 and Figure 4-1). The suspicion here was motivated by the bench measurements, in which we could clearly see that running the DACs would cause liquid nitrogen to boil off at an increased rate. The chips were also observed to heat the carrier boards to the point that they were hot to the touch when operating at room temperature. Indeed, this was the reason for including the PPD feature in the revised design. When the full 16 electrode trap is running, the power consumed by the circuit is ~ 500 mW, which must be dissipated as heat, and even with the PPD running, the power dissipation was still ~ 400 mW. Due to the limited cooling power of the cold finger, and the thermal resistance of the CPGA heatsink, we expect the temperature of the chip to increase from the baseline level in both of these cases.

5.4.1 Thermometry Results

To quantify the temperature increase of the trap when it is operating in situ, we attached a temperature-sensing diode (Lakeshore DT-670E-BR) to the surface of the chip (see Figure 5-7). The diode is attached to the chip using a small dab of silver paint. Electrical connections to the diode are made via wirebonds to unused control voltage feedthroughs. These lines have significant resistance, due in part to the phosphor-bronze wiring, but mostly to the resistors on the PCB octagon. Due to this high probe resistance, we arrange the diode measurement in a 4-wire configuration with two wires to supply current and two wires to sense voltage. The measured voltage is converted to a temperature using a spare channel of our temperature controller.

We found that the temperature of the chip increases from 4 K to just above 50 K when powered on in the CPGA mount. When the PPD is active, the temperature is increased to 43 K, which is close enough to the temperature when the chip is fully enabled that it is difficult to distinguish between these two cases with an ion measurement. At around



Condition	Temperature (K)
All Off	7.0
PPD, RF Low	41.7
PPD, RF High	43.4
Normal, RF Low	51.0
Normal, RF High	51.8

Figure 5-7: (Left) Photograph of the temperature-sensing diode affixed to the surface of the trap. The diode is mounted to one of the outer endcap electrodes. (Right) Table of measured temperatures for different configurations of the trap. With the trap totally powered down, we measure a temperature of 7 K, corresponding to the temperature setpoint of the cold stage heater. With the PPD active, we measure lower temperatures than when the trap is in normal operation. We also measure the temperature increase when the applied RF power is changed over the range used during the experiments described above. The effect of heating from the RF is not as significant as the Joule heating in the DACs.

50 K, we expect anomalous electric field noise to limit the heating rate to 30–130 quanta/s [BKR15, BSC15], which is in agreement with our measurement in this work. The range of expected levels of anomalous heating is rather wide, and it depends significantly on the material and state of the surface during the measurement. The top metal layer is aluminum, but there are several other materials which are deposited and stripped away during the fabrication process which could have left some residual contamination which may affect the noise seen by the ion. An independent test, with the trap temperature varied by an external heater and the voltages supplied by an independent source, may illuminate the actual limit of anomalous heating. Finally, it is possible that pre-treating the surface, for instance with ion milling (see section 4.2.1), could have reduced the anomalous heating background, though care will be required to not destroy the CMOS electronics when this is done.

5.4.2 In-Situ Power Down

To further mitigate the effect of trap temperature increase due to power dissipation, a more aggressive power down ability was added which switches off additional current sources

inside the op amp circuits. Adding features such as this requires additional high-voltage transistors and thus increases the size of the device, so in practice not everything can be shut off. Even in this compromised design, the amount of power drawn by the chip is still greatly reduced; we observe a power consumption decrease from ~ 1300 mW to ~ 170 mW when activating the power down on the full 16 DAC device at 295 K. Note that the power consumption in this version of the chip is higher when fully enabled, due to an increased number of total high-voltage transistors when this and the other new feature, the on-chip multiplexer, were added. We anticipate that the temperature increase when the chip is operated in the chamber will be even lower when the new power down is active—perhaps low enough that the effect can be measured with an ion’s heating rate. It remains to be tested how quickly the power down can be operated while performing experiments with ions.

5.5 Output Bandwidth

Serialization of input data comes with a speed tradeoff, since trap voltages may only be updated after sending 12 bits to each of the 16 DAC channels. From SPICE simulation [NP73], we determine the highest data rate of the SPI bus to be $\sim$ GHz (maximum voltage update rate ~ 10 MHz, due to the length of the 192-bit input string); however, in our experiment, we have achieved a maximum SPI data rate of 50 MHz (voltage update rate of 250 kHz). Above this speed, there is a sharp decrease in the fidelity of the digital-to-analog conversion. We presume this limit comes from reflections in the data lines into the vacuum chamber, since this has previously been a source of error in our experience with this SPI bus design. Using controlled-impedance lines and terminations at the chip would help to reduce these reflections.

In practice, any voltage updates on the trap electrodes are ultimately limited by the pole at 48 kHz created by the "on" resistance of the EIS ($R_{\text{EIS}} = 3.3$ k Ω) and the capacitance on the filter board ($C_{||} = 1$ nF). The relatively large resistance of the EIS is not fundamental, and can be reduced in subsequent design revisions. We could also improve this bandwidth by reducing the parallel filter capacitance, $C_{||}$, but this requires careful consideration. This capacitor also attenuates the Johnson noise associated with the large ($\sim T\Omega$) "off" resistance of the EIS; however, at relevant ion frequencies around 1 MHz, for

parallel capacitance of even a few aF, this noise source is much smaller than that due to the resistance of the electrode metal and the equivalent series resistance of the capacitor. More significantly, the additional capacitance to ground shunts leakage from the RF electrode onto the control electrode surfaces by forming a capacitive voltage divider with the effective capacitance between the electrodes [DAW⁺12]. From electrostatic simulation, we find that approximately 10% of the RF voltage, if present on the other trap electrodes, can significantly reduce the depth of the potential well seen by the ion. Acting conservatively, we might choose to keep the RF leakage below 1%. The simulated capacitance between the RF and DC electrodes is between 0.1 and 0.5 pF, so this requires a shunting capacitance of 10 to 50 pF. With these values, we would expect the bandwidth with the EIS closed to be 1-5 MHz and the isolation with the EIS open to be >160 dB, which should suppress the DAC technical noise to the same level as the Johnson noise due to the resistance of the trap metal, far below the level of electric field noise due to anomalous heating for traps of this size [BKRB15]. On our chip, a 10 pF capacitor could be integrated in the lower metal layers for each of the trap electrodes without increasing the total device area.

Chapter 6

A Compact Narrow-Linewidth Qubit Laser

“This was a triumph.

I’m making a note here: HUGE SUCCESS.”

—GlaDOS (Jonathan Coulton)

The integrated electronics we have discussed, though challenging to implement, had a clear path to implementation within an ion trap. But what about a component where the path to a smaller size is not so clear? Arguably, the largest and most cumbersome part of the experiment, aside from the chamber itself, is the qubit laser, which requires its own vacuum chamber in order to isolate the reference cavity from the environment. Even working as hard as we can to reduce the system size, [LTBR11, DRBL⁺17, KGV⁺17, DMM⁺18] there is a limit, given by the size of the cavity and the surrounding chamber, and whatever components are needed to maintain vacuum. There is no clear path for how this might be reduced to a chip-scale component.

Here, we discuss an alternative approach to laser stabilization which does not require as high of a quality factor. Instead, an optical nonlinearity, stimulated Brillouin scattering (SBS), will be used to suppress frequency fluctuations from a commercial ECDL laser, resulting in a laser with similar performance to that of a ULE cavity stabilized laser. Since this method of laser stabilization does not require exceptionally-high finesse (and a correspondingly-low loss), it can be made using integrated photonics, particularly microfabricated resonators [Vah03], which offers the possibility of a greatly-reduced cavity volume

and operation without vacuum. Here, we present the operation of an SBS laser, constructed from a ring of optical fiber, which is used as the local oscillator for an atomic clock, with short-term performance approaching that of the best single ion optical atomic clocks constructed to date¹ [BCH⁺19]. This demonstration establishes the SBS laser as a potential candidate for integration into a future trapped-ion quantum computer.

In section 6.1, we introduce the concept of the SBS laser and describe the setup used in our experiments. Our work was enabled by a novel feedforward scheme which suppressed frequency fluctuation due to temperature drift; this scheme is described and demonstrated in section 6.2. Finally, in section 6.3, we present the algorithm we use to operate an optical atomic clock with the SBS laser, including a procedure for initializing the lock when the frequency of the SBS laser is not known.

6.1 Stimulated Brillouin Scattering Laser

Stimulated scattering in optical fiber has been studied for many years [IS72]. It is often encountered as a limiting factor to the amount of power that can be transmitted in a fiber. The phenomenon arises due to light scattering off thermally-generated acoustic waves in the fiber. When the input light to a fiber (the pump beam) scatters on one of these phonons, a new photon is produced; in order to satisfy energy and momentum conservation, the following equations must be satisfied [Sch04, KSC10]:

$$\vec{k}_S = \vec{k}_P - \vec{k}_A \quad (6.1)$$

$$f_S = f_P - f_A \quad (6.2)$$

where $|\vec{k}_S| = 2\pi n/\lambda_S$ is the wave vector for the scattered wave, $|\vec{k}_P| = 2\pi n/\lambda_P$ is the wave vector for the pump wave, $|\vec{k}_A| = \omega_A/\nu_A$ is the wave vector for the acoustic wave, and ν_A is the speed of sound in the fiber; f_S , f_P , and f_A are the corresponding frequencies. Here, we have chosen to focus on the Stokes process, in which the frequency of the scattered photon is downshifted from the frequency of the pump; though the anti-Stokes process is energetically allowed, it is attenuated. Since the magnitude of the optical wavevectors are much greater than the magnitude of the wavevector of the acoustic mode, we can approximate $|\vec{k}_S| \approx |\vec{k}_P|$,

¹Compare the Allan deviation of $1.2 \times 10^{-15}/\sqrt{\tau}$, reported in the state-of-the-art single ion clock described in [BCH⁺19] with the Allan deviation of $2.5 \times 10^{-14}/\sqrt{\tau}$, described below in section 6.3.

from which we can obtain $f_A = 2\nu_A n / \lambda_P$, which is also equal to the amount by which the Stokes wave is downshifted.

When the Stokes wave interferes with the pump wave, a pattern of intensity modulation emerges, which has a frequency equal to $f_P - f_S = f_A$. Due to electrostriction in the fiber, the intensity modulation leads to a corresponding density modulation, which has the same frequency as the acoustic wave. The generation of this pressure wave reinforces the SBS process, and when the rate of scattering exceeds the loss, the production of SBS light increases exponentially. The generation of significant SBS typically occurs at high powers, in the vicinity of 1 W, depending on the exact fiber material, configuration and wavelength of light. High intensities can be reached with more modest input power with the use of an optical cavity. This configuration, known as an SBS laser [SZE91], can be constructed with a variety of optical cavities, though here we focus on a configuration in which the cavity is a ring of optical fiber, into which power is coupled via an evanescent directional coupler.

The Stokes wave, the primary output of the SBS laser, has interesting and useful noise properties. Frequency fluctuations in the pump laser are primarily transferred into the acoustic mode, and thus the SBS laser frequency noise is suppressed by a factor approximately equal to the ratio of the loss in the acoustic mode to the loss in the optical mode, which in our setup is about 34 dB of suppression² [DRZ00, LPD15]. Since the linewidth of the pump laser is already reduced by locking to the fiber cavity, the further suppressed noise in the SBS light can lead to low linewidths, competitive with those achieved with ultra-high finesse optical cavities. One of the biggest advantages of the SBS laser is that the choice of wavelength can be quite flexible; as long as a low-loss medium is available for the chosen wavelength, and a resonator can be made with appropriate acoustic properties, then SBS lasing can be achieved. A variety of different material platforms have been used to make SBS lasers, some of which can be made quite small, down to the level of a few cm² [GSW⁺06, GMM09, SAMI⁺11, LCL⁺12, TLF⁺12, KPC⁺13, LGB⁺15, LBC⁺16, OBK⁺18, GBP⁺19].

One of the primary drawbacks of the SBS laser is the instability of the resonance frequency of the resonator. While the SBS signal has its frequency fluctuations suppressed

²One way to intuitively see the noise suppression is to consider the energy balance in equation 6.2; when the frequency of the pump f_P fluctuates, some combination of f_A and f_S must change to compensate. Since the time constant of the acoustic mode is much faster, it is able to compensate for most of the shift, before f_S has to change.




Figure 6-1: (a) Diagram of the SBS laser setup. An AOM, electro-optic modulator (EOM), semiconductor optical amplifier (SOA), polarization rotator (Rot), polarization beam splitter (PBS), photodetector (PD), and local oscillator (LO) enable the independent locking of a pump laser to two orthogonal polarization modes of the resonator (Res). Two counter-propagating SBS lasers are created from the pump lasers. (b) Photograph of the SBS resonator resting within a copper enclosure. The SBS resonator including the enclosure and lid occupies a total space of 3.1 in. $\times$ 3.5 in. $\times$ 1 in.

with respect to the center frequency of the resonator, if the resonance itself shifts, then this noise can be written directly onto the SBS laser. Thus some care must be taken, as with a ULE cavity, in order to isolate the resonator from the environment. Unlike the ULE cavity, which has very low temperature sensitivity, an optical resonator can have a sharp dependence on temperature, due both to length and index effects [NTHS93]. One of the major innovations of this work was the scheme used to suppress temperature fluctuations, which will be discussed in detail in section 6.2. First, we turn our attention to the generation and vibration stabilization of our SBS laser system.

6.1.1 Setup for Generating SBS Laser Light

In our setup, we use a laser and resonator at 1348 nm, half the frequency of the optical transition in strontium. Most components are commercially-available PM fiber-coupled devices which we either splice together or connect with fiber sleeve couplers. The primary laser is a commercial external-cavity diode laser with 30 mW of power. Before being sent into the fiber resonator, the light is phase modulated using a fiber EOM. The resonator itself is constructed using a fixed ratio evanescent coupler (Evanescent Optics 954P), where

two ends are spliced together to create a loop of fiber approximately 2m, wound around a 2-inch diameter mandrel that rests within a $3.1 \times 3.5 \times 1$ -inch temperature-controlled copper enclosure. The resonator is slightly overcoupled in order to improve the output SBS power; the loaded Q in our setup is 1.9×10^8 . With 10 mW of pump power, we observe about 5 mW of SBS laser light. The SBS light can be separated from the pump with an optical circulator placed before the resonator. The pump laser is fixed onto resonance by demodulating the transmitted signal via the PDH technique [DHK⁺83], with feedback to the current and piezo of the ECDL. We pick off 10% of the output SBS power coming from the resonator in order to be able to monitor and control the optical power in the resonator. Note that, for our temperature compensation scheme, it will be important to generate SBS light in two distinct polarizations, which is why we show two of many of the components in Figure 6-1(a); this will be explained further in section 6.2.

The output of the circulator is amplified by an SOA and then frequency doubled through a fiber-pigtailed KTP waveguide doubler for interrogation of an $^{88}\text{Sr}^+$ ion at 674 nm. For an amplified SBS power of ~ 60 mW at the input of the frequency doubler, the 674 nm output after doubling reaches $\sim 60 \mu\text{W}$, corresponding to an efficiency of 1.7%/W. In the future, it may be possible to achieve higher output power by fabricating the frequency doubler onto the same substrate as the resonator (while also integrating the other required components). For our near-term purposes, this amount of power is sufficient to seed the injection system that we normally use for our qubit laser (see section 3.4.2). We are thus able to use all of the same downstream fiber and free space components as in our standard ULE-reference operation when interrogating the ion with the SBS-based light.

6.1.2 Vibration Isolation

Though we do not require a vacuum environment to stabilize our resonator, we must still employ some amount of vibration and acoustic shielding. Most of our components sit on top of an active vibration isolation table; however, the bulk of our isolation comes from a steel box, approximately an 8 in. cube, which sits on top of a ring of sorbothane (see Figure 6-2). This lid isolates the interior from both sound in the room and from temperature fluctuations due to convective airflow in the room. Inside of this box, we have a ~ 1 in. thick steel plate which sits on a layer of sorbothane, to act as a miniature optical table (called a "Baffle" in Figure 6-2). On this plate, we have the majority of the passive optical fiber components,




Figure 6-2: SBS vibration isolation setup. The resonator box (Figure 6-1b) sits at the top of a stack. The fiber mandrel is thermally separated from the outside of the small copper box by a layer of Viton, in order to reduce the thermal time constant of the resonator response. A Peltier device and copper heatsink, along with a thermistor embedded in the copper block, allow us to control the temperature of the enclosure to about the 1 mK level. Layers of sorbothane bumpers and a steel plate (baffle) provide vibration isolation from the optical table. The stack is surrounded by a steel box which sits on a gasket made of a sorbothane sheet; we find that the addition of this shield is important to limit acoustic noise in the room and to improve the temperature stability. The whole assembly shown here sits on top of an active vibration isolation stage, though we did not find a noticeable difference in operation when the vibration cancellation was running. It is possible that many of these layers could be reduced in size or perhaps eliminated with further engineering.

including circulators, couplers, and polarization beam splitters—all spliced together.

The resonator itself sits just above this, within a separate, smaller copper enclosure (Figure 6-1b), comprising a copper platform that rests on a layer of Viton. The enclosure isolates the resonator from the environment but permits the ability to control the resonator’s temperature through the copper. After winding the resonator coil, the inner enclosure is hermetically sealed using Torr seal. The temperature of the resonator enclosure is controlled using a precision thermistor and a peltier device thermally connected to a copper heatsink.

6.2 Temperature Feedforward Scheme

Our method for sensing the temperature of the resonator is based on the differential temperature sensitivity of two orthogonal polarization modes within the resonator [STB⁺11, FAS⁺12, WAS⁺14]. The ability to sense temperature with this method has been limited by noise in the probe lasers and issues with temperature control to around $10\,\mu\text{K}$ [LLS⁺19]. Here, we use the naturally-low linewidth of the SBS laser as a temperature sensor with resolution below 100 nK [LYOJ19]. To circumvent problems with the bandwidth and resolution




Figure 6-3: (a) Diagram of the SBS stabilization circuitry. The beat note of two SBS lasers is first converted to a voltage signal through a voltage-controlled oscillator (VCO) operating in a phase-locked loop (PLL). Afterwards, a linear ramp subtraction, low-pass filter (LPF), and factor of 39 multiplication is applied before this signal is used for correcting the SBS laser drift. (b) Simplified diagram illustrating the correction of the SBS frequency drift via subtraction of the free-running SBS laser and a $39\times$ multiplied polarization beat signal.

of temperature control (e.g. with a heater element)³, we instead use a feedforward system in order to remove the temperature-dependent component of the frequency fluctuation from the laser [LSR⁺20]. This suppression of temperature fluctuation brings the stability of the laser into the regime that we are able to perform spectroscopy and then execute a clock protocol in order to lock the laser frequency to the atomic resonance.

6.2.1 Extracting the Temperature Signal and Applying the Correction

As indicated in Figure 6-1(a), our SBS laser is composed of two separate control loops, which are used to generate SBS laser light for both polarization modes. In order to easily separate the two lasers, we send the two different polarizations in separate directions, such

³Using feedforward instead of controlling the temperature also avoids unintentional length shifts from thermal expansion and contraction, which can also affect the resonance frequency.

that, with the use of polarizing splitters and circulators, the two pump signals and the two SBS signals can all be differentiated. The two pumps each come from the same primary laser, which is split at a 90:10 ratio between two separate paths. The 90% path passes through an additional AOM, which is used both to shift the light by ~ 180 MHz and to enable independent control of its frequency. Unlike the first pump, the second pump is passed through a polarization-rotating fiber before it is launched into the resonator. Both pumps are phase modulated for PDH locking. The first pump is controlled by modulating the current of the ECDL, and the second pump is locked via the AOM.

As shown in Figure 6-3, the two SBS outputs are directed to a photodetector, where they interfere to produce a beat note at 180 MHz. Since vibration noise is common to both SBS laser signals, the primary difference is the difference in their shift of optical index. The frequency shift of the beat note is thus a sensitive proxy for the temperature of the resonator. The frequency shift of the beat note is related to the temperature-induced frequency shift of either SBS laser by a multiplicative factor, which we determine empirically. If we were to naively apply the beat note signal as a correction, after multiplication by some method, then we may end up writing extra noise onto our laser, since the phase noise on the beat note will be increased by frequency multiplication. In order to limit the bandwidth of the correction, we first convert the frequency to a voltage signal using a PLL. Low passing the voltage signal isolates the part of the signal which is related to temperature fluctuation. In our experiment, we low-pass the signal at 0.03 Hz and amplify the result by a factor of 2 using an SR560 low-noise amplifier. At this point, linear drift in the output signal, discussed in more detail below, can also be removed by a voltage ramp; this is done using the inverting input of the SR560. The resulting signal is then converted back into the RF domain using a VCO. For the VCOs used both at this stage and in the PLL, we use function generators (Rigol DG4062) with frequency modulation of adjustable depth. We found that standard crystal VCOs were too sensitive to ambient temperature and that the function generators were much more stable. The adjustable modulation depth also allows us to fine-tune the overall multiplication factor to match the desired $39\times$.

6.2.2 Feedforward Simulation

To tune up this procedure, both the frequency of the free-running SBS laser's beat note with a reference ULE laser and the frequency of the dual-polarization beat note are experimen-




Figure 6-4: (a) Experimental time series of the SBS laser (red) and the polarization beat-note (blue) drift. A numerical subtraction of the two (black) produces a residual linear drift. (b) Numerical removal of the linear drift from the SBS and polarization beat note time series revealing the underlying temperature-induced frequency drift. The excellent agreement between the two demonstrates that the correction signal accurately tracks the SBS temperature drift. (c) Subtraction of the SBS and polarization beat-note signals in (b). The drift is numerically canceled through the linear ramp and temperature correction procedure. (d) Fractional frequency noise measurements of the SBS laser. The green shaded section indicates the region where the SBS laser noise is gradually averaging down. We lock the SBS laser to the ion at these time scales for the demonstration of the optical atomic clock. The blue section indicates where the SBS laser experiences drift that can be mostly compensated for via subtraction. The red shaded section indicates where the SBS laser drift dominates. After subtraction, the resulting noise of the SBS laser reveals a long-term frequency variation of 220 Hz, which can be further reduced by the lock to the ion

tally measured at time intervals of 1 ms and plotted for the wavelength of 1348 nm over a period of 12 minutes (Figure 6-4(a)). The polarization beat is also numerically multiplied by the empirically-determined correction factor of 39 in order to account for the common-mode suppression of the extracted temperature shifts⁴. The resulting multiplied signal serves as

⁴The multiplication factor is determined by measuring the frequency of the beat note before and after applying a 0.25 K temperature shift using the temperature controller for the resonator. The factor is then fine-tuned by adjusting it until no visible correlation remains in the corrected signal. We find that the

the feedforward correction to temperature drift, which upon numerical subtraction from the free-running SBS laser yields a residual linear drift of 160 kHz/minute⁵. This linear drift results from the individual linear drifts of the SBS and polarization beat signals, which is in opposing directions for the case shown in Figure 6-4(a), and represents a parasitic frequency shift that our temperature sensing technique cannot account for. We attribute the observed linear drift to a slow relaxation of the SBS resonator over time [SBJ⁺98]. When the linear part of the drift of the SBS and polarization-beat traces is removed, the underlying temperature-induced fluctuation can be seen, as in Figure 6-4(b). The multiplied polarization beat signal tracks the SBS signal closely, and when the two signals are subtracted, the resulting stabilized SBS signal is limited to frequency fluctuations at about 220 Hz, shown in Figure 6-4(c).

Another way to visualize the effect of the feedforward correction is with an Allan deviation. In Figure 6-4(d), we see the fractional frequency noise derived from the same data plotted throughout the rest of Figure 6-4. The free-running 1348 nm SBS Laser reaches a minimum noise level of 1.4×10^{-13} (corresponding to a linewidth of 30 Hz) at 10 ms but becomes unbounded at longer time scales and increases to 5.8×10^{-10} at 100 s. When the measured drift is subtracted to account for both the linear and temperature-induced drift, the SBS laser maintains its performance of 1.0×10^{-13} (22 Hz) at short time scales and experiences a >500-fold reduction in drift over the long-term. The SBS laser frequency excursions become bounded at the value of 1.0×10^{-12} or 220 Hz, which corresponds to an effective temperature stabilization of the SBS laser at a level below 120 nK.

6.2.3 Feedforward Results

Next, we use our feedforward apparatus to try to match the simulated behavior described in the previous section. The drift of the SBS laser is monitored by a beat note between the doubled output of the SBS system and the ULE-cavity stabilized laser in our lab (described in section 3.4.2). We use a function generator and mixer combination in order to shift the beat note frequency into the range of the various instruments that we use to characterize the signal. In Figure 6-5(a), the SBS laser starts initially in a free-running state for the first 6.5 minutes of measurement and drifts by 400 kHz within this time. Once the feedforward

accuracy of the correction signal does not degrade significantly for scale factors of 39 ± 2

⁵This rather high drift was recorded shortly after the resonator was constructed. After a few months, the drift of the same device had reduced to 30 Hz/s.




Figure 6-5: (a) Measurement of the SBS laser's frequency drift (solid line) with and without feedforward stabilization. The measurement interval was 1 ms. (b) Zoomed-in plot of the SBS frequency drift. The average long-term frequency deviation is 310 Hz. (c) Measured lineshape of the stabilized SBS laser (red solid line) and corresponding Voigt fit with $R^2 = 0.93$ (dashed line). The spectrum is taken with a sweep time of 225 ms and a resolution bandwidth of 20 Hz. (d) Fractional frequency noise comparison between the stabilized SBS laser (red line) and the ideal drift-canceled SBS laser of Fig. 2d. Long-term temperature stabilization at the level of 170 nK is experimentally achieved.

correction is applied, the SBS frequency drift flattens significantly. The zoomed-in trace of the stabilized SBS frequency drift (Figure 6-5(b)) shows close agreement to Figure 6-4(c) (accounting for the frequency doubling) and indicates that our experimental implementation of the feedforward stabilization accurately compensates for both linear and temperature-induced SBS laser drift.

The lineshape of the SBS laser subsystem (Figure 6-5(c)), measured at 674 nm via beating with the ULE laser, demonstrates further the laser's exceptional short-term noise. The spectrum exhibits a linewidth of 50 Hz. This value of linewidth is confirmed by the measured fractional frequency noise of the stabilized SBS laser (Figure 6-5(d)), which reaches a minimum of 1.1×10^{-13} at 60 ms and corresponds to a frequency deviation of 48 Hz at




Figure 6-6: (a) Schematic of laser beam path to the clock chamber. A flipper mirror allows the SBS laser subsystem to be interchanged with a BCS laser as the master frequency source. The injected laser and tapered amplifier represent a series of two injection locking and amplification stages, respectively. (b) Measured Ramsey fringes on the clock transition. Two $\frac{\pi}{2}$ pulses are applied around a $\tau = 1$ ms interrogation period. The population in the lower clock state is measured as a function of the phase of the second pulse. (c) Spectroscopy of the $|5S_{1/2}, m_J = -1/2\rangle \rightarrow |4D_{5/2}, m_J = -3/2\rangle$ transition in $^{88}\text{Sr}^+$ measured with the SBS laser subsystem. Spectroscopic measurements are interleaved with frequency corrections to keep the SBS laser subsystem locked to the atomic transition. The vertical blue bars represent 1σ error derived from photon counting statistics.

674 nm. At long time scales, the noise level becomes 1.4×10^{-12} (620 Hz), which is slightly larger than the frequency excursions found with the numerical drift-canceling procedure of 6-4(d). We attribute this difference in drift to noise in the electronics used for feedforward stabilization. At the level of 620 Hz, our achieved frequency drift corresponds to an effective temperature stabilization of the SBS laser subsystem at a level below 170 nK.

6.3 Clock Operation

After temperature feedforward, the drift of the SBS laser is low enough that it can be used in a clock protocol [MDZ⁺12]. The clock protocol, while interesting in itself, is also a good test of the control that would be required for quantum computation. By stabilizing the SBS laser to an ion, we achieve a laser source that is on par with the laser we would normally use for our other experiments.

6.3.1 Measurement Process

For our clock transition, we use the $S_{1/2} \leftrightarrow D_{5/2}$ quadrupole transition in $^{88}\text{Sr}^+$ (0.4 Hz natural linewidth). As mentioned previously, and shown in Figure 6-6(a), the 674 nm out of the SBS system is used to seed the injection and amplification stages in route to our vacuum

chamber. In this setup, we can use a single flipper mirror to select either an existing the transmission of the ULE cavity or the SBS laser subsystem as the initial seed for the injection stages. The chamber that we use for this experiment employs all of the frequency noise mitigation techniques discussed in section 4.3.1, including superconducting rings and active vibration cancellation. To baseline the apparatus, we use the ULE laser as a seed for the injection, and perform a series of Ramsey experiments to deduce a coherence time of 2.9 ms, which we take as an upper bound for the performance of the clock experiment. Assuming that frequency fluctuations in the laser are the dominant decoherence mechanism, we infer an effective laser linewidth of $1/(2\pi \cdot 2.9 \text{ ms}) = 55 \text{ Hz}$, which is close to the SBS laser subsystems's linewidth in Figure 6-5(c). Thus we conclude that uncompensated noise in our experiment (e.g. from optics subject to acoustic vibrations or magnetic field instability at the ion) contributes at about the same level as noise from the SBS laser subsystem. All together, these noise sources limit the clock interrogation time we achieve, and therefore set the ultimate limit in clock performance.

An AOM is used both for scanning the SBS laser subsystem over the clock transition and for locking the laser's frequency to the atomic resonance. To discipline the SBS laser subsystem to the atomic resonance frequency, we create an error signal via a Ramsey experiment consisting of two $\frac{\pi}{2}$ pulses on the clock transition⁶, separated by an interrogation time $\tau = 1 \text{ ms}$. As the phase of the second $\frac{\pi}{2}$ pulse is varied, the population in the lower clock state traces out a sine curve (Figure 6-6(b)). The slope of this signal reaches a maximum when the second pulse is 90° out of phase with the first pulse. Variations in the laser frequency during the interrogation time result in an additional phase shift and are thus mapped to the ion's state distribution. Between interrogation cycles, the state of the ion must be detected and then re-prepared into the lower clock level; this leads to a 1.85 ms dead time, during which we are insensitive to frequency fluctuations of the laser.

After completing this Ramsey sequence, we measure the ion state via state-dependent scattering of photons on the cycling $S_{1/2} \rightarrow P_{1/2}$ ion transition. After determining which of the two clock states the ion is in, we update a frequency correction applied to the ion as follows:

$$\Delta f_{i+1} = \Delta f_i + \frac{1}{2\pi} \frac{\alpha m_i}{\tau}, \quad (6.3)$$

⁶Both of the $\pi/2$ pulses actually consist of SCROFULOUS pulse sequences which mitigate the effect of amplitude fluctuations [CLJ03].

where $\tau = 1$ ms is the interrogation time, $\alpha \approx 0.2$ is a gain term optimized to achieve the best stability, and m_i represents the measurement result, with $m_i = +1$ (-1) corresponding to finding the ion in the $D_{5/2}$ ($S_{1/2}$) state. The measured noise of the SBS laser can be used to create a model of the system's behavior when the clock is running, which allows us to tune the value of α and estimate the performance. In the simulation, the laser frequency is updated once per clock cycle based on simulated ion measurements that take into account quantum projection noise and dead time effects.

As a demonstration of the efficacy of the atomic lock, we perform Rabi spectroscopy of the clock transition using the SBS laser subsystem, interleaved with frequency corrections following the protocol above (Figure 6-6(c)). These spectroscopic measurements are performed with pulse time of 10 ms (to avoid Fourier broadening of the line), which leads to a longer clock dead time of ~ 11 ms for these measurements. The width of the measured feature, 370 Hz, is in close correspondence with the simulated value of 250 Hz, derived from numerical application of the clock protocol to the free-running SBS laser noise, given the effective dead time of 11 ms. This provides further evidence that noise mechanisms downstream from the laser source do not substantially affect measurements with the SBS laser.

6.3.2 Finding the Line from Zero

In our initial experiments with the SBS laser subsystem, we used our ULE laser as a stable reference to monitor the linear drift of the SBS laser by tracking the beat note between the two optical signals. This use of the ULE laser was merely a convenience, and is not required for the clock application. When nothing about the drift of the SBS laser is known, the atom itself can be used as a stable frequency reference. The frequency of the full SBS laser subsystem must first be brought to within a few tens of MHz of the atomic resonance frequency; this degree of frequency determination is afforded by commercial wavelength meters based on Fizeau interferometers. When the laser frequency is close to the clock transition, we may perform Rabi spectroscopy by scanning the frequency of the laser using an AOM and search for atomic transitions. With our knowledge of the magnetic field, the laser polarization and the frequency of the ion's motion we are able to uniquely identify our clock transition. Next, we perform repeated spectroscopy across the clock transition (see Figure 6-7); since the drift of the SBS laser's frequency is the dominant source of frequency




Figure 6-7: (a) Rabi spectroscopy of the $|5S_{1/2}, m_J = -1/2\rangle \rightarrow |4D_{5/2}, m_J = -3/2\rangle$ clock transition and the first-order motional sidebands at $\nu_{\text{clock}} \pm \nu_{\text{trap}}$ is taken at regular intervals of approximately 25 seconds. After performing fits to the symmetric sidebands, we can average the two frequencies to obtain an accurate measure of the frequency of the central feature. Only the spectroscopic data for the final experiment (black line) is shown; for all other data sets, the Gaussian peak fits to the sidebands (red and blue curves) are shown with progressively darkening color to illustrate the movement of these features over time. For these data, an intentional linear drift of 5 kHz/sec (at 674 nm; equivalently 2.5 kHz/sec at 1348 nm) was applied to demonstrate the efficacy of this method in cases of high drift. (b) Rabi spectroscopy of the clock transition and sidebands after applying a linear drift correction to null out the natural drift of the resonator's frequency. Over the course of 20 minutes of measurements, very little deviation in the center frequency is observed. (c) Linear drift determined from the data presented in (a) and (b). The linear drift can be obtained from a fit to the apparent frequency of the clock transition as a function of time. In the first case, with the large drift intentionally applied to the laser frequency, we obtain a drift of 5.2 kHz/sec (at 674 nm) from the fit. After few iterations of applying a correction and measuring the resulting drift, the drift is driven down to 17 Hz/sec, which is lower than the natural drift observed in this resonator at this time. (d) Integrated clock correction signal applied to the laser to keep the frequency resonant with the atom's transition. In this case, we use a simplified clock protocol with an interrogation time of $\tau = 100 \mu\text{s}$.

error, any apparent shift in the clock transition is attributed to the SBS laser. See also Ref. [ANK⁺15] where a similar method is used to track the drift of a BCS laser.

After waiting long enough the linear drift can also naturally decrease to a level that is suitable for locking with no further correction. We performed these experiments several months after the other experiments discussed in the rest of this chapter, so the resonator had a lot of time to relax to a lower value of linear drift, down to approximately 30 Hz/sec (at 1348 nm), at which level we have been able to perform our clock protocol. In order to show that our method here would be effective even when the drift is high, we intentionally applied a frequency shift using a mixer and a function generator with a programmed frequency ramp to collect the data in Figure 6-7.

Once the linear drift is brought down to an acceptable level, the error signal of the clock itself can be used to track the frequency of the SBS laser subsystem with respect to the frequency of the atomic transition. By accumulating the frequency corrections applied to the final frequency-shifting AOM, the frequency difference between the SBS laser subsystem and the clock transition can be monitored in real time (see Figure 6-7(d)). For this real-time frequency tracking, we run our clock protocol with a reduced interrogation time, which leads to a decreased frequency resolution but an increased capture range for the lock signal. For future operations, we may construct our control software such that a limited version of the clock protocol runs between all other calibration routines and experiments. The frequency difference could thus be constantly tracked and programmatically fed forward into each experimental operation.

6.3.3 Results

To assess the stability of the clock when running with the SBS laser subsystem, we perform a self-comparison measurement via two independently-operated clock signals generated by interleaving distinct sets of correction signals applied to the laser. This technique has been previously used to characterize clock performance when two independent clocks are not available [NMW⁺12, NCH⁺15] and is known to accurately capture the short-term clock stability, including ion projection noise and the Dick effect due to dead time in the clock protocol. To obtain the two clock signals, we run two independent control loops according to the protocol described above. During odd-numbered interrogations, we use and update the first frequency correction $\Delta f_i^{(1)}$; during even-numbered interrogations we use and update




Figure 6-8: (a) Pictorial representation of the interleaved clock procedure. Here the Doppler segments represent the $700\,\mu\text{s}$ duration in which the ion is Doppler cooled. During the OP segments, the ion undergoes $450\,\mu\text{s}$ of optical pumping in order to prepare the electron in the lower level of the clock transition. The Interrogate segments are each $1\,\text{ms}$ of interrogation time, bounded by composite $\pi/2$ pulses. Lastly, the Detect segments are $700\,\mu\text{s}$ of detection time, during which the photons emitted by the ion are detected on a photomultiplier tube and counted by our timing controller. During the Update segment, and depending on the number of photons collected, the state of the ion is determined, and the frequency of the clock is either increased or decreased. As discussed in the text, two separate clock signals, $f^{(1)}$ and $f^{(2)}$, are maintained; here these are indicated as Clock #1 and Clock #2. While the frequency of either clock is updated, the experiment begins to prepare the state for the next measurement, as indicated by the black arrows. Each of these clocks is sensitive to laser frequency fluctuation only during the $1\,\text{ms}$ interrogation period of the total $5.7\,\text{ms}$ cycle time; during all other times, the frequency of the laser must stay within the capture range of the lock. (b) Allan deviation of the difference frequency between interleaved clocks. The measured Allan deviation is divided by $\sqrt{2}$ to estimate the error of a single clock, assuming even distribution of error in the correction signals. Here we report the so-called "totdev" measure of frequency variance, which we find has more stable behavior at long sample times [How00]. The vertical blue bars indicate the 1σ error in the Allan deviation.

the second frequency correction $\Delta f_i^{(2)}$. The timing of the interleaved sequences is shown in Figure 6-8(a). After the experiment is over, we can then take the Allan deviation of the frequency difference $\Delta f_i^{(2)} - \Delta f_i^{(1)}$.

The self-comparison technique is insensitive to long-term frequency drifts of the ion, which if present would be common to both clock signals, but at the same time under-

estimates the performance of the clock due to the extra dead time incurred through the interleaving operation. We note that our clock is interrogated on the magnetically-sensitive $|5S_{1/2}, m_J = -1/2\rangle \rightarrow |4D_{5/2}, m_J = -3/2\rangle$ transition, which would limit long-term stability if magnetic field drifts were not canceled; however, a simple protocol for $^{88}\text{Sr}^+$ clocks has been developed [DMB⁺05] which eliminates both first-order Zeeman and electric quadrupole shifts of the clock transition and can be implemented in future measurements.

For a 1 ms interrogation time followed by a 1.85 ms recoiling and state preparation period, the effective dead time for each clock is 4.7 ms and results in a measured clock stability of $3.9 \times 10^{-14}/\sqrt{\tau}$ (Figure 6-8(d)). This value of frequency stability agrees well with numerical simulations we performed of the optical clock using the measured SBS laser subsystem noise, which predict the clock to operate at the level of $4.1 \times 10^{-14}/\sqrt{\tau}$ for 4.7 ms of dead time. The good correspondence between measurement and simulation suggests that under normal clock operation with 1.85 ms of dead time, the clock stability would reach $2.5 \times 10^{-14}/\sqrt{\tau}$. This stability is beyond the short-term stability achievable by the best microwave clocks [JMS⁺00] and close to an order of magnitude away from state-of-the-art laboratory single-ion clocks [BCH⁺19].

Chapter 7

Ion Traps with Integrated Photonics

*“Work it harder. Make it better.
Do it faster. Makes us stronger.”*

—Daft Punk

We now turn our attention to another classical technology which can be readily integrated into a surface electrode ion trap, namely integrated photonics. In this case, while the co-fabrication of integrated photonics was a huge step in itself [MBM⁺16], there was a great deal of work left to be done to couple light into these devices in a scalable way. On top of this, close investigation of the behavior of ions revealed many non-ideal features of the photonics-enabled chip, some of which could be mitigated and others can be addressed in future fabrications.

Here we present the construction, packaging and installation of chips with multi-color integrated photonics. Coherent control of the electronic state of trapped ions reveals a benefit inherent to integrated photonics, arising from the natural vibration insensitivity of registering light sources to the trap chip. We also take an early look at a trap design with multiple photonic zones which offers the potential for performing multiple, distinct operations simultaneously using global beam operations and local changes to the voltages controlling the ions.

In section 7.1, we introduce two major integrated-photonic components, the waveguide and the diffractive grating coupler, and then describe how these can be used to construct an

ion trap with on-chip laser routing and focusing. The procedure and techniques for installing an ion trap with integrated photonics into our cryogenic system are presented in section 7.2. A few technical challenges which contribute to the noise that we observe on our photonics chips, along with some potential solutions, are discussed in section 7.3. In section 7.4, we overview some single-ion results that were achieved using laser light routed to the ion via integrated photonics, and here we focus on an improvement to qubit coherence attributed to the use of photonic structures. Finally, in section 7.5, we present a new trap design with multiple photonic zones which was used to demonstrate the way in which integrated optics can be used along with control of the ions' positions.

7.1 Integrated Photonics Toolkit

Over the past few decades, the field of integrated photonics has flourished with many technologies which have pushed the boundary of modern telecommunication. Most commercial work in photonics is driven by applications in the near-infrared and infrared, particularly in the so-called C-band near 1550 nm, where high-power optical amplifiers are readily available at low cost. Recently, many of the photonic structures are finding their way into platforms which are more amenable to operation in the visible [SABK⁺18, SAKY⁺19, WLK⁺19]. Many of these devices, if fabricated at wavelengths compatible with our ion species of interest, could have immediate benefits for overcoming the size requirements of free-space optics. To date, a chip-integrated single-mode waveguide has allowed coherent operations on a trapped ion [MBM⁺16]. However, the monolithic integration of optical components for delivery of all light required for basic ion-qubit operations has remained an outstanding challenge.

Here we will focus on structures which are fabricated primarily in a silicon-nitride (SiN) process. Our photonics chips are fabricated in a foundry at Lincoln Laboratory, with the help of photonics designers and fabrication specialists. All devices consist of a silicon substrate onto which additional layers are deposited and patterned with optical lithography. Photonic elements are created by patterning a high index material, in this case SiN deposited by plasma-enhanced chemical vapor deposition (PECVD), and then surrounding the SiN with a lower index material, in this case SiO₂ deposited by PECVD. The SiN and SiO₂ have indices of refraction of $n_{\text{SiN}} = 1.89$ and $n_{\text{SiO}_2} = 1.5$ measured at $\lambda = 633$ nm.

7.1.1 Waveguides

Optical waveguides are structures consisting of a silicon nitride (SiN) guiding layer surrounded by silicon dioxide (SiO₂) cladding; each waveguide is patterned to have a width that ensures single-mode operation at its design wavelength (ranging from a 250 nm width for a design wavelength of 405 nm, to 1.1 μ m for 1092 nm). These polarization-maintaining, single-mode waveguides allow flexible routing of light to arbitrary locations below the chip surface.

One of the major challenges to overcome with integrated optics is loss in optical power, which can arise due to absorption in the material or scattering off roughness in the sides of the waveguide [BHJ⁺11]. The SiN waveguides have propagation losses below 0.5 dB/cm for wavelengths above 633 nm, with losses increasing to $\sim$ 10 dB/cm at 405 nm (for more details see ref. [SAKY⁺19]). Since most of the optical path lengths on our chips are on the order of 1 cm, loss in the waveguide can be an appreciable part of the total loss for the bluer wavelengths. A new process has recently been developed which uses Al₂O₃ as the waveguide material, which has around an order of magnitude lower loss at blue frequencies of interest for strontium and calcium [WLK⁺19]. The final experiment described in section 7.5 is done with a chip from this new, low-loss process, but everything else is done with only SiN structures.

Light is coupled from optical fiber onto the chip via inverse-taper waveguide couplers, sections of waveguide that are narrowed such that the spatial mode more closely matches the optical fiber mode [APL03]. The light is sent in linearly polarized in the plane of the chip, and the quasi-transverse-electric-field mode excited in the waveguides maintains this polarization in the plane of the chip and perpendicular to the propagation direction. The polarization is maintained due to the substantial birefringence in the waveguides, which arises from the difference of the spatial modes associated with orthogonal polarizations of the light. The on-chip coupling (OCC) loss, arising from the optical-mode-size mismatch between that of the fibers and the integrated waveguides, is one of the main sources of total power loss in our devices. This could be improved by expanding the mode in the waveguides, e.g., via thinning or further lateral tapering of the SiN core or via the use of an input coupler that has a waveguide core material with lower optical index.




Figure 7-1: (a) Plan view of one of the grating couplers superimposed on the trap electrode metal (partial). The single-mode waveguide is tapered to provide a wider beam at the grating coupler. The grating teeth are formed via a partial etch into the waveguide material beneath a window in the electrode metal. The gratings used in this work are forward-emitting, designed to emit a beam to the ion location as depicted, but 55 μm above the surface of the metal. (b) Photonic ion-trap cross section (not to scale), showing the integrated waveguides beneath the surface-trap electrodes and light diffracting out of the chip by means of a grating coupler. (c) Side view of light coming out of a grating [MBM⁺16]

7.1.2 Grating Couplers

Light can be routed around on a chip within waveguides; to couple light back out of the plane of the chip, we use diffractive grating couplers [DW75]. At the point where we want to couple light out of the chip, the width of the waveguide is tapered up to 18 μm in order to expand the spatial mode of the light. The gratings are created by etching a periodic pattern into a widened waveguide along the direction of propagation (Figure 7-1(a)). This creates a periodic variation of the effective refractive index, causing light to diffract out of the plane of the chip at an angle dependent on the wavelength, grating period, and effective index. To increase the light intensity at the ion location, the grating teeth are curved, focusing the beam to a few-micron waist in the direction parallel to the chip surface and perpendicular to the grating emission.

The gratings presented here are designed with a uniform period so that they do not focus in the vertical transverse direction (the direction perpendicular to the beam and closest to z) and maintain $\sim 11\ \mu\text{m}$ beam diameters ($1/e^2$ intensity) at the ion location. However, the gratings are designed with a slight curvature to focus the beams in the horizontal transverse direction (the direction perpendicular to the beam and parallel to the chip surface) to $\sim 4\text{--}7\ \mu\text{m}$ beam diameters.

In the most of experiments described here, the gratings are all designed to emit light in the forward direction, at angles around 45° from normal. Forward emitting gratings generally make the layout of waveguides simpler, since waveguide crossing can be avoided. Shallow emission angles also simplify routing since gratings can be placed further away from the ion location, which helps avoid crowding close to the middle of the trap. The tradeoff for using gratings like this is the presence of higher-order diffraction modes in the output, which results in several other spurious beams coming out of the grating. To mitigate this effect, a partial etch of the grating teeth was implemented during fabrication, which reduces the index contrast and thus reduces the efficiency into all diffraction orders. For the designs in this work, the grating efficiencies are approximately 10%; along with OCC loss, the grating is the other primary source of total optical loss in our process.

The first round of gratings we used in this work were designed using an index of refraction based on low pressure chemical vapor deposition (LPCVD) SiN. However, the devices were actually fabricated using plasma-enhanced chemical vapor deposition (PECVD), which was chosen for two reasons. First, in the time between grating design and the fabrication of the chips in this work, a PECVD SiN with substantially lower propagation loss than LPCVD SiN at wavelengths below 460 nm was developed. Second, the PECVD SiN can be deposited at temperatures compatible with pre-existing metal layers, which allows us to place a continuous metal ground plane between the photonics and the Si substrate. The index of refraction for PECVD SiN is different from that for LPCVD, which means that the emission angle from the grating differed from the expectation. In current designs, the tolerance on the index of refraction has been identified as the primary property that must be controlled during fabrication in order for our devices to hit our target location.

Grating couplers can be characterized by directly profiling the emitted beam with a high numerical aperture (NA) microscope objective¹ and projecting the beam onto a CMOS

¹The profiler consists of an infinity-corrected 0.9 NA, 60 $\times$ (magnification) microscope objective and a 1X




Figure 7-2: (a) The fibers are arranged in a fiber-array block which is aligned to photonic waveguides within the chip. The waveguides route the light to the chip center where vertical grating couplers diffract the light toward the ion trapped above the surface of the electrodes (the paths depicted are only notional). (b) Scanning-electron micrograph (SEM) of the center of the trap showing square windows in the electrodes for the underlying grating couplers. Inset: SEM of a grating coupler during fabrication showing the curvature of the grating lines, leading to transverse beam focusing. (c) 3D Reconstruction of the different lasers diffracting out of the surface of the chip. The beam profiles are based on collected data, but the colors are false.

detector [MBM⁺16, MR17]. The object plane of the imaging system is stepped vertically away from the chip surface while images are taken. The resulting data of 2D optical intensity versus height can be used to reconstruct a 3D map of the beam coming out of the grating. The beam profile can be registered to the chip surface using alignment features on the surface, photographed with the same microscope setup, and the results can be compared with simulation.

7.1.3 Combining into an Ion Trap

In this work, we follow on from an experiment in which a single wavelength of light was coupled onto a chip with waveguides and grating couplers [MBM⁺16] and extend to an ion trap which has six wavelengths of light which are all coupled onto the chip directly out of optical fiber. The six lasers needed for ionization, cooling and coherent control of strontium are coupled onto the edge of a chip using a glass array of optical fibers (Figure 7-2a). On the chip, waveguides route lasers from the edge towards the middle where a trap is patterned in the top metal layer with a geometry similar to our standard trap² (see section 3.6.1). Near the ion's location, there are just four grating couplers, so two of the optical channels must

tube lens, followed by a CMOS sensor based camera located at the image plane of the lens system.

²The primary difference with this particular geometry is the lack of segmentation in the thin electrode which runs between the two branches of the RF electrode. The trap height is also a little higher, at around $z = 55 \mu\text{m}$.

carry two separate wavelengths of light (see Figure 7-2b-c). We combine the 1033 nm and 1092 nm into one of the channels and the 405 nm and the 422 nm into another one of the channels. The angle of diffraction will be different for the two wavelengths in the combined channels, but since the gratings only focus in the transverse direction (the direction parallel to the plane of the chip and perpendicular to the emission direction), we find the beams are sufficiently large at the ion such that the discrepancy in pointing does not significantly affect operation.

The four grating couplers are placed around the ion location in the cardinal directions of the trap geometry. The 422 nm Doppler cooling beam must have a projection along the ion's mode of motion, which requires this light to come through one of the gratings along the center axis of the trap. Since we do many qubit operations using the axial mode of motion, we choose to use the other axial grating coupler for the 674 nm. The 461 nm, which does the first step of photoionization, works best when the propagation direction is orthogonal to the direction of the neutral atoms; for our typical trap orientation, this works out naturally when the beam is sent in through one of the radial gratings. The direction of the two infrared beams does not matter, so they are simply placed on the remaining radial grating. The polarization of the light in the waveguide, predominantly in the transverse direction, is maintained in the grating, and thus the light which comes out of the chip is polarized parallel to the plane of the chip. For our optical qubit, polarization is most important in the 674 nm, the 1033 nm and the 1092 nm. For the two infrared beams, we desire a polarization which is at least partially perpendicular to the quantization axis (defined by the applied magnetic field which is normal to the trap), so this condition is satisfied. For the 674 nm the polarization, quantization axis, and the propagation direction determine the allowed transitions and set the respective Rabi frequencies (see section 2.5.1).

These traps are fabricated using the PECVD process described above. Beneath the photonics layer, a 0.5 μm -thick sputtered niobium metal layer is deposited and patterned via optical lithography to form a ground plane for the ion trap [MEB⁺14]. The photonics are defined in a 100 nm layer of SiN which is patterned with optical lithography and sandwiched between two 5 μm thick layers of SiO₂. The trap electrodes are defined in a 1 μm thick layer of sputtered niobium; above the gratings, four $20 \times 20 \mu\text{m}^2$ apertures are opened to allow light to propagate out of the chip. Above these windows, we deposit and pattern a 20 nm thick film of the transparent, conductive material indium tin oxide (ITO) to reduce the

exposed area of dielectric that could potentially become charged [WHL⁺11] and adversely affect the electric field at the ion location (see section 4.1.3).

The total optical loss from fiber to ion depends on the channel. For the infrared channel, the combined loss of the coupling onto the chip, the propagation and the diffraction efficiency is about 17 dB of loss, and for the other channels, the combined loss is about 22 dB. As mentioned previously, a significant portion of the loss occurs in the diffraction grating, which is intentionally defined with only a partial etch of the SiN. Another source of effective power loss comes from beam pointing error. As mentioned previously, there was a known difference between optical index used in simulation and the one achieved in fabrication. This led to the beams propagating a steeper angle than intended, such that the beams intersect each other at a height of $65\,\mu\text{m}$ above the chip surface, $10\,\mu\text{m}$ above the ion’s position. Accounting for all of these loss mechanisms, and using optical profiles of the light coming out of the gratings, we can predict the optical intensity of the various wavelengths of light at the ion’s location as it is translated along the axial direction. Despite the losses, we still achieve beam intensities at the ion location that are comparable to those obtained when using free space beam paths with similar optical power, due to the fact that the cross-sectional area of the focused beam is about 100 times smaller than in the free-space case. With improved waveguide materials [WLK⁺19], alternative packaging techniques [MZM⁺20], and improved grating designs [MR17, MZM⁺20] it should be possible to achieve even higher beam intensity at the ion.

7.2 Installing a Photonics Chip into a Cryogenic Apparatus

Translating the finely-aligned array of optical fiber from the bench to the chamber requires some engineering in our packaging procedure. Outside of the vacuum chamber, we use a six-axis positioning stage in combination with a custom-machined vacuum chuck to precisely align the fiber array to the edge of the chip (Figure 7-3b). To obtain optimal coupling, the fiber array must be aligned to within about $1\,\mu\text{m}$. Any deviation from this optimal coupling point will lead to excess loss between the fiber and the grating coupler. The method for fixing this alignment will be described in the following subsections.

Light is coupled into the fiber array via standard bayonet connections on the other ends of the fibers. Avoiding fiber splices increases the ease with which traps can be installed,




Figure 7-3: (a) Light is coupled to the integrated photonic trap chip via optical fibers which enter the cryogenic vacuum chamber through a fiber feed-through. In vacuum, the fibers are coupled to the chip with the use of a fiber array block. (b) Six-axis positioning stage used to hold and align fiber block array near the edge of the chip before and during the epoxy application. The coupling efficiency can be monitored using a power meter to measure the power transmitted through the loop-back or propagating out of the grating. (c) Photonic ion-trap chip packaged and mounted with strain relief. (d) A later version of the strain relief bracket, including a 90° bend to help improve the ease with which fibers can be routed up and out of the 50 K shield.

but comes at a cost of additional loss at the fiber-to-fiber interface. Inside the vacuum chamber, we use screw-on fiber-fiber connections when installing the packaged chip. These connections are done with standard threaded sleeve couplers which we mount in the chamber around the outside of the the 50 K shield. The fibers from the fiber array are routed outside of the 50 K shield during trap installation. On the other end of the couplers, laser light is routed into the chamber via a custom fiber feed-through. The feedthrough is constructed by drilling holes into a conflat blank and then threading a cleaved fiber for each of the optical channels through the holes. After threading the fibers, they are then re-spliced and the holes are sealed using Torr Seal, which we find to be a sufficiently vacuum-tight bond.

When the chip is installed with the fiber feedthrough, we observe an additional 3 dB of loss arising from fiber splices and fiber-to-fiber connectors. Fibers in the vacuum system are all unjacketed to reduce their overall size and make routing easier, but they all have a polymer coating and plastic housings around the connectors. With the whole fiber assembly present, we still observe lifetimes on the order of a few tens of minutes³.

Since most of our lasers are coupled into fiber before being routed to the final free-space positioning stage, switching from external optics to integrated optics is as simple as moving a fiber connection. The polarization of light into each fiber is controlled using half-wave plates in the free-space path. We record the input power into each fiber and, when possible, measure the light diffracted out of the chip, through viewports in our chamber.

7.2.1 Attachment Procedure

To prepare a chip for attachment, the facet on which the waveguides terminate is first optically polished. During dicing and polishing, the top surface of the chip is protected with photoresist to prevent damage to the trap electrodes. We clean and bond chips to the CPGA package as we do with our other chips, but we use three additional trap-sized pieces of silicon to stand the chip off by an additional 1.75 mm further than the usual distance; this extra height ensures that the waveguides are sufficiently proud of the CPGA that the fiber array does not collide with the CPGA during alignment.

The packaged chip is mounted on a translation stage to do coarse alignment while the fiber array is brought in on the vacuum chuck which was previously mentioned. During this procedure, we use an additional “loop-back” waveguide designed for light at 1033 nm to monitor coupling efficiency into the chip as we align the fiber array to the waveguide inputs. The loop-back makes a loop around the rest of the photonic structures, connecting the outer two waveguides together. We send light into one of these waveguides and monitor the output power at the other; when the coupling is optimized, the position and rotation of the fiber block is set and all of the fiber channels in between the ends are also , to within a precision set by the fabrication of the fiber array and the concentricity of the cores of the fibers. We find that the optimum chip-to-array distance varies depending on the wavelength, so the total coupling of the 674 nm (typically the most power-starved wavelength) is checked and

³This lifetime is lower than the several hours we would expect for a standard configuration, but to install the fiber-packaged chip, we must also use a modified 4 K shield (see section 3.1) which does not fully surround the trap, thereby exposing ions to a greater solid angle of higher-temperature surfaces and gases than usual.

optimized before securing the alignment into place.

To secure the alignment, we use two stages of epoxy. The first stage is a UV-curable electronic adhesive (Norland NEA123) which starts out as a viscous gel and then cures to a firm hold with some slight give. We apply just a small dab of the adhesive to either side of the fiber array block and then shine some UV light in order to hold the alignment in place⁴. It takes less than a minute of UV exposure for the adhesive to begin curing. With the alignment then temporarily secured by the first adhesive, we then apply a two-part cryogenically-serviceable epoxy (Masterbond EP21TCHT-1) which starts out as a paste and then cures over 48 hours to a hard ceramic material. We apply enough of the two-part epoxy to completely encase the first layer of glue, and we ensure that the second epoxy has plenty of contact to both the chip and the fiber array.

In some commercial optical packaging procedures, a thin epoxy is applied to the interface between the components to be joined. We had initially attempted this with a common optical adhesive (Norland NOA81), but we found that the epoxy underwent a process known as "yellowing" when we attempted to pass violet light (at 422 nm) through the glue joint, at which point the cured epoxy became opaque to the input laser. This is a known limitation in many UV-curable glues. In our case, the issue is exacerbated by the high optical intensity at the core of the fiber where the joint is made.

The cryogenic epoxy we use was chosen for cryogenic compatibility, relatively low out-gassing and low thermal expansion. Despite the low coefficient of thermal expansion (CTE), we still observe some shift in the alignment of the fiber array during cool down. In most of the results discussed here, there was an additional loss of about 7 dB during the cool-down. To better match CTEs, in later iterations of our process, we have employed a fiber array made from silicon, rather than glass, to match the material in the chip stack. We glue both the silicon fiber array and the chip to the same dummy silicon wafer, which we then mount in our CPGA. Better matching of materials in the stack has resulted in coupling losses less than 3 dB during cool-down.

⁴We had initially used this adhesive as the only bond, but we discovered that the cured material would tend to crack after thermal cycling.

7.2.2 Strain Relief

In the first few cool-down attempts with the fully packaged chip, we observed failures in which the epoxy would cleanly break free of the chip surface, causing the fiber array to separate. As mentioned above, some improvements were made to the epoxy used to make the attachment, but we also created a support structure to relieve some stress at the fiber block interface. Our primary concern was to avoid a torque force which could be exerted on a long lever arm if the fibers were pushed laterally during chip installation. We also wanted to allow the fibers to have some room to contract and shift slightly during the cool-down, so we chose to avoid a solution in which the fibers were glued down at another position on the CPGA or PCB octagon.

In our most recent successful cool-downs, we use a bracket assembly (Figure 7-3c-d) which is machined from PEEK⁵. The bracket clamps to either side of the CPGA, hugging the top and bottom edge of the ceramic package, which fixes it in place. Since the bracket gets so close to the metal pins of the CPGA, it cannot be made of a conductive material. On the end of the chip with the fiber array, an arm extends from the package with a small channel to support the fibers as they run away from the chip. The fibers are secured in place with an aluminum top plate, which holds the fibers down, but leaves enough room for the fibers to be able to slide back and forth during expansion and contraction. Around the sides of the CPGA, grooves are cut into the material to allow laser (and atom) beam access.

The latest design of the bracket also bends the fibers at a 90° angle as they exit the chip, such that they end up oriented vertically in the chamber. The design supports the bend that the fibers are required to make as they route out of the vacuum system, and it also makes installation into the vacuum system easier.

7.3 Technical Challenges

Even after getting chips with photonics aligned and installed in the system, there were still a few issues encountered when doing operations with ions. The primary non-ideal behavior that we observed was a heating rate of 640 ± 200 quanta/s with a trap frequency

⁵Polyetheretherketone (PEEK) is a chemically-resistant plastic material which is known to have low outgassing properties that are suitable for installation in UHV systems. We chose to use a UHV compatible material since the bracket would be so close to the ion's position.

of 1.4 MHz [NSSA⁺20]. As discussed in section 4.2, this level of heating is higher than we would naturally expect for one of our standard materials at a temperature of ~ 6 K, which would be about 10 quanta/s for a strontium trap at this frequency.

7.3.1 Trap Temperature

In early experiments, we began to suspect that the temperature of the chip might have been increasing due to the applied RF potential. We observed that as we increased the RF power, the resonance frequency changed by significantly more than the typical few tens of kHz shift. To check the trap temperature *in situ*, we used a temperature sensing resistor installed on the chip surface, as in section 5.4. For typical RF powers required for trapping, we found that the trap temperature increased to around 24 K. According to the data in Figure 4-1, this temperature is likely not enough to explain the measured heating rate, if the anomalous heating was dominated by the effect of the niobium metal⁶.

The reason for the increased temperature of this chip is probably due to the higher loss of the SiO₂ dielectric below the RF electrode over the standard sapphire dielectric, which has very low loss. We do not observe a significantly lower quality factor in the RF resonator. Thus the added loss does not impact the performance of the RF circuit even though it locally causes the chip to heat up due to the added power dissipation. To remedy this problem, we made a few attempts at better heat sinking from the chip to the cold head, but this was tricky due to the several layers of interconnects from the interposers and chip carriers. It should be possible, with more direct thermal links between the chip and cold head [Tod20], to obtain more cooling power.

7.3.2 Charging and Technical Noise

As noted in section 4.1, it is possible that our control electronics can introduce heating. In this trap, we were particularly cognizant of technical noise given the noted higher-than-usual heating rate. Given the various materials and processing steps, there was some concern about dielectric charging, even with the ITO. However, after many days of operation with several iterations of these chips, we did not observe any of the usual hallmarks of laser-induced charging, like trap frequency variation or changes in compensation fields.

⁶ Assuming, of course, that the niobium on the photonics chip behaves in the same way as the niobium on our standard single-metal-layer on sapphire chips. In reality, there could be effects from other, unknown materials left over on the trap surface from fabrication steps.

Even so, over the many photonics chips that were installed, we observed that some chips required very large compensation fields to even get suitable potentials for trapping ions. One difficulty with these chips is in the modeling of the unit potentials. Since part of the ITO covers the RF electrode, it will carry part of the voltage which creates the pseudopotential. Unlike the metal in the rest of the electrode, the ITO is a resistive conductor, so it may not be valid to model these patches with the same boundary conditions as the rest of the surfaces, which are modeled with zero resistance. A full time-dependent model of the fields above the trap which takes into account the ways in which charge flows along the electrodes may be required to accurately capture the potential seen by the ion.

Even without a perfect model of the trap, we can still employ the noise mitigation techniques discussed in section 4.1.2, including pseudopotential gradient minimization and filtering of the phase noise on the RF source. We also use very low bandwidth versions of the control voltage filters, since most characterization took place without requiring ion motion. To the best of our ability to mitigate noise, we found that the excess heating was not due to any aspect of our voltage control system.

7.3.3 Surface Treatment

Several of the materials in the packaging and in the fabrication (like the ITO) had not been carefully examined in our experiments. We also had reason to believe that the ITO had something to do with the high heating rate since we had observed especially high heating with the ions positioned just above the windows in a few cases. We performed a few experiments with more simple configurations to try to isolate the effects of the different materials. Without the fiberblock packaging, we found that some chips had heating rates as low as 100 quanta/s, though we did not collect enough statistics to know for sure if chip-to-chip variation or packaging variation was the dominant factor in increasing the final observed heating rate. More recent iterations of the photonics chip have higher "baseline" heating rates when tested without packaging, which points to a difference in chip fabrication. For these newer chips, we have found that ion milling (see section 4.2.1) can reduce heating rates back down to the ~ 100 quanta/s level, which is further evidence of contamination on the surface of the chip (presumably left over from a fabrication step) leading to the increased heating.

Milling a chip with fiber packaging is a challenge in itself. Either the chip must be put

into the ion mill chamber with the fibers attached or the chip must be milled before the attachment. Without doing some engineering to find a way to safely install and remove the fiber-attached chip into the mill, we had to live with about a full day of air exposure after milling while the epoxy from the attachment set before the chip could be installed in the chamber. This additional re-exposure after milling is undesirable since surface contaminants can enter in while the chip is packaged.

We were unable to get a chip prepared which reached the baseline ~ 100 quanta/s heating rate with all of the photonics installed. There remain many open questions as to which processing steps and materials might most play a role in the excess heating and to what extent the geometry of the windows affects the observed heating which were beyond the scope of this work.

7.4 Results

After getting the photonics chip installed in the system, we were able to perform characterizations to compare the behavior with external optics to that with integrated photonics. We were able to use ion measurements to match data recorded on the bench, and we also observed the intrinsic stability afforded by using light sources registered to the ion trap.

7.4.1 Single Qubit Operations

We characterized the profiles, as well as the positions of the beams relative to the trap electrodes, *in situ* using the ion by measuring the strength of the laser-ion interaction as a function of ion position. The location of the ion was varied by changing the voltages applied to the DC electrodes (see section 2.3) to perform rudimentary ion shuttling along the axial direction of the trap. The position of the ion is advanced in single μm steps, and a laser interaction is performed at each position. We were also able to use each beam path to demonstrate key ion-control operations. Most characterization was performed with one beam at a time, due to the noted optical index discrepancy which resulted in the point of intersection of all of the beams being vertically displaced from the RF null line.

We profiled the 674 nm integrated beam by measuring the frequency of Rabi oscillations on the qubit transition as a function of the ion's position along the trap axis (see Fig. 7-4a). The ion-interaction profile resolves some of the non-Gaussian beam structure evident in the




Figure 7-4: Beam profiles measured *in situ*. The profiles shown are measured with different wavelengths: (a) 674 nm, (b) 1092 nm, (c) 422 nm, and (d) 1033 nm. In each case, the ion position is adjusted in the $\hat{y}$ direction by calculating and applying appropriate control voltages, and at each position, the intensity is measure via ion-laser interactions described in the test. While both the 422 nm and 674 nm beams are focused transversely along $\hat{x}$, profiling with the ion along $\hat{y}$ measures the unfocused cross-section. We note that all the ion interaction profiles agree very well with the beam profiles measured using the microscope objective (shown in solid lines), which gives us confidence that the integrated photonic components can be accurately characterized, independent of the ion, using conventional optical techniques and equipment. In addition, this implies that there are not significant effects on the integrated optics from cryogenic or UHV operation.

line cut of the laser intensity from images taken with the microscope *ex situ* and shows a beam diameter ($1/e^2$ intensity) of $13 \pm 0.7 \mu\text{m}$ along the axial direction, which is offset to the right of the trap center ($y = 0$) at $y = 13 \mu\text{m}$ (see Fig. 7-4a).

We used the integrated 674 nm beam path to perform spectroscopy on the qubit transition, as well as quantum state preparation (including OP and RSC to $\bar{n} < 1$ motional quanta in the axial mode), Rabi oscillations, and Ramsey interferometry. With 10 mW of

optical power coupled into the fiber attached to the chip, we were able to perform π pulses on the optical qubit in $6.5 \mu\text{s}$. This measured pulse time agrees with the value we obtain through a first-principles calculation of the Rabi frequency [Roo00] given the measured beam dimensions, the beam power input to the chip, and the total optical loss of 31.4 dB in this channel.

The 422 nm integrated beam was profiled by measuring the fluorescence of an ion illuminated by 422 nm light emitted from the grating and a free-space 1092 nm beam. In these measurements, the beam intensities were kept below saturation so that the ion fluorescence rate was approximately proportional to both the 422 nm and 1092 nm intensity. We measure a beam diameter of $8.5 \pm 1 \mu\text{m}$ along the axial direction, offset to the left of the trap center at $y = -11 \mu\text{m}$ (Fig. 7-4c). With the integrated 422 nm beam path, we were able to demonstrate ion detection with 99.6% fidelity in 3 ms [NSSA⁺20]. This measurement was limited by the effect of background scatter from the grating and waveguide which was not focused onto the ion and instead made its way into our imaging system. We were able to overcome some of this scatter with spatial filtering by inserting beam blocks into the imaging path to more-selectively pass photons from the ion.

The 1092 nm integrated beam was profiled by measuring ion fluorescence while illuminated by a free-space 422 nm beam (again using intensities below saturation) and shows that the beam focuses along the axial direction ($\hat{y}$) to a beam diameter of $5.5 \pm 0.25 \mu\text{m}$ (Fig. 7-4b). We demonstrated repumping from the $D_{3/2}$ state via the integrated 1092 nm beam with a sufficient amount of power for Doppler cooling and detection.

The infrared integrated beam path can also be profiled using using 1033 nm light. First we prepared the ion in the upper (dark) state by applying a π pulse with an external 674 nm beam. Then we applied a short quench pulse of 1033 nm light via the integrated beam path to partially quench the dark state population. This increased the probability for the ion to be in the lower qubit (bright) state but maintained finite dark state probability (to avoid saturation). Finally, we measured the probability that the ion was in the bright state at various positions along $\hat{y}$ and obtained a beam diameter of $6.7 \pm 0.3 \mu\text{m}$ (see Fig. 7-4d). With light from the integrated 1033 nm pathway, we were able to quench the $D_{5/2}$ state in less than $10 \mu\text{s}$.

There was not an easy way to measure the intensity of the 461 nm or 405 nm integrated pathways directly. The 405 nm could however be approximately characterized using a nearby

transition at 408nm [NSSA+20]. By separately adjusting the ion position to maximize the intensity of either photoionization beam, we were able to demonstrate ionization using the integrated beam paths. Ion loading was achieved within a few seconds using the integrated 461 nm beam and a free-space 405 nm beam, and in approximately 1 minute via the integrated 405 nm beam and free-space 461 nm beam. The longer loading time of the latter results from the broad 405 nm transition being unsaturated and from a lower 405 nm beam intensity relative to that available when using the free space beam.

Even though we were not able to position the ion at the location where all of the beams intersect, we were still able to demonstrate a few operations involving multiple beams, by positioning the the ion near the intersection of two or more beams, although this effectively reduced the available laser intensity for various operations. Doppler cooling and state detection with the integrated 422 nm and 1092 nm beams were performed, enabling trap lifetimes exceeding multiple hours and measurement of the ion state without the use of free-space beams; we observed ion detection with 99.0% fidelity in 5 ms. We also used the integrated 674 nm, 1033 nm, and 1092 nm beams simultaneously to perform qubit state preparation via OP and spectroscopy of the $S_{1/2} \rightarrow D_{5/2}$ carrier transition and its first order motional sidebands.

7.4.2 Improved Coherence with Integrated Photonics

The inherent stability of optical paths integrated into the trap chip can provide vibration tolerance to trapped ion qubits, clocks, and sensors. Effective qubit decoherence due to optical phase variation and amplitude modulation arising from vibration of the ion with respect to the laser beam (Fig. 7-5a) should be significantly reduced, as the vibrations of the ion and the light delivery optics are common-mode due to monolithic integration with the trap (Fig. 7-5b). To test this hypothesis, vibrations of varying amplitude were intentionally introduced, and we measured the effect on trapped-ion qubit coherence.

By clamping the vibrating cryocooler head to the trap mount via the upper portion of the vacuum chamber, we can incrementally increase the vibrational coupling, causing the chip and ion to oscillate in space with a significant amplitude, as shown in Fig. 7-5c-d. Here images of the ion fluorescence are observed as a function of time to determine the approximate induced vibration amplitude.

With the cryocooler head mechanically isolated from the trap, a $1/e$ Ramsey contrast




Figure 7-5: (a) Vibrations introduced from the cryocooler attached to the vacuum chamber cause the ion to vibrate relative to the free space optical path. (b) Integrated photonic beams emitted from the chip will vibrate in common with the ion. (c-d) Long-exposure images of increasing vibrational coupling, collected using a high-NA lens and an electron-multiplying CCD camera. Increasing the mechanical coupling increases the amplitude of ion oscillation, and hence ion acceleration for fixed oscillation period (we observe an oscillation period of approximately 15 ms). (e) The Ramsey-contrast decay time ($1/e$) measured via the free-space external beam path (red circles) decreases rapidly as the acceleration experienced by the vibrating ion increases. However, the coherence time measured when using delivery via the integrated beam path (blue diamonds) is unchanged. Lines are exponential fits to the data in both cases and error bars indicate the standard error of the mean.

decay time of approximately 600 μs is measured using either the free-space 674 nm beam path or the grating-coupler-delivered 674 nm beam (see Fig. 7-5e, points at zero acceleration). The contrast decay in these cases is limited by a combination of magnetic-field noise and uncompensated acoustic/thermal noise in the ~ 30 m of fiber used to transmit the light from the laser to the vacuum chamber.

Increasing the vibration coupling between the cryocooler and the trap mount as described above has a strong effect on the measured Ramsey-contrast decay when using the free-space qubit-control beam. As shown in Fig. 7-5e, the Ramsey decay time as a function of the ion acceleration (as extracted from the ion motion, see [NSSA⁺20] for more discussion) drops rapidly with increasing ion motion due to the vibration-induced randomization of the laser optical phase between the Ramsey pulses. In contrast, the Ramsey decay measured using the grating-coupler-delivered beam remains unchanged as a function of ion acceleration. At this coherence level, the integrated beam path renders the ion-laser interaction immune to even very strong vibrations and the corresponding time-varying Doppler shifts; the peak Doppler shift corresponding to the highest acceleration measured here is ~ 9 kHz. This suggests that coherence limited by such perturbation, for instance in a fieldable sensor




Figure 7-6: (a) Illustration of a trap with two zones with diffractive grating couplers. (b) Example experiment of controllable addressing using ion motion. We execute a Rabi experiment with the global channel and move one of the ions (here the left one) out of the beam in order to hide it from the effect of the qubit laser. From a fit to equation 4.3, we obtain a Rabi frequency of 180 kHz for the right ion and a frequency <1 kHz for the left ion.

or clock platform, may be improved by using integrated photonics for quantum control of trapped-atomic systems.

7.5 Multiplexed Operations in Two Zones

Recently, we have performed experiments with a new generation of photonics chips, which have a few interesting geometric features. On one version of the chip, depicted in Figure 7-6a, there are two zones, at $y = -120 \mu\text{m}$ and $y = +120 \mu\text{m}$ which both have grating couplers for the 674 nm; each zone has an independent, "local" channel, and both zones share a "global" channel, which is derived from a single input waveguide followed by a 3 dB power splitter. This arrangement allows for some tests of selective addressability.

After characterizing the behavior of the laser interaction at both zones independently with a single ion, we then moved on to experiments with two ions—one in each of the zones. Since we currently only have the ability to read out the state of the ion in a single zone⁷, we used ion motion sequences to move ions to the middle of the trap to perform detection and cooling and then moved the ions back out to the photonic interaction zones to perform qubit pulses. In all cases, the ions were kept in two separate harmonic wells; these wells needed to stay separated by about $240 \mu\text{m}$ in order for the voltages at one well to not interfere with the other. Ions are continuously juggled around as experiments take place; the bright

⁷This is due to the fact that we have just one PMT which collects photons from a single zone. It is possible to achieve more flexible multi-zone readout with either a PMT array or by quickly reading out frames from the EMCCD.

probability of the left and right ion are tracked independently.

To demonstrate controlled addressing, we did an experiment in which the ion in one zone was held firmly within the path of the light coming out of the grating coupler and the other ion was held slightly outside of the beam. We then applied a laser pulse to the global channel, followed by detection of the state of each ion. From the Rabi oscillations in Figure 7-6b, we can see that the ion in the beam is flopping while the other ion is unaffected by the laser field. The ratio of the observed Rabi frequency is bounded on the low end by about 180:1. Since the optical intensity is proportional to the square of the Rabi frequency, this corresponds to a reduction in intensity at the ion position of more than 4×10^4 . This high degree of isolation, provided by motion control of the ions, opens up various possibilities for controlling and multiplexing qubit operations applied to global channels. Such experiments could be controlled using combined voltage sources like those discussed in chapter 5, and the motion control techniques discussed in the next chapter can help to increase the speed of these motion-mediated operations.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 8

Ion Motion and Planar Trap Designs for Array Architectures

*“Up late for six nights in a row
Figures in the dark, oh, they move so slow”*

—Snail Mail

So far we have looked at a few different technologies which can all have potential benefits to future complex arrays of ion traps. In order to transmit quantum information around in such a device, the ions themselves will eventually need to be moved around the trap. As we have seen again and again throughout this thesis, ions tend to be persnickety sensors in the presence of noise or non-ideal control. Significant care is required to control without destroying carefully prepared states of motion. In this chapter, we present some control techniques and trap designs which have been implemented to improve ion motion operations.

In section 8.1, we discuss the generation of voltage waveforms that are used to control the position of ion(s), and we consider the ways in which these motion operations can be optimized and applied with high fidelity. Control of the ions’ potential can also be improved by making optimizations in the control electrode geometry, so we present a trap design which was optimized for splitting chains of ions in section 8.2. Lastly, in section 8.3, we overview some initial results obtained with the optimized trap design.

8.1 Generating Voltage Waveforms

As discussed in section 2.3, it is possible to control many aspects of the potential seen by the ion with the manipulation of the voltages applied to the trap electrodes. Since we have so many electrodes, typically 10 or more, which have an effect, it can be possible to control many terms in the potential. By dynamically adjusting the control voltages, we can change the potential in order to manipulate the positions and order of chains of ions [RBKD⁺02]. The voltage signals applied to the control electrodes are typically called "voltage waveforms."

In order to preserve the carefully cooled motional states, we generally want to perform these operations in ways that do not excite vibration in the modes we use for quantum information. One way to achieve this is to make all changes very slowly, in the so-called *adiabatic* regime. Changes are considered adiabatic when they occur on timescales which are much longer than the period of the ion's motion (i.e. much greater than $1\ \mu\text{s}$, on the order of a few hundred μs); the ion tends to remain in the initial motional state during operations on these timescales. In few-ion experiments, long time scales may be acceptable, but in larger scale experiments, which will involve motion operations interleaved with quantum gates, the total time spent in transit can add up to a significant part of the total time. Ideally, we desire motion operations that are on the same time scale as the gates, on the order of a few μs . Ion motion operations at this speed are higher are said to be in the *diabatic* regime.

Devising and applying voltage waveforms that achieve this level of performance requires a combination of theory, engineering and measurement. First we consider what it takes to make an optimal trajectory, and then we consider how we can go about realizing this in the experiment.

8.1.1 Ion Motion Primitives

The ion motion requirements in a large array can be decomposed into a few basic desiderata. These include the ability to transport single ions or small chains of ions from one location to another, to join ions into a chain¹, and to reorder a chain of ions. Usually, operations designed for one or two ions can be performed on groups of ions (i.e. moving a small crystal of a few ions). In the work presented here, we focus only on operations involving ions of a

¹The waveform to join ions into a chain can also be played in reverse in order to split a chain into discrete ions.

single atomic species. Generally, ion chains with mixed masses experience dynamics that are more difficult to model analytically. While the methods presented here may someday be applicable in mixed species experiments, they are beyond the scope of the present work.

The most simple operation is transport, in which we want to move the position of the potential minimum. We take the direction of motion to be the y direction, along the minimum of the pseudopotential. The potential in the axial direction is:

$$\phi(y, t) = \frac{1}{2} \frac{m}{q} \omega^2 (y - q(t))^2 \quad (8.1)$$

where $q(t)$ is the trajectory of the potential minimum. The functional form will be discussed further in the next section.

During a split, we consider the fourth-order term in the expansion of the potential:

$$\phi(y, t) = \alpha(t)y^2 + \beta(t)y^4 \quad (8.2)$$

Typically, we operate with a purely harmonic potential in which $\alpha > 0$ and $\beta = 0$. In order to separate a chain of ions, we want to invert the sign of α such that the ions experience a repulsive force at the center of the trap. We can maintain confinement by ramping up the value of β .

Reordering chains of ions can be done in a few different ways². One method is to "swap in place" by rotating a pair of ions about one of the radial axes such that the ions change position along the trap axis. This can be achieved by locally increasing the strength of the confinement in the axial direction while simultaneously introducing a force in the radial direction to encourage the rotation. Rotations like this have been recently demonstrated with low heat gain and speed approaching the adiabatic limit [KRS⁺17, vMMG⁺20, WJE⁺20, PDF⁺21]. Another way to reorder a chain of ions is to move through a junction, alternately allowing ions to pass through or turn a corner at the intersection. But motion near a junction is tricky. The change in direction requires a projection of the RF potential along all three dimensions, so micromotion becomes difficult to avoid and ions are sensitive to noise coupling during the transit (see section 4.1.2). Motion in junctions has been explored in both 3D [Bla12, HOS⁺06] and planar traps [WAF⁺13, SVB⁺14]. Ion reorder-

²It's worth pointing out that it is also possible to reorder qubit states in chains of ions through the use of teleportation, but here we focus only on techniques which involve moving ions via external fields.




Figure 8-1: Plot of the ion position versus time for a few different protocols [HYO⁺08]. Here we see a linear, sinusoidal and two hyperbolic tangent (parameterized by a term N which controls the slope at the inflection point) functional forms; see [HYO⁺08] for the precise functions.

ing will not be explored in this work, though motion in junctions is an obvious extension [BOV⁺09, BOV⁺11] (discussed further in section 9.1.1).

8.1.2 Determining Optimal Trajectories

Before we consider what we mean by the distinction of an "optimal" trajectory, we first consider what it might take for a trajectory to be "pretty good." In the case of linear transport (see Figure 8-1 for a depiction of the adiabatic transport discussed here), the most straightforward $q(t)$ we can employ is a linear ramp, with constant velocity, from the start to the end point. The linear protocol has the advantage of having no acceleration during transit, but the two discontinuous momentum kicks at the start and end of the sequence can be problematic. Increasing the complexity, we might then seek to smooth out the sharp corners at the ends of the transport by moving to a $\sin^2$ or $\tanh$ shape for $q(t)$. These styles both have the advantage of having the starting and ending impulse spread out over time. These transport schemes can all be considered analytically [RLB⁺06, HYO⁺08], and for transport times on the order of a few trap periods, the heat gain is well below one quanta of axial motion.

These transport protocols are still in the adiabatic regime. One way to push into the realm of *diabatic* transport is through well-timed kicks of momentum. In one approach, a large coherent state is created by switching the trap voltages suddenly to a different configuration which puts the ion in a highly displaced position of a trap at the midpoint between the ends of the desired transport [ALS⁺16]. After half a period of oscillation in this intermediate trap, the ion will be at the endpoint. Here, the voltages are quickly switched again to a well centered on the endpoint. This can be achieved by quickly switching between fixed potentials. More generally, any large coherent state created from an initial kick of momentum can in principle be undone with a subsequent kick which serves to stop the ion at the endpoint while also undoing the initial coherent displacement [WZR⁺12]. These procedures are typically referred to as "bang bang" transport.

A more general approach to cold, fast transport is called shortcut to adiabaticity (STA) [GORK⁺19]. In this method, the trajectory is written as an expansion of either polynomial or trigonometric basis functions with adjustable coefficients. By assigning certain boundary conditions to the trajectory, designed to ensure that the ion's vibrational eigenstate remains the same after transport, the coefficients in the expansion can be fixed. In the time domain, the STA protocols look very similar to the other smoothed transport methods; the subtle difference can be teased out by looking at the Fourier decomposition of the acceleration during the transport. For STA protocols, the component of the acceleration at the trap frequency is minimized, such that the ion will experience a minimal resonant force.

The methods of STA can also be employed to determine optimized protocols for splitting chains. In this case, the parameter which is smoothly varied is the inter-ion spacing, d , given by [HS06]:

$$\beta(t)d(t)^5 + 2\alpha(t)d(t)^3 = \frac{q}{2\pi\epsilon_0} \quad (8.3)$$

which results from the interaction of the confining potential and the Coulomb repulsion between the two ions. From the optimized function for $d(t)$, corresponding functions for $\alpha(t)$ and $\beta(t)$ can be determined. There is an extra degree of freedom in this determination, which we resolve by requiring that the frequency of the center-of-mass (COM) mode stay

constant during the split. The frequencies of the two modes of motion during the split are:

$$\omega_{COM}^2(t) = \frac{q}{m}(2\alpha(t) + 3\beta(t)d^2(t)) \quad (8.4)$$

$$\omega_{STR}^2(t) = \omega_{COM}^2(t) \left(1 + \frac{q^2}{\pi\epsilon_0 m \omega_{COM}^2(t) d^3(t)} \right) \quad (8.5)$$

The COM frequency is minimal when $\alpha = 0$, which is called the critical point. At this point, the inter-ion distance assumes the value:

$$d_{CP} = \left(\frac{q}{2\pi\epsilon_0\beta} \right)^{1/5} \quad (8.6)$$

and the COM frequency is given by:

$$\omega_{COM}|_{CP} = \left(\frac{q}{2\pi\epsilon_0} \right)^{0.2} \left(\frac{3q}{m} \right)^{0.5} \beta^{0.3} \quad (8.7)$$

Thus, we see that for a given ion species, the minimum frequency during the split is determined by the value of β , which is tied to the trap geometry and discuss further in section [8.2.1](#).

The STA trajectories considered here are optimal in the sense that they minimize excitation given a certain trap frequency and protocol time, but they may not be the best by every conceivable metric. One way in which STA is not "optimal" is in how finely tuned it is to the application. We find, through simulation, that small perturbations in the trajectory due to imperfect control, can remove much of the anticipated benefit of an STA technique. There exist some protocols which are optimized to certain kinds of experimental imperfections, but it remains an open research area to determine if there are control methods which are robust to the kinds of experimental imperfections we typically observe in the lab. The very narrow target of the STA method is one of the reasons that the phenomenally-low heat gain expected from STA techniques has not been demonstrated experimentally.

8.1.3 Creating and Simulating Voltage Waveforms

Whether we have determined a trajectory for linear transport or an optimized protocol for a split, we end up with a set of target coefficients in the ion's potential. At any instant in time, the coefficients can be used to solve for trap voltages as described in section [2.3](#)




Figure 8-2: LTSPICE model circuit for the full voltage source, filter network, amplifier and feedthrough into the chamber, including realistic values of stray capacitance. The model for the Fastino is in the upper left part of the schematic. The high voltage amplifier is in the lower middle (compare with Figure A-1). The PI filter and octagon filter, discussed in section 3.2.2, are in the lower right. The three cables in the pathway are represented by their equivalent parallel capacitance, determined from separate measurements.

and appendix B. With this approach, the target coefficients can be converted to voltage waveforms suited to a particular trap geometry.

The challenge comes in trying to capture the time dependence of these waveforms in our experiment. Due to the finite sample rate of our voltage sources, we can only sample a few points of the voltage waveform. Particularly for diabatic schemes, we may be limited to only a few samples of the fastest voltage source that we use (see section 3.2.1). This discretization, in combination with filtering effects, can lead to voltage waveforms on the trap which are quite different from those which we intended to apply. Fortunately, the behavior of the circuit can be studied and well-modeled using circuit simulation software such as SPICE [NP73]; in Figure 8-2, we show such a model. In addition to known component values, both in the circuit for the voltage source and in the high voltage amplifier (see appendix A), we must also consider the effect of parasitic capacitance and resistance, due to the cables and feedthroughs between the voltage source and the electrodes of the trap.

Careful studies were undertaken to measure the parasitic capacitance at each leg of the path. In each of these measurements, an input sine wave is sent in through a known resistance along one of the control lines and then fed back to one of the other control lines with an intentional short. We sweep the frequency of the input sine wave and fit a simple




Figure 8-3: Classical simulation of the trajectory of an ion's motion during a transport operation. The position of the minimum of the harmonic well is shown in the dashed black line and is referenced to the left y axis. The ion's position relative to the minimum is shown in the blue line and is referenced to the right y axis. We see in this case that there is residual oscillation when the protocol is complete. The amplitude of this oscillation can be measured and converted into an equivalent amount of energy [Qi21].

RC low pass filter model to the result in order to determine the parasitic capacitance. These parasitic capacitances are usually on the order of a few hundred pF, which can combine with resistances on the order of 1 k Ω to create poles that are significant for our motion operations (i.e. around 1 MHz). We have determined that filtering effects due to stray inductance and antenna effects in our cables both contribute at frequencies that are outside of a bandwidth that is significant ($\gg 10$ MHz).

Our circuit model can also capture nonlinear behavior, for instance finite voltage slew rates in the amplifiers due to saturation effects in transistors. We can plug in a particular waveform, in the form of a few voltage samples separated by some constant sample time dt . We can then be fairly confident that the output of the model is a close representation of the voltages that would be applied to the electrodes of the trap.

From these simulated voltages, we can easily find the potential above the trap using the same grid files that we use to calculate the trap voltages. Even though we must apply discrete voltages with our voltage sources, we can evaluate the actual applied voltage continuously in simulation. We do this by sampling the output waveform on a fine timescale




Figure 8-4: An example target voltage waveform (dashed black curve) along with a pre-distorted input (solid blue lines). The simulated output of the circuit with the pre-distorted input is shown in the orange curve [Qi21].

and then interpolating the result. With the smoothed output potential, we can estimate the classical heat gain that the ion would experience by integrating the equations of motion of the harmonic oscillator as it is dragged along during the protocol (Figure 8-3). Even though we must formally consider the heat gain quantum mechanically, it has been shown that there is a close correspondence between the classical and the quantum result for this system [HYO⁺08].

8.1.4 Predistortion

The filtering in our system tends to distort the waveform before it reaches the trap electrodes. This is due to the fact that a significant amount of spectral content in the voltage waveform lies beyond the cutoff frequency of our lowpass filters. As voltage waveforms approach the diabatic limit, this effect becomes more pronounced. One way that we can overcome this distortion is by inverting the effect of the lowpass filters in a process called predistortion [Tod20].

The majority of the filtering in our system can be modeled analytically, and these effects are straightforward to invert. In this work, we have used an additional loop of optimization, comparing the target waveform to the numerically-simulated output, while optimizing the input voltages, in order to account for higher-order effects in the circuit, like nonlinearity




Figure 8-5: Surface trap geometries with (left) outer segmentation and (right) inner segmentation [NH12]. Here the variable a refers to the width of the segment between the RF electrodes. The variables W , C , and E refer to the width of the control electrodes; see [NH12] for more description of the optimization of these parameters.

in the op amps [Qi21]; see Figure 8-4 for a simulated example of a pre-distorted waveform. In our experiment, we have achieved early results with this predistortion, but we have not used the predistorted waveforms to push ions faster than the speed limit set by our filters.

8.2 Optimizing Trap Designs

Many aspects of ion motion are set by the voltage sources and filters, but it can also behoove us to improve the control of electric fields by adjusting the geometry of the trap electrodes. In our standard trap design (section 3.6.1), the highest trap frequency that we can achieve given our voltage range of ± 40 V is limited to around 2.4 MHz for calcium. The largest value of β that we can achieve is limited to $1.1 \times 10^{15} \text{ V/m}^4$, which leads to a critical COM frequency of $\omega_{COM} = 124 \text{ kHz}$. A low critical frequency is an issue due to anomalous heating, which usually has ω^{-2} dependence. From our simulations, we have also seen that higher motional frequencies are better behaved during transport.

Larger electric fields for a given voltage can be achieved by reducing the effective distance from the electrode to the ion. From inspection of the standard electrode layout presented in section 3.6.1, one way we might think to do this is by moving the segmented control electrodes from the outside of the trap to the center rail; we'll refer to this style as "center segmentation" (see Figure 8-5 for a comparison). Some theoretical analysis has been performed to determine the optimal electrode geometry for this configuration [NH12]. Here we follow along with this analysis and justify the choices for the dimensions of a trap that was fabricated to meet these specifications.

8.2.1 Choosing the Right Pitch

When designing this trap, we first start out by choosing a trap height. For this, we stick with the $50\text{ }\mu\text{m}$ height of the standard geometry, since this is where most of our systems have been characterized. Any closer to the trap and anomalous heating will start to be an issue, and any further away will begin to require higher RF voltages. From an analytic model of a "five wire" surface trap [Hou08], the width of the RF electrodes and the space between them can be chosen to maximize the trap depth. The ideal values for these parameters are $a \approx 35.0\text{ }\mu\text{m}$ and $b \approx 128.7\text{ }\mu\text{m}$; however, there is some freedom to adjust these values while staying within a few percent of the ideal depth [NH12].

With the geometry of the RF electrode fixed, it then only remains to determine the pitch of the control electrodes. This can be optimized by choosing a geometry which allows for the highest quartic confinement (i.e. highest value of β). From numerical simulation of unit potentials, it has been shown that, for the center-segmented case, this is optimal when the control electrodes all have approximately the same width, and when that width is approximately equal to the gap between the RF electrodes. Thus, the ideal center-segmented electrodes are a series of squares between the two RF electrodes.

When designing a trap with this kind of segmentation, the limits of fabrication ability must be taken into account. Electrodes between the RF rails cannot be electrically connected without internal wiring within the chip substrate, which requires the use of vias (conductive pathways between metal layers of the fabrication). In our process, there is a minimum size to these vias, so there is also a minimum allowable electrode size. Additionally, we want to avoid having too many control electrodes³, so for our final design, we backed off the value of a to $45\text{ }\mu\text{m}$, and the pitch of the electrodes was increased to $45\text{ }\mu\text{m}$ to match.

8.2.2 Tilt Considerations

In the analysis of [NH12], the required tilt of normal modes axes is considered by changing the widths of the RF electrodes. This approach is nice in that it does not require the control electrodes to obtain the tilt, but it also builds in a choice of tilt angle. Since we are

³For this, we chose a limit of 32 independent control electrodes, to match the 32-channels of a single Fastino (see section 3.2.1). At the time of the trap design, we were only anticipating that one voltage source would be available.

considering a geometry which we would like to be extensible to larger arrays, which might possibly have junctions, we wanted to choose a design which left the tilt as something which could be controlled by applied voltages.

In order to achieve tilt, the control electrodes in the optimization loop must have a projection along the transverse $\hat{x}$ direction, and there must be enough degrees of freedom to independently be able to adjust the coefficient of the x and xz terms in the potential expansion. One way to achieve this would be to further divide the center electrodes, splitting each down the y axis to create pairs of electrodes. However, this would greatly increase the number of control electrodes, and it would also require electrodes which are smaller than the via size limit. Another way to introduce this field control is to place additional control electrodes outside the RF electrodes, which also have the advantage of being able to introduce a global tilt for a larger chain of ions. From trap symmetry, we find that we need at least three independent electrodes (two on one side of the axis, and one on the other) to control the tilt, so we choose to place two on both sides for greater flexibility.

When considering the size of these electrodes, we must take into account the limited range of voltages that we can apply (± 40 V); if the electrodes are too far away from the ion, then the voltage optimizer may not be able to find a solution for the desired tilt angle. From some explorations with our unit potential solver, we found that we were unable to find suitable tilt electrode widths with the RF electrode width which maximized the trap depth. As noted above, the RF electrode width can be tuned away from the optimal value without a huge effect. We reduce the RF electrode width to $b = 90 \mu\text{m}$ and find that we are still able to achieve a trap depth of $110 \text{ meV} \approx 1300 \text{ K} \cdot k_B$, which is well above the initial temperature of the neutral atoms, especially considering the pre-cooling from the 2D MOT. This level is also higher than the depth we would obtain with our standard geometry, given the same RF voltage, which would be $\sim 50 \text{ meV}$.

From further exploration, we found an inner tilt electrode width of $30 \mu\text{m}$ and an outer tilt electrode width of $70 \mu\text{m}$ to be a good compromise. With this geometry, we were able to obtain tilt values as high as 30° (a standard value for us) while keeping control voltages within an acceptable range.




Figure 8-6: (a) Picture of the trap as fabricated. Five electrodes with a dense pitch, optimized to give a high value of β , comprise each of four major "zones" along the trap axis. Between these zones, there are two wider electrodes; these areas are primarily meant to transport ions between the splitting zones. Outside of the RF electrode, there are four total electrodes used to help generate tilted potentials (discussed in section 8.2.2). The RF and tilt electrodes are all symmetric about the trap axis, which lies along a vertical line in this view. (b) Drawing showing the internal layers; the purple segments are internal routing layers which are used to bring in voltage signals from the outside of the trap to the middle. The width of the wiring is decreased as it passes under the RF electrode to reduce the capacitance to ground on the RF. Most electrodes are electrically connected to bond pads around the outside of the trap via this internal wiring. However, the RF electrode is done in a contiguous "tuning fork" shape entirely in the top metal layer, similar to geometry of our standard trap.

8.2.3 The MaxBeta Trap

All of the dimensions given above were taken in the simplified case in which all electrodes directly abut each other. In reality, we must allow gaps between these electrodes, and for this we use a standard value of $4\mu\text{m}$, which we have found to be compatible with the RF voltages that we typically apply. All dimensions we found from theory were reduced by $2\mu\text{m}$ to account for the gaps.

In the fabricated traps, we use three metals layers. The top layer is $1\mu\text{m}$ thick niobium, in which we define the trap electrodes and bonding pads which are used to connect to external voltage sources (see Figure 8-6 and Table 8.1 for geometry details). Below this is a $5\mu\text{m}$ layer of SiO_2 followed by another layer of $1\mu\text{m}$ niobium which is used to define

Table 8.1: Dimensions of the MaxBeta trap

Segmented electrode pitch (split areas)	41 μm
Segmented electrode pitch (transport areas)	70 μm
RF electrode width	86 μm
Space between RF electrodes	49 μm
Inner tilt electrode width	26 μm
Outer tilt electrode width	66 μm
Gaps between electrodes	4 μm
Pseudopotential minimum height	50 μm

the internal routing from the bond pads to the the trap electrodes. In order to limit the capacitance of the RF electrode, we try to minimize the area of the overlap between the RF electrode and the conductors in this layer. Below this, there is an additional layer of 5 μm of SiO_2 followed by a final layer of 1 μm niobium. In the bottom metal layer, we define a ground plane as with the chips with integrated photonics (see section 7.1.3).

Using the as-fabricated dimensions, a final round of simulation was performed with a realistic geometry which took into account these gaps, the finite conductor thickness, the shape of the vias and the dielectric layers. We found that the performance was similar to our expectation given the idealized geometry, and we did not anticipate any effect from gradients in the pseudopotential to lead to severely limiting technical noise from this source. We do note that the decreased distance from the control electrodes to the ion increases the sensitivity to voltage noise on these electrodes (see section 4.1.1), so additional care must be taken to filter voltage sources when using this trap.

From the simulated unit potentials, we were able to calculate voltages to generate values of β as high as $2 \times 10^7 \text{ V/m}^4$, which corresponds to a critical frequency of $2\pi \times 1.3 \text{ MHz}$. The fine electrode pitch also allowed us to generate simple harmonic frequencies at least as high as 4 MHz without reaching the limits of our voltage range. From further exploration, there are other interesting potentials which can be realized with this geometry, including shapes which result in equally spaced chains of ions, which is not usually the case for long chains of ions in simple harmonic wells [Qi21].




Figure 8-7: (a) A large cloud of ions loaded in the MaxBeta trap (b) The scatter from the ions on the trap, near the middle of one of the storage zones

8.3 Experimental Results

Using the center-segmented trap, and the voltage distribution hardware described in section 3.2 and in appendix A, we were able to perform some initial investigations into ion motion with the improved geometry with promising results.

8.3.1 Baseline Characterization

With the MaxBeta trap, we were able to load large clouds of calcium ions (over 50) by decreasing the RF voltage to around 16 V and by dropping down the axial frequency to around 500 kHz (see Figure 8-7). The voltage sets that we use for loading do not have any applied tilt. We presume there must be enough residual rotation of the motional axes, due to asymmetry in the trap geometry or to other background fields, in order to get a sufficient projection of the radial modes with the Doppler cooling beam. In all experiments presented here, we start by loading in the second zone of the trap.

To avoid dealing with charging effects from the UV PI lasers, we shuttle ions away from the loading area before proceeding with experiments. After loading a large cloud of ions, we then ramp the RF voltage up to 36 V and then increase the axial frequency to 2.5 MHz. We then perform a slow (~ 1 s) transport from the second trap zone to the third. When the ion arrives at this final position, we ramp the RF voltage up to the final value of 72 V, at which point the radial mode frequencies are both in excess of 8 MHz. During the ramp up, the number of ions in the cloud usually dropped down to around 5 or less.

When generating transport voltages for moving between zones, we had difficulty keeping

some voltages from blowing up; distant electrodes, which have only a negligible effect on the field at the ion's location, would tend to take on large voltage values. To limit this behavior, we manually fix the voltage of electrodes when they are beyond a certain distance from the ion. This approach had been suitable for the standard geometry, but on the MaxBeta trap, the problem is exaggerated, and we see additional discontinuities in our voltage waveforms as electrodes are dropped from and added to the optimization. In the end, we simply tweaked the voltage generation by adding or removing electrodes from optimization and adjusting other parameters of the optimizer. We ended up finding a sequence of 20 voltage sets which we could fill in using linear interpolation. We found this set of voltages worked well over many rounds of loading. For small transport operations within a zone, we did not have the same difficulties generating voltage waveforms.

The entire procedure of turning on PI beams, ramping voltages and frequencies, and transporting to the final zone was all automated into a single script which could be re-run any time ions were lost. By adjusting the amount of time that the PI beams were left on, we could change the average number of ions that were loaded, though the final result was always random. In order to do the experiments with multiple ions described below, we simply kept reloading until we obtained the desired number⁴.

At the final position, we measured some single-ion heating rates as a function of frequency (see Figure 8-8). Actually, initial ion heating results with this trap were at a much higher level than this, which turned out to be a limitation of voltage noise on the control electrodes. Since the MaxBeta geometry has electrodes with a much lower effective distance to the ion, the voltage noise to electric field noise conversion of equation 4.5 is much harsher on this trap. Thus the $1 \text{ nV}/\sqrt{\text{Hz}}$ rule which we apply to our other traps must be adjusted to less than $0.1 \text{ nV}/\sqrt{\text{Hz}}$ for voltage noise to not be a limitation. In initial tests, we had been omitting the filter motherboards discussed in section 3.2.2, so the voltage noise on trap electrodes was at a level of $3 \text{ nV}/\sqrt{\text{Hz}}$, well above the limit to cause issues and leading to a heating rate in excess of 200 quanta/s. To avoid technical noise due to phase noise on the RF source, we performed all of these measurements with the bandpass filters in place in the RF distribution (see section 4.1.2).

With all of the filters in place, we have more measurable levels of heating that agree

⁴A ion chain assembly script, which could shuttle in loaded ions and combine them with an existing crystal, would be a nice addition. But without more control on the exact number of ions that are loaded, it would have been difficult in practice.




Figure 8-8: Heating rates on the MaxBeta trap, measured at three different axial frequencies. The dashed blue line is a fit to a function of the form $\bar{n}(f) \propto \omega^{-\alpha-1}$ (see [BSC15] for similar analysis). From the fit, we obtain $\alpha = 1.0 \pm 0.2$, which is in good agreement with an electric field noise density that scales as $1/f$ (see section 4.2).

well with the $1/f^2$ behavior that we would expect from anomalous heating limited behavior (see section 4.2). By accounting for the mass and trap frequency difference, these heating rates can be converted to equivalent heating rates for strontium at 1.3 MHz, for which we have much more historical data. For all of the datapoints above, the equivalent level of heating for strontium would be about 40 quanta/s, which is about a factor of 4 higher than we would normally expect (see Figure 4-1). Since the dielectric below the RF electrode of the MaxBeta trap is SiO_2 , which is lossier than the usual sapphire dielectric, it is possible that heating from the RF voltage is causing the temperature of the trap to increase, as was observed on the photonics chips in section 7.3; it is also possible that these chips could benefit from a round of ion milling before installing them in the chamber.

8.3.2 Linear Transport

With the trap frequency set to 4 MHz, the level of heating was low enough that ion's axial motion could be reliably initialized to a level consistent with 0 quanta. This allowed us to perform some experiments to measure the amount of motional excitation gained from a linear transport operation. To do this experiment, we first cool the ion to the ground state of motion at the center of the third trap position. From here, we produce a voltage

waveform to translate the ion from the middle of the zone to the middle of the next electrode to the right (i.e. $+45\text{ }\mu\text{m}$). We allow the ion to dwell here for a few tens of μs , and then we perform a second voltage waveform to translate the ion back to the starting position. The average axial motion can then be measured in the usual way; the one-way increase in motional excitation is half of the resulting value.

The speed of our transport is limited by our ability to predistort voltages, which we were not able to fully debug. The fastest motion which was achievable was then to simply apply a single-length "waveform" in which the ending voltage set was instantaneously applied. Rather than get kicked immediately, the ion is transported more smoothly due to the effect of the filter network between the voltage sources and the control electrodes. In simulation, we have seen that the time-domain effect of all of these filters is to naturally convert a step function to something which looks similar to a $\tanh$ ⁵. Thus this very naive approach still gives us an increase in excitation which is consistent with zero. The fact that this works is likely owed to both the natural smoothing of the filters and the stiffness of the 4 MHz trap.

To verify that transport has happened as fast as we expect, we can take advantage of the fast timing of laser pulses. Given a Doppler-cooled initial state, the ion will be in a mixed state of equal probability in the $|S_{1/2}, m_j = -1/2\rangle$ and $|S_{1/2}, m_j = +1/2\rangle$ state. If we then turn off Doppler lasers and hit the ion with a π pulse with the qubit laser (tuned to the carrier), we can prepare a mixed state of $|S_{1/2}, m_j = +1/2\rangle$ and $|D_{5/2}, m_j = -5/2\rangle$. Subsequent detection will result in measurements that are 50% fluorescent and 50% dark. In these experiments, we typically operate with the qubit laser focused on the starting position of the transport operation; by aligning the qubit beam to different positions along the trap, we can effectively use the π pulses to determine whether or not the ion was in the beam at the time of the pulse.

The full experiment for transport verification (shown in Figure 8-9) will be to start with the ion Doppler cooled and then switch off all lasers. The forward transport waveform is then played. Once this completes, we do a π pulse on the qubit beam, followed by a delay of a few tens of μs . Finally, we perform the reverse transport operation to get the ion back to the starting position and then detect the state of the ion in the usual way. The bright

⁵The transfer function of the linear part of the voltage circuit can be used to write an analytic expression for an ideal step response, but nonlinear effects, particularly the finite slew rate of the op amps, can distort the result. In our application, it is most convenient to simply use the SPICE model of Figure 8-2 to calculate the output waveform numerically.




Figure 8-9: [Top] Timeline of the transport verification experiment. The sequence begins with a period of Doppler cooling. Before the forward transport sequence, the current time is recorded. The first wait time, before the π pulse, is variable and can be scanned in a loop. After the π pulse, there is a fixed delay, included to avoid timing errors as the backward transport waveform is prepared. After the backward transport, there is an additional fixed delay and then the stopping time is recorded. [Bottom] Results of a transport verification experiment for a simple two-step transport protocol. The time reported on the x axis is $t_{stop} - t_{start}$. In the green dataset, the verification beam is focused at the endpoint of the transport. In the red dataset, the verification beam is focused at the starting point. When the probability to be in the bright state decreases, we know that the ion is at least partially within the beam. The transit time can be determined from the width of the segment where both datasets are in the bright state, meaning that the ion is between the two locations (the area highlighted in light blue is a conservative bound). In the experiment shown here, this width is about $19 \mu s$.

population is measured as a function of delay time between the start of the experiment and the start of the π pulse. This scheme has the advantage that all of the Doppler cooling lasers and detection optics do not need to be translated. The only beam which needs to be moved is the qubit beam.

To verify transport, we do this experiment in two different ways. In the first case, we start with the qubit laser aligned to the starting position. In this way we can effectively determine the time at which the ion has left the starting point. We note that this is not necessarily the same as the time at which the commands are sent to the voltage source, since




Figure 8-10: Frequency of the normal modes of motion during a split on the MaxBeta trap (see also [Qi21]). Blue lines represent the stretch mode and black lines represent the COM mode. The dashed lines are the expected values, given our target STA protocol and the solid lines were derived from measured spectroscopy.

there is both finite processing time and latency, in addition to propagation delay from the filters. After this, we perform the same experiment sequence with the qubit beam aligned to the final position. In this way, we determine the time at which the ion has arrived and the transport has concluded. For the single-step voltage waveform, we find a transport time of $19\ \mu\text{s}$, which corresponds to an average speed of $45\ \mu\text{m}/19\ \mu\text{s}$.

The realized transport in this work is an order of magnitude slower than the record values for linear transport [WZR⁺12, BGL⁺12]. Yet, the total time for our operation is still short enough that it is not the slowest operation in an experiment, and it is much lower than the coherence time of our best systems. It is still useful to know that transport at this level is attainable with no predistortion or special timing required.

8.3.3 Ion Behavior during Splits

We concluded by performing some experiments to demonstrate the main function that our MaxBeta trap was designed to do: crystal splitting. We generated voltage sets to split a chain of two ions following an STA protocol to smoothly vary the inter-ion distance. During the split, the frequency of the COM mode was held constant at 1.2 MHz, which we expected to make a significant difference over split experiments which were performed with trap with

segmented electrodes on the outside of the RF electrodes [RWK⁺14].

With this voltage waveform, we were able to reliably split chains of ions, after careful tuning of the the applied electric field in the y direction [KRS⁺14]. We did not use any laser pulse schemes to verify splitting, so from the frame rate of the camera, we can only be sure that we were able to split as fast as ~ 30 ms. Ions survived several thousand such split/join cycles before we would lose one.

To verify how well the voltage waveforms were matching our expectation and attempt to calibrated the effect of stray fields [KLS⁺20], we performed some step-by-step spectroscopy using the qubit laser to measure the motional sidebands due to the stretch and COM modes of motion, and ensure that these were evolving as expected during the split. The results of this experiment, shown in Figure 8-10, were surprising. It turned out that that the mode frequencies remained consistent with the frequencies of a single harmonic well, i.e. $\omega_{STR} = \omega_{COM}\sqrt{3}$, for the majority of the initial steps in the sequence, in contrast to the smooth decrease of ω_{STR} that was expected. We also found that the COM frequency did not remain constant and actually decreased and increased during the operation. At one point in the sequence, both motional frequencies decrease sharply, and, on the other side of this, the split has apparently completed with mode frequencies that match the expectation of two separate wells.

There are likely multiple mechanisms at work which are interfering with our control of the ions' motion. In simulation, we were able to explain some of this discrepancy by modeling the effect of a background electric field in the z direction, which could be due to effects from laser-induced charging. In the experiment, we saw that the addition of small voltage offsets applied to electrodes near the point where the frequencies dropped was able to change the mode frequencies significantly. The center-segmented geometry is advantageous in that it helps us to achieve strong potentials, but the price is that small voltage error can introduce big deviations. In future work, we will try to develop quartic voltage sets which are more robust to errors in the applied voltage. We will also proceed with a heat-gain analysis of the split, since it is possible that the smoothing from the filters might still limit the heat gain in the same way we have seen for the single-step transport described above.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 9

Conclusions and Future Directions

*“Now here I go again, I see the crystal visions
I keep my visions to myself”*

—Fleetwood Mac

In the course of this work, we have examined a few different control systems which have been either successfully integrated into an ion trap or at least shown to have a path to this end. The goal of this thesis was to develop technologies and control techniques which could be applied to increase the density of control systems in an ion trap so that larger and more complex ion arrays can be constructed. While we did not see this development through to the eventual completion of the unit cell, we still examined the performance of several novel components and found many promising results.

The integrated voltage sources presented in chapter 5 mark the first time that voltage sources have been integrated into an ion trap, and the switch isolation architecture presented promises a low-noise, high-bandwidth system which can be used to greatly increase the density and speed of voltage control. These voltage sources can now be used to construct more complicated trap designs than was previously possible. Adding a digital control interface into the vacuum system can enable high-speed motion operations, which can take advantage of the predistortion tools presented in chapter 8. The demonstration of multi-wavelength photonics presented in chapter 7, the first time that all wavelengths required to operate a trapped-ion qubit have been implemented in a photonic platform, along with complimentary work that used integrated photonics to perform multi-qubit logic [MZM⁺20], establishes photonics as an important and viable tool in more complex trap designs. Ion

motion techniques allow for increased flexibility and the ability to multiplex qubit control, which will be valuable in larger arrays. The ion motion operations in the multi-zone photonics trap shown in section 7.5 opens the pathway to more sophisticated examples. The passive vibration stability of integrated photonics which was demonstrated here will likely be applied in portable atomic systems, like gyroscopes or clocks. Finally, the laser design presented in chapter 6 breaks a paradigm of requiring a laser based on a bulky stabilization cavity in order to control the quadrupole transition that we use for qubit or clock experiments. It's possible that an SBS laser could be used in a miniaturized optical clock that could be operated on a satellite in space or in a rack-sized quantum computer based on ions.

9.1 Future Experiments

With the various components and techniques that have been presented here, there are a number of exciting applications that can now be investigated. Here are a few areas in which the work may be extended in future research.

9.1.1 Junction Design and Operation

One critical component of an array architecture which was not discussed in detail yet is a junction [BOV⁺09, BOV⁺11]. These crossings are important for being able to reorder chains of ions and for distilling down to the few ions required for each operation in an array trap. Since junctions necessarily break the symmetry of the RF electrode, they are typically challenging to use with ions, due to the effects of micromotion that cannot always be compensated away.

In the course of the trap design undertaken in section 8.2, we also explored some designs which included X-junctions (see Figure 9-1), based closely on the work in [WAF⁺13]. We improved upon this design by increasing the trap depth with a modified geometry and then adopted the center segmentation employed in the previous chapter. These modifications should improve our ability to control ions as they pass through the junction.

The numerical tools developed in section 8.1 are well suited to studying motion in junctions, since in these cases there is not an analytic prescription for transport. Instead, the heat gain through a junction could be modeled in simulation and optimized using




Figure 9-1: Diagram of a trap design with a junction. Two junctions were combined to make a simple array building block, which can be compared with the design shown in Figure 1-1. The rectangular holes in the labeled photonics zones have grating couplers of various wavelengths like those presented in section 7.1.3. The large, contiguous electrodes with the squiggle edges are the RF electrodes. Note the segmented control electrodes in between these electrodes, which are similar to those discussed in section 8.2. The circles in the control electrodes are vias which connect to interior metal layers of the trap design, used for routing.

numerical tools. The noise mitigation techniques presented in section 4.1 can also be used to limit the amount of technical noise in these experiments.

9.1.2 Faster Transport

The machinery for predistorting voltage waveforms was developed in the course of this work, but it was just beyond our reach to apply these with ions. When eventually performing these experiments, it will be particularly illuminating to determine how precisely-tuned our model of the voltage sources must be to get reliable pre-distortion. If it turns out that small

discrepancies in the model, due to variation in parasitic capacitance for instance, change the behavior of the waveforms enough that low heat gain is no longer guaranteed, then it may be better to stick with more naive transport methods, which we have found to be naturally filtered to an acceptable level by our hardware. The tradeoff between speed of transport and simplicity of design will be an important consideration in an array with many such operations.

9.1.3 Architecture Studies

In this work, and a few other recent studies [RHS⁺20, MZM⁺20, TVM⁺21, ISK⁺21, SIS⁺21], many components of integrated technologies required for scalable control within an ion trap array architecture have been independently characterized. An exciting horizon is now opening up in which we can begin to study the interactions and tradeoffs between different components. By combining several of the components developed here into a single trap, we may finally begin to achieve benefits of scaling that overcome the effort required to miniaturize and integrate systems into ion traps.

The multi-zone trap discussed in section 7.5 has the potential to enable some rudimentary architecture studies. We present now an outline for an experiment to explore the limitations of an array-based scheme versus a linear register. In a facsimile of a linear register, we can trap a chain of two ions, which we keep in the same well for the duration of the experiment. To get individual addressing in this case, we use one of the "local" grating coupler channels which primarily has a projection in the radial direction. To test a broad range of functions, we can perform a CNOT gate, which includes both a two-qubit gate and individual qubit rotations (see [DLF⁺16] for a description of such an operation). An alternative approach to this operation is to obtain individual addressing by splitting the chain of ions and then moving the two to separate zones such that individual rotations can be performed independently via the local channels. There is room to learn about the space of tradeoffs by exposing the limitations due to the heat gained during a split, the extra time required to shuttle a chain back and forth to do individual addressing, and the effects of optical crosstalk.

9.2 Final Thoughts

Looking ahead, as ion control technology is advanced and as more components are integrated into traps, new benchmarks of performance will be desirable. Explorations such as the ones in this thesis, which evaluate the performance of quantum systems with advanced classical control, are key in this process of assessing the utility of modern control technology for future quantum sensors, simulators, and computers. While we may not have unlocked the path to a quantum version of Moore’s law yet in the course of this and the other, related work, we have made significant strides towards future systems that are becoming increasingly more complex while also increasing the number of qubits that can be controlled.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A

High-Voltage Amplifier Design

Given the ± 10 V limit of the voltage sources that we use, our typical trap geometries, and desired trap frequencies, it is useful to have an amplifier to extend the range of available voltages. Since we care about the bandwidth of the voltage circuit for doing fast ion motion operations, special consideration must be taken in the design of this amplifier and the choice of components.

A.1 Schematic and Layout

The high voltage amplifier is designed around the LTC6090-5, which is a high-voltage (± 70 V range) op amp with the highest voltage slew rate ($37 \text{ V}/\mu\text{s}$) in a viable package available at the time of this writing¹. With this op amp, we construct a standard non-inverting amplifier with voltage gain of $4\times$. The particular variant of the op amp chip that we use is only guaranteed to be stable for gains >5 , so we employ an additional resistor, placed in between the inputs to the op amp in a technique called "noise gain manipulation." On the output of the amplifier, we include a small series resistance (30Ω) to reduce the effect of instability and prevent the circuit from oscillating when driving capacitive loads [Jun05, HH20]. When evaluating the system performance, we also take into account the effect of subsequent filter stages and confirm that the entire circuit does not produce unstable behavior when driving trap electrodes directly.

At the input to the amplifier, we include a single-pole low-pass filter consisting of a discrete resistor and capacitor combination. The cutoff frequency (~ 1.5 MHz) is chosen

¹Higher slew rate amplifiers exist, but they are much larger ($>1 \text{ cm}^2$) and more expensive (>50 \$).




Figure A-1: Schematic of the high-voltage amplifier.

to limit the voltage noise at frequencies greater than our trap frequencies of interest (up to 4 MHz). For typical transport protocols, we do not expect distortion from this filter to play a large role. Two other filters are used. One is a high frequency electromagnetic interference (EMI) suppression filter (NFE61PT472C1H9L), designed to reduce noise at frequencies above a few MHz, and the other is a common mode choke (Würth Elektronik 744229), which helps to reduce noise coupled onto the output of the voltage amplifier.

Some effort is made to allow the voltage amplifier to function even with a voltage source which may not share the same exact ground connection as the amplifier's power supply. To do this, all return paths in the amplifier circuit are referenced to a negative input pin for each channel from the voltage source. This arrangement is useful in cases where we want to keep the ground of the voltage sources separate from the mains ground to help avoid unwanted ground loops. To give the design some flexibility, a ground connection between the power supply and the negative rail of the input voltage can be made by populating some resistor footprints which are normally left open.

Each of our DAC voltage sources has 32 channels, so we combine 32 voltage amplifiers like the one shown in Figure A-1 in the full design (Figure A-2). Voltages are routed around on these boards in a carefully designed pattern so that the pins in the output connector correspond directly to the pins of the input connector. This allows us to "drop in" the voltage amplifier system whenever we want to use it without making any other changes; no re-mapping of control voltages is required.




Figure A-2: Layout of the high-voltage amplifier.

A.2 Testing Results

After assembling the voltage amplifier and all of the other filter components discussed in section 3.2.2, we were able to test the performance of the amplifier against our expectation from a model of the circuit (see Figure A-3 for the data and Figure 8-2 for an example SPICE model). To do this measurement with realistic stray capacitance, we connected the full system to the chamber, and then used a special trap configuration which had an additional wire for routing the signal from a control electrode back outside of the chamber. (The stray capacitance of this additional wire was also characterized and modeled.)

At first we found that the agreement between the model and the measurement was not quite good; the measured data indicated that the source was slewing faster than expected. We found that adjusting the values of linear components (e.g. tweaking capacitor values) was not enough to explain this behavior. However, by modifying the model of the LTC6090-5 slightly in our circuit simulation, we were able to obtain the agreement shown in Figure A-3. More characterization will be required to determine if this same tweak is required




Figure A-3: Response to a pulse applied using the high-voltage amplifier and all of the other filters in the chain.

for all of our LTC6090-5 chips or if there is variation between devices. With the modified model, we will be able to run our predistortion algorithm discussed in section [8.1.4](#).

Appendix B

Ion Trap Voltage Generation

The software that we use for generating voltages is divided into a few components. At the heart of the system is a numerical optimization loop which adjusts the applied voltages to bring the total potential above the trap as close as possible to the desired configuration. Most of this optimization is independent of the geometry of the trap. In this appendix, we will describe the different components of voltage generation.

B.1 Generating Grid Files

Given some trap geometry, we can use electrostatics to determine the electric field at any position above the trap, for a particular configuration of voltages. In general, this can't be done analytically, so we must use some numerical tools to solve for the potential. There are a few different methods to do this, including boundary element methods and an electric field analog of the Biot-Savart law [OM01], both of which will not be used here. We choose to use a finite element method (FEM) with COMSOL [SPM⁺10, PH17] since this allows us to take into account different dielectric materials and complex geometries which consist of multiple layers that are not necessarily rectilinear.

To create the model of our trap, we start with the 2D geometry which comes from the photolithography masks. Using a 3D drafting tool (Solidworks), we add thickness to this drawing and also include internal layers, like ground planes. Generally, we do not model the internal layers in full detail, though it is in principle possible to do so. From experience tweaking such things, we have found that structure far away ($> 500 \mu\text{m}$) from the ion does not have a large effect in the potential. In Solidworks, we create separate bodies for all

internal layers and for each control electrode to which we will apply a voltage, including a ground plane, if applicable.

FEM proceeds by numerically solving partial differential equations from Maxwell's equations in steady state and subject to boundary conditions which we can specify. The area over which we want to solve the equation is discretized into many elements, each with an independent set of basis functions; the coefficients of these basis functions can then be adjusted to try to find a smooth composite solution to the equation which satisfies the boundary conditions. Finding the right discretization is part of the finesse of using FEM. We use a mesh generation tool in COMSOL to split up the geometry. Near the ion, we want a precise description of the electric potential, and so here we require a fine mesh. If we use a fine mesh throughout, this will result in an optimization problem with too many elements, so we allow the mesh to become progressively more coarse further away from the ion. This is done by introducing non-physical shapes into the 3D model in which we can apply different meshing rules [HPKM12]. In each of the elements, we use fourth-order basis functions, which we have found to give smoother results than the default second-order functions.

We apply boundary conditions using a parametric sweep. For each electrode in the model, we perform one simulation in which that electrode has a constant 1 V potential on the entire outer boundary, and all other conductors are at ground. The result of each step in the sweep is a solution for the potential at each point in the region of interest; we call these simulated datasets "grid files." By the linearity of the Laplace equation, we can determine the potential for an arbitrary voltage configuration by multiplying the potential in each grid file by the voltage applied to the corresponding electrode and then summing the result together. Since we use an electrostatic simulation, this does not capture the time dependence of the potential above the RF electrode, so we use the pseudopotential approximation (see section 2.7), taking the norm squared of the electric field due to the RF electrode and substituting this as an effective static potential.

B.2 Determining Taylor Coefficients

We typically have some desired configuration for the potential above the trap, given some trap frequency, ion position and tilt of the normal modes. It is possible to convert these

requirements into a potential field above the trap (i.e. a matrix of coordinates with some potential at each point). In principle, the voltage applied to each of the grid files can be adjusted using numerical optimization to realize a potential field that matches the target. But in practice, the number of degrees of freedom in the target potential (one for each point in the field) is quite large and can make optimization tricky. As an intermediate step to reduce the degrees of freedom of this field, we can convert each of the grid files into coefficients of a Taylor expansion, centered about a point which is given by the target ion location. We need only perform this expansion up to the highest order coefficient in our target potential, which may simply be second-order if we just want to define the potential of a single ion or a small chain.

From each electrode/grid file, we get an array of coefficients of the form:

$$\mathbf{C}_i = [E_x \ E_y \ E_z \ \alpha_{xx} \ \alpha_{xy} \ \alpha_{xz} \dots] \quad (\text{B.1})$$

with a term for each permutation of x, y and z. From the orthogonality of these basis functions, the coefficient array for the total realized potential is $\mathbf{C} = \sum_i V_i \mathbf{C}_i$, where V_i is the voltage applied to each electrode. The number of coefficients can be much smaller than the number of points in the field.

B.3 Voltage Optimization

The main task is to determine a set of voltages V_i which realizes the desired potential. This can be expressed as minimizing a cost function $\|\sum_i V_i \mathbf{C}_i - \mathbf{C}_{target}\|$. A choice for V_i can be obtained with the use of numerical optimization tools. We use a Sequential Least Squares Programming (SLSQP) algorithm written in SciPy. There are a number of parameters of the optimization which can be tuned (discussed more in [Qi21]) in order to speed up the process or make it more reliable. We find that, since the problem is in general underconstrained (more free voltages V_i than there are desired constraints in $\mathbf{C}_{target}$), optimization is prone to hopping between degenerate solutions (solutions which have similar values of $\|\sum_i V_i \mathbf{C}_i - \mathbf{C}_{target}\|$ but not similar voltages V_i). This can be addressed by limiting the number of free voltages to only a few nearby the ion; we do this using heuristic "relevance" ranges to decide which electrodes participate in optimization and which are held at fixed voltages. Solution degeneracy can also be avoided by introducing more elements to the cost function,

for instance restrictions on how large voltages can be overall and how much they are allowed to change between solution steps, in the case of transport [BOV⁺11].

While the process of importing grid files takes around a few minutes, depending on the number of electrodes, the voltage optimization step can be much faster. For second-order expansions, trap voltages can be determined in less than a few seconds.

B.4 Calculating the Total Potential and Generating Tilt

In the voltage optimization above, only the control electrodes were participating. The simplest case, in which we want to specify a trap of a certain frequency at a certain position, requires only a target set of coefficients $[E_x \ E_y \ E_z \ \alpha_{yy} \dots]$, where again we are assuming that the trap axis is in the y direction. Specifying $E_x, E_y, E_z = 0$ at a point is equivalent to defining a local minimum in the DC part of the total potential¹. We choose the position of the DC minimum to be co-aligned with the pseudopotential minimum. In the control potential, we initially do not constrain α_{xx} and α_{zz} , so these are free to wander. By Laplace's equation, at least one of these two will be negative. But the strength of the ponderomotive confinement overcomes this, and a net positive radial trap remains. In simulation, we can sum the control potential, given a certain voltage set, together with the pseudopotential, determined from the grid file for the RF electrode along with the RF voltage and frequency, in order to get a total potential. This can then be fitted and plotted in order to visualize the potential seen by the ion.

The directions of the normal modes in the radial direction are usually tied to the geometrical axes of the trap, but it is possible to rotate the modes by adjusting the control potential. For a given tilt angle θ_{target} , measured between the most vertical normal mode and the direction perpendicular to the trap, we can adjust the terms α_{xx} , α_{xz} , and α_{zz} in $\mathbf{C}_{target}$ in order to get the *total* potential to match the expectation. In our system, we do this with an optimization loop, which adjusts the angle $\theta_{control}$ between α_{xx} and α_{zz} in order to minimize $|\theta_{target} - \theta_{total}|$ ². With the addition of tilt, we have seven total constraints in $\mathbf{C}_{target}$; we will often also constrain $\alpha_{xy}, \alpha_{yz} = 0$, though these are often naturally met

¹It is possible that other, higher-order odd terms in the DC part of the potential could defeat the minimum, but the cubic terms and beyond are typically weak in our trap geometries.

²In order to determine the values for α_{xx} and α_{zz} , we use the heuristic $\alpha_{xx} = 3\alpha_{yy}$ and $\alpha_{xx} = -4\alpha_{zz}$. This naturally satisfies Laplace's equation, and we find that it gives good results for our typical trap geometries and frequencies of interest.

by symmetries in the trap.

B.5 Translating and Applying Voltages

At the end of the optimization, we have an array of voltages $\mathbf{V}$. The index of the voltage naturally corresponds to the index of the grid file. To simplify the accounting, we name each of the electrodes (and thus grid files), using a naming scheme based on the CPGA pad to which the electrodes are usually bonded. We can then use these names as keys in a python dictionary structure (dicts). The dicts can be stored to disk or transmitted between programs.

To apply the voltages to the trap electrodes, we unpack the dicts in software. Here we have a hard-coded array of electrode names, which specifies the order in which the electrodes are connected to the physical DAC hardware. We have functions written which can convert an input dict into an array of voltages to apply, also accounting for the gain of the amplifiers. We also have functions for taking a list of dicts and converting to a timed array of voltages, for use when doing ion motion operations.

Since the grid file import step takes a long time, we have recently started running the voltage optimization in a separate thread³. Experiments which require new sets of voltages can send socket requests to the voltage optimizer and quickly get back dicts. It is also possible to save certain waveforms to disk and read them back, in cases where we want to repeatedly apply the same sequence.

³Specifically, we use the Network Device Support Package (NDSP) protocol in ARTIQ in order to get native compatibility with our existing experiments.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

- [AGH⁺11] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas. Electric-field noise above a thin dielectric layer on metal electrodes. *New J. Phys.*, 13:123023, 2011.
- [AHM⁺95] R. Alheit, C. Hennig, R. Morgenstern, F. Vendel, and G. Werth. Observation of instabilities in a Paul trap with higher-order anharmonicities. *Appl. Phys. B*, 61:277–283, 1995.
- [ALKH13] J. Alonso, F. M. Leupold, B. C. Keitch, and J. P. Home. Quantum control of the motional states of trapped ions through fast switching of trapping potentials. *New J. Phys.*, 15:023001, 2013.
- [ALS⁺16] J. Alonso, F. M. Luepold, Z. U. Solèr, M. Fadel, M. Marinelli, B. C. Keitch, V. Negnevitsky, and J. P. Home. Generation of large coherent states by bang-bang control of a trapped-ion oscillator. *Nat. Commun.*, 7:11243, 2016.
- [AMUH19] Da An, Clemens Matthiesen, Erik Urban, and Hartmut Häffner. Distance scaling and polarization of electric-field noise in a surface ion trap. *Phys. Rev. A*, 100:063405, Dec 2019.
- [ANK⁺15] N. Akerman, N. Navon, S. Kotler, Y. Glickman, and R. Ozeri. Universal gate-set for trapped-ion qubits using a narrow linewidth diode laser. *New J. Phys.*, 17(11):113060, nov 2015.
- [APL03] Vilson R. Almeida, Roberto R. Panepucci, and Michal Lipson. Nanotaper for compact mode conversion. *Opt. Lett.*, 28(15):1302–1304, Aug 2003.
- [AUW⁺10] J. M. Amini, H. Uys, J. H. Wesenberg, S. Seidelin, J. Britton, J. J. Bollinger, D. Leibfried, C. Ospelkaus, A. P. VanDevender, and D. J. Wineland. Toward scalable ion traps for quantum information processing. *New J. Phys.*, 12:033031, 2010.
- [BB14] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theor. Comput. Sci.*, 560:7–11, Dec 2014.
- [BCH⁺19] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibbrandt. $^{27}\text{Al}^+$ quantum-logic clock with a systematic uncertainty below 10^{-18} . *Phys. Rev. Lett.*, 123:033201, 2019.
- [BCMS19] Colin D Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M Sage. Trapped-ion quantum computing: Progress and challenges. *Appl. Phys. Rev.*, 6(2):021314, 2019.

- [BCSH21] K. R. Brown, J. Chiaverini, J. M. Sage, and H. Häffner. Materials challenges for trapped-ion quantum computers. *Nat. Rev. Mater.*, 2021. <https://doi.org/10.1038/s41578-021-00292-1>.
- [BGL⁺12] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. Coherent diabatic ion transport and separation in a multizone trap array. *Phys. Rev. Lett.*, 109:080502, Aug 2012.
- [BHJ⁺11] Jared F. Bauters, Martijn J. R. Heck, Demis John, Daoxin Dai, Ming-Chun Tien, Jonathon S. Barton, Arne Leinse, René G. Heideman, Daniel J. Blumenthal, and John E. Bowers. Ultra-low-loss high-aspect-ratio Si₃N₄ waveguides. *Opt. Express*, 19(4):3163–3174, Feb 2011.
- [BHL⁺16] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117:060504, Aug 2016.
- [BKRB15] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. Ion-trap measurements of electric-field noise near surfaces. *Rev. Mod. Phys.*, 87:1419–1482, Dec 2015.
- [BKW18] Ivan A. Boldin, Alexander Kraft, and Christof Wunderlich. Measuring anomalous heating in a planar ion trap with variable ion-surface separation. *Phys. Rev. Lett.*, 120:023201, Jan 2018.
- [Bla12] R. B. Blakestad. *Transport of Trapped-Ion Qubits within a Scalable Quantum Processor*. PhD thesis, University of Colorado, 2012.
- [BMB⁺98] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *J. Appl. Phys.*, 83(10):5025–5033, 1998.
- [BMCS16] C. D. Bruzewicz, R. McConnell, J. Chiaverini, and J. M. Sage. Scalable loading of a two-dimensional trapped-ion array. *Nat. Commun.*, 7:13005, 2016.
- [BMS⁺19] C. D. Bruzewicz, R. McConnell, J. Stuart, J. Chiaverini, and J. M. Sage. Dual-species, multi-qubit logic primitives for Ca⁺/Sr⁺ trapped-ion crystals. *npj Quantum Inf.*, 5:102, 2019.
- [BOV⁺09] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. High-fidelity transport of trapped-ion qubits through an **X**-junction trap array. *Phys. Rev. Lett.*, 102:153002, Apr 2009.
- [BOV⁺11] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, J. H. Wesenberg, M. J. Biercuk, D. Leibfried, and D. J. Wineland. Near-ground-state transport of trapped-ion qubits through a multidimensional array. *Phys. Rev. A*, 84:032314, Sep 2011.
- [BSC15] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. Measurement of ion motional heating rates over a range of trap frequencies and temperatures. *Phys. Rev. A*, 91:041402, Apr 2015.

- [BSWL10] A. H. Burrell, D. J. Szwer, S. C. Webster, and D. M. Lucas. Scalable simultaneous multiqubit readout with 99.99% single-shot fidelity. *Phys. Rev. A*, 81:040302, Apr 2010.
- [BURL⁺21] Maya Berlin-Udi, Clemens Matthiesen, P. N. Thomas Lloyd, Alberto Alonso, Crystal Noel, Chris Orme, Chang-Eun Kim, Vincenzo Lordi, and Hartmut Häffner. Changes in electric-field noise due to thermal transformation of a surface ion trap. *arXiv:2103.04482*, 2021.
- [Bur10] A. H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis, University of Oxford, 2010.
- [BWJ⁺18] S. Bourdeauducq, Whitequark, R. Jördens, Y. Sionneau, Enjoy-digital, C. Ballance, D. Nadlinger, T. Harty, J. Britton, D. Slichter, D. Rissinger, Mntng, R. Srinivas, Apatura-iris, S. Mackenzie, Z. Smith, Paweł K., M. Weber, K. Stevens, F. Held, and D. Leibbrandt. m-labs/artiq: 4.0, November 2018. <https://doi.org/10.5281/zenodo.1492176>.
- [BXN⁺17] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller. Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation. *Phys. Rev. X*, 7:041061, Dec 2017.
- [CBB⁺05] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Surface-electrode architecture for ion-trap quantum information processing. *Quantum Inf. Comp.*, 5:419, 2005.
- [Chw09] M. Chwalla. *Precision spectroscopy with $^{40}\text{Ca}^+$ ions in a Paul trap*. PhD thesis, Universität Innsbruck, 2009.
- [CLJ03] H. K. Cummins, G. Llewellyn, and J. A. Jones. Tackling systematic errors in quantum logic gates with composite rotations. *Phys. Rev. A*, 67:042308, Apr 2003.
- [CLO15] X. Chen, Y.-W. Lin, and B. C. Odom. Doppler amplification of motion of a trapped three-level ion. *New J. Phys.*, 17(4):043037, apr 2015.
- [CS96] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Phys. Rev. A*, 54:1098–1105, Aug 1996.
- [CS14] J. Chiaverini and J. M. Sage. Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range. *Phys. Rev. A*, 89:012318, Jan 2014.
- [DAW⁺12] S. C. Doret, J. M. Amini, K. Wright, C. Volin, T. Killian, A. Ozakin, D. Denison, H. Hayden, C.-S. Pai, R. E. Slusher, and A. W. Harter. Controlling trapping potentials and stray electric fields in a microfabricated ion trap through design and compensation. *New J. Phys.*, 14:073012, 2012.
- [DB00] M. Drewsen and A. Brøner. Harmonic linear Paul trap: Stability diagram and effective potentials. *Phys. Rev. A*, 62:045401, Sep 2000.

- [DBIW89] F. Diedrich, J. C. Bergquist, Wayne M. Itano, and D. J. Wineland. Laser cooling to the zero-point energy of motion. *Phys. Rev. Lett.*, 62:403–406, Jan 1989.
- [Deh68] H.G. Dehmelt. Radiofrequency spectroscopy of stored ions I: Storage. In D. R. Bates and Immanuel Estermann, editors, *Adv. At. Mol. Phys.*, volume 3, pages 53–72. Academic Press, 1968.
- [DeP21] M. DePalatis. Ion trapping periodic table, 2021. <https://mike.depalatis.net/ionptable/> [Online; accessed 20-May-2021].
- [Deu85] D. Deutsch. Quantum theory, the Church–Turing principle and the universal quantum computer. *Proc. R. Soc. Lond. A*, 400:97–117, 1985.
- [DGB⁺14] N. Daniilidis, S. Gerber, G. Bolloten, M. Ramm, A. Ransford, E. Ulin-Avila, I. Talukdar, and H. Häffner. Surface noise analysis using a single-ion sensor. *Phys. Rev. B*, 89:245435, Jun 2014.
- [DHK⁺83] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31(2):97–105, Jun 1983.
- [DHL⁺05] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. Double-pass acousto-optic modulator system. *Rev. Sci. Instrum.*, 76(6):063112, 2005.
- [DiV00] David P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.
- [DJ92] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proc. R. Soc. A*, 439(1907):553–558, 1992.
- [DL18] M. Delehaye and C. Lacroûte. Single-ion, transportable optical atomic clocks. *J. Mod. Opt.*, 65:622–639, 2018.
- [DLF⁺16] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536:63–66, 2016.
- [DMB⁺05] P. Dubé, A. A. Madej, J. E. Bernard, L. Marmet, J.-S. Boulanger, and S. Cundy. Electric quadrupole shift cancellation in single-ion optical frequency standards. *Phys. Rev. Lett.*, 95:033001, Jul 2005.
- [DMM⁺18] A. Didier, J. Millo, B. Marechal, C. Rocher, E. Rubiola, R. Lecomte, M. Ouisse, J. Delporte, C. Lacroûte, and Y. Kersalé. Ultracompact reference ultralow expansion glass cavity. *Appl. Opt.*, 57(22):6470–6473, Aug 2018.
- [DRBL⁺17] J. Davila-Rodriguez, F. N. Baynes, A. D. Ludlow, T. M. Fortier, H. Leopardi, S. A. Diddams, and F. Quinlan. Compact, thermal-noise-limited reference cavity for ultra-low-noise microwave generation. *Opt. Lett.*, 42(7):1277–1280, Apr 2017.
- [DRZ00] A. Debut, S. Randoux, and J. Zemmouri. Linewidth narrowing in Brillouin lasers: Theoretical analysis. *Phys. Rev. A*, 62:023803, Jul 2000.

- [DW75] D. G. Dalgoutte and C. D. W. Wilkinson. Thin grating couplers for integrated optics: An experimental and theoretical study. *Appl. Opt.*, 14(12):2983–2998, Dec 1975.
- [EGS⁺16] A. M. Eltony, D. Gangloff, M. Shi, A. Bylinskii, V. Vuletić, and I. L. Chuang. Technologies for trapped-ion quantum information systems. *Quantum Inf. Process.*, 15:5351–5383, 2016.
- [FAS⁺12] I. Fescenko, J. Alnis, A. Schliesser, C. Y. Wang, T. J. Kippenberg, and T. W. Hänsch. Dual-mode temperature compensation technique for laser stabilization to a crystalline whispering gallery mode resonator. *Opt. Express*, 20(17):19185–19193, Aug 2012.
- [Fey82] R. P. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21:467–488, 1982.
- [FHWJ19] M. Fan, C. A. Holliman, A. L. Wang, and A. M. Jayich. Laser cooling of radium ions. *Phys. Rev. Lett.*, 122:223001, Jun 2019.
- [GBP⁺19] S. Gundavarapu, G. M. Brodnik, M. Puckett, T. Huffman, D. Bose, R. Behunin, J. Wu, T. Qiu, C. Pinho, N. Chauhan, J. Nohava, P. T. Rakich, K. D. Nelson, M. Salit, and D. J. Blumenthal. Sub-Hertz fundamental linewidth photonic integrated Brillouin laser. *Nat. Photonics*, 13:60–67, 2019.
- [GFH⁺14] N. D. Guise, S. D. Fallek, H. Hayden, C.-S. Pai, C. Volin, K.R. Brown, J. T. Merrill, A. W. Harter, J. M. Amini, L. M. Lust, K. Muldoon, D. Carlson, and J. Budach. In-vacuum active electronics for microfabricated ion traps. *Rev. Sci. Instrum.*, 85:063101, 2014.
- [Gho95] Pradip K. Ghosh. *Ion Traps*. Oxford University Press, 1995.
- [GMM09] I. S. Grudinin, A. B. Matsko, and L. Maleki. Brillouin lasing with a CaF₂ whispering gallery mode resonator. *Phys. Rev. Lett.*, 102:043902, Jan 2009.
- [GORK⁺19] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga. Shortcuts to adiabaticity: Concepts, methods, and applications. *Rev. Mod. Phys.*, 91:045001, Oct 2019.
- [Gre16] Amy Greene. Experiments towards mitigation of motional heating in trapped ion quantum information processing. Master’s thesis, Massachusetts Institute of Technology, 2016.
- [Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.
- [GSW⁺06] J. Geng, S. Staines, Z. Wang, J. Zong, M. Blake, and S. Jiang. Highly stable low-noise Brillouin fiber laser with ultranarrow spectral linewidth. *IEEE Photon. Techno. Lett.*, 18(17):1813–1815, Sep. 2006.

- [GTL⁺16] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-fidelity universal gate set for $^9\text{Be}^+$ ion qubits. *Phys. Rev. Lett.*, 117:060505, Aug 2016.
- [HAB⁺14] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Phys. Rev. Lett.*, 113:220501, Nov 2014.
- [HBHB10] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. Trapped-ion probing of light-induced charging effects on dielectrics. *New J. Phys.*, 12:093035, 2010.
- [HCW⁺12] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-fold reduction of electric-field noise in an ion trap cleaned with in situ argon-ion-beam bombardment. *Phys. Rev. Lett.*, 109:103001, Sep 2012.
- [HCW⁺13] D. A. Hite, Y. Colombe, A. C. Wilson, D. T. C. Allcock, D. Leibfried, D. J. Wineland, and D. P. Pappas. Surface science for improved ion traps. *MRS Bull.*, 38(10):826–833, 2013.
- [HH15] Paul Horowitz and Winfield Hill. *The Art of Electronics*. Cambridge University Press, 3rd edition, 2015.
- [HH20] Paul Horowitz and Winfield Hill. *The Art of Electronics: The x-Chapters*. Cambridge University Press, 1st edition, 2020.
- [HLBH11] M. D. Hughes, B. Lekitsch, J. Broersma, and W. Hensinger. Microfabricated ion traps. *Contemp. Phys.*, 52:505–529, 2011.
- [HMK⁺17] D. A. Hite, K. S. McKay, S. Kotler, D. Leibfried, D. J. Wineland, and D. P. Pappas. Measurements of trapped-ion heating rates with exchangeable surfaces in close proximity. *MRS Adv.*, 2(41):2189–2197, 2017.
- [HOS⁺06] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk. T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation. *Appl. Phys. Lett.*, 88(3):034101, 2006.
- [Hou08] M. G. House. Analytic model for electrostatic fields in surface-electrode ion traps. *Phys. Rev. A*, 78:033402, Sep 2008.
- [How00] D. A. Howe. The total deviation approach to long-term characterization of frequency stability. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 47:1102–1110, 2000.
- [HPKM12] N. Herschbach, K. Pyka, J. Keller, and Mehlstäubler. Linear Paul trap design for an optical clock with Coulomb crystals. *Appl. Phys. B*, 107:891–906, 2012.
- [HRB08] H. Häffner, C.F. Roos, and R. Blatt. Quantum computing with trapped ions. *Phys. Rep.*, 469(4):155–203, 2008.

- [HS06] J. P. Home and A. M. Steane. Electrode configurations for fast separation of trapped ions. *Quantum Inf. Comp.*, 6:289–325, 2006.
- [HYO⁺08] D. Hucul, M. Yeo, S. Olmschenk, C. Monroe, W. K. Hensinger, and J. Rabchuk. On the transport of atomic ions in linear and multidimensional ion trap arrays. *Quantum Inf. Comput.*, 8(6):501–578, 2008.
- [IS72] E. P. Ippen and R. H. Stolen. Stimulated Brillouin scattering in optical fibers. *Appl. Phys. Lett.*, 21(11):539–541, 1972.
- [ISK⁺21] M. Ivory, W. J. Setzer, N. Karl, H. McGuinness, C. DeRose, M. Blain, D. Stick, M. Gehl, and L. P. Parazzoli. Integrated optical addressing of a trapped ytterbium ion. *arXiv:2011.12376*, 2021.
- [ITU11] Y. Ibaraki, U. Tanaka, and S. Urabe. Detection of parametric resonance of trapped ions for micromotion compensation. *Appl. Phys. B*, 105:219–223, 2011.
- [Jam98] D. F. V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Appl. Phys. B*, 66(2):181–190, Feb 1998.
- [JLL⁺11] Y. Y. Jiang, A. D. Ludlow, A. D. Lemke, R. W. Fox, J. A. Sherman, L.-S. Ma, and C. W. Oates. Making optical atomic clocks more stable with 10^{-16} -level laser stabilization. *Nat. Photonics*, 5:158–161, 2011.
- [JMS⁺00] S. R. Jefferts, D. M. Meekhof, J. H. Shirley, M. Stepanovic, and T. E. Parker. Accuracy results from NIST-F1 laser-cooled cesium primary frequency standard. In *Proceedings of the IEEE/EIA International Frequency Control Symposium and Exhibition*, pages 714–717, 2000.
- [Jun05] Walter G. Jung. *Op Amp Applications Handbook*. Newnes/Elsevier, 2005.
- [KBZ⁺09] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and Blatt R. Deterministic entanglement of ions in thermal states of motion. *New J. Phys.*, 11:023002, 2009.
- [KGV⁺17] S. B. Koller, J. Grotti, St. Vogt, A. Al-Masoudi, S. Dörscher, S. Häfner, U. Sterr, and Ch. Lisdat. Transportable optical lattice clock with 7×10^{-17} uncertainty. *Phys. Rev. Lett.*, 118:073601, 2017.
- [KHR⁺16] M. Kumph, C. Henkel, P. Rabl, M. Brownnutt, and R. Blatt. Electric-field noise above a thin dielectric layer on metal electrodes. *New J. Phys.*, 18:023020, 2016.
- [Kie01] D. Kielpinski. *Entanglement and Decoherence in a Trapped-Ion Quantum Register*. PhD thesis, University of Colorado, 2001.
- [KKM⁺00] D. Kielpinski, B. E. King, C. J. Myatt, C. A. Sackett, Q. A. Turchette, W. M. Itano, C. Monroe, D. J. Wineland, and W. H. Zurek. Sympathetic cooling of trapped ions for quantum logic. *Phys. Rev. A*, 61:032310, Feb 2000.

- [KLS⁺20] V. Kaushal, B. Lekitsch, A. Stahl, J. Hilder, D. Pijn, C. Schmiegelow, A. Bermudez, M. Müller, F. Schmidt-Kaler, and U. Poschinger. Shuttling-based trapped-ion quantum information processing. *AVS Quantum Sci.*, 2(1):014101, 2020.
- [KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, 2002.
- [KPC⁺13] I. V. Kabakova, R. Pant, D.-Y. Choi, S. Debbarma, B. Luther-Davies, S. J. Madden, and B. J. Eggleton. Narrow linewidth Brillouin laser based on chalcogenide photonic chip. *Opt. Lett.*, 38(17):3208–3211, Sep 2013.
- [KRS⁺14] H. Kaufmann, T. Ruster, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger. Dynamics and control of fast ion crystal splitting in segmented Paul traps. *New J. Phys.*, 16:073012, 2014.
- [KRS⁺17] H. Kaufmann, T. Ruster, C. T. Schmiegelow, M. A. Luda, V. Kaushal, J. Schulz, D. von Lindenfels, F. Schmidt-Kaler, and U. G. Poschinger. Fast ion swapping for quantum-information processing. *Phys. Rev. A*, 95:052319, 2017.
- [KSC10] Andrey Kobayakov, Michael Sauer, and Dipak Chowdhury. Stimulated Brillouin scattering in optical fibers. *Adv. Opt. Photon.*, 2(1):1–59, Mar 2010.
- [KSNH⁺17] E. Kim, A. Safavi-Naini, D. A. Hite, K. S. McKay, D. P. Pappas, P. F. Weck, and H. R. Sadeghpour. Electric-field noise from carbon-adsorbate diffusion on a Au(110) surface: First-principles calculations and experiments. *Phys. Rev. A*, 95:033407, Mar 2017.
- [Lab08] J. Labaziewicz. *High Fidelity Quantum Gates with Ions in Cryogenic Microfabricated Ion Traps*. PhD thesis, Massachusetts Institute of Technology, 2008.
- [LBC⁺16] W. Loh, J. Becker, D. C. Cole, A. Coillet, F. N. Baynes, S. B. Papp, and S. A. Diddams. A microrod-resonator Brillouin laser with 240 Hz absolute linewidth. *New J. Phys.*, 18:045001, 2016.
- [LBMW03] Dietrich Leibfried, Rainer Blatt, Christopher Monroe, and David Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75(1):281, 2003.
- [LCL⁺12] H. Lee, T. Chen, J. Li, K. Y. Yang, S. Jeon, O. Painter, and K. J. Vahala. Chemically etched ultrahigh-Q wedge-resonator on a silicon chip. *Nat. Photonics*, 6:369–373, 2012.
- [LGB⁺15] W. Loh, A. A. S. Green, F. N. Baynes, D. C. Cole, F. J. Quinlan, H. Lee, K. J. Vahala, S. B. Papp, and S. A. Diddams. Dual-microcavity narrow-linewidth Brillouin laser. *Optica*, 2(3):225–232, Mar 2015.
- [LGL⁺08] J. Labaziewicz, Y. Ge, D. R. Leibbrandt, S. X. Wang, R. Shewmon, and I. L. Chuang. Temperature dependence of electric field noise above gold surfaces. *Phys. Rev. Lett.*, 101:180602, 2008.

- [LGRS04] V. Letchumanan, P. Gill, E. Riis, and A. G. Sinclair. Optical Ramsey spectroscopy of a single trapped $^{88}\text{Sr}^+$ ion. *Phys. Rev. A*, 70:033419, Sep 2004.
- [LHC11] Guang Hao Low, Peter F. Herskind, and Isaac L. Chuang. Finite-geometry models of electric field noise from patch potentials in ion traps. *Phys. Rev. A*, 84:053425, Nov 2011.
- [LLS⁺19] J. Lim, W. Liang, A. A. Savchenkov, A. B. Matsko, L. Maleki, and C. W. Wong. Probing 10 μK stability and residual drifts in the cross-polarization dual-mode stabilization of single-crystal ultrahigh-Q optical resonators. *Light Sci. Appl.*, 8:1, 2019.
- [LoEOoSI16] Sandia National Laboratories, United States. Department of Energy. Office of Scientific, and Technical Information. *High Optical Access Trap 2.0*. United States. Department of Energy., 2016.
- [LPD15] William Loh, Scott B. Papp, and Scott A. Diddams. Noise and dynamics of stimulated-Brillouin-scattering microresonator lasers. *Phys. Rev. A*, 91:053843, May 2015.
- [LRH⁺04] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J.-P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-selective photoionization for calcium ion trapping. *Phys. Rev. A*, 69:012711, Jan 2004.
- [LSR⁺20] W. Loh, J. Stuart, D. Reens, C. D. Bruzewicz, D. Braje, J. Chiaverini, P. W. Juodawlakis, J. M. Sage, and R. McConnell. Operation of an optical atomic clock with a Brillouin laser subsystem. *Nature*, 588:244–249, 2020.
- [LTBR11] D. R. Leibbrandt, M. J. Thorpe, J. C. Bergquist, and T. Rosenband. Field-test of a robust, portable, frequency-stable laser. *Opt. Express*, 19(11):10278–10286, May 2011.
- [LWF⁺17] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger. Blueprint for a microwave trapped ion quantum computer. *Sci. Adv.*, 3(2):e1601540, 2017.
- [LYOJ19] W. Loh, S. Yegnanarayanan, F. O’Donnell, and P. W. Juodawlakis. Ultra-narrow linewidth Brillouin laser with nanokelvin temperature self-referencing. *Optica*, 6(2):152–159, Feb 2019.
- [LYS07] D. Leibbrandt, B. Yurke, and R. Slusher. Modeling ion trap thermal noise decoherence. *Quantum Inf. Comput.*, 7(1):52–72, January 2007.
- [MBCS15] Robert McConnell, Colin Bruzewicz, John Chiaverini, and Jeremy Sage. Reduction of trapped-ion anomalous heating by in situ surface plasma cleaning. *Phys. Rev. A*, 92:020302, Aug 2015.
- [MBM⁺16] Karan K. Mehta, Colin D. Bruzewicz, Robert McConnell, Rajeev J. Ram, Jeremy M. Sage, and John Chiaverini. Integrated optical addressing of an ion qubit. *Nat. Nanotechnol.*, 11(12):1066, 2016.

- [MDZ⁺12] A. A. Madej, P. Dubé, Z. Zhou, J. E. Bernard, and M. Gertsch. $^{88}\text{Sr}^+$ 445-THz single-ion reference at the 10^{-17} level via control and cancellation of systematic uncertainties and its measurement against the SI second. *Phys. Rev. Lett.*, 109:203002, Nov 2012.
- [MEB⁺14] K. K. Mehta, A. M. Eltony, C. D. Bruzewicz, I. L. Chuang, R. J. Ram, J. M. Sage, and J. Chiaverini. Ion traps fabricated in a CMOS foundry. *Appl. Phys. Lett.*, 105:044103, 2014.
- [MJYH19] L.-S. Ma, P. Junger, J. Ye, and J. L. Hall. Delivering the same optical frequency at two places: Accurate cancellation of phase noise introduced by an optical fiber or other time-varying path. *Opt. Lett.*, 19(2):1777–1779, 2019.
- [MK13] Christopher Monroe and Jungsang Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.
- [MKW⁺16] M. Mielenz, H. Kalis, M. Wittmer, F. Hakelberg, U. Warring, R. Schmied, M. Blain, P. Maunz, D. L. Moehring, D. Leibfried, and T. Schaetz. Arrays of individually controlled ions suitable for two-dimensional quantum simulations. *Nat. Commun.*, 7:11839, 2016.
- [MMB98] A. Madej, L. Marmet, and J. Bernard. Rb atomic absorption line reference for single Sr^+ laser cooling systems. *Appl. Phys. B*, 67:229–234, 1998.
- [MMK⁺95] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Phys. Rev. Lett.*, 75:4714–4717, Dec 1995.
- [MNM⁺16] Thomas Monz, Daniel Nigg, Esteban A. Martinez, Matthias F. Brandl, Philipp Schindler, Richard Rines, Shannon X. Wang, Isaac L. Chuang, and Rainer Blatt. Realization of a scalable Shor algorithm. *Science*, 351(6277):1068–1070, 2016.
- [MR17] Karan K. Mehta and Rajeev J. Ram. Precise and diffraction-limited waveguide-to-free-space focusing gratings. *Sci. Rep.*, 7(1):2019, 2017.
- [MRR⁺14] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Phys. Rev. A*, 89:022317, Feb 2014.
- [MSJ⁺21] Galan Moody, Volker J. Sorger, Paul W. Juodawlkis, William Loh, Cheryl Sorace-Agaskar, Marcelo Davanco, Lin Chang, John E. Bowers, Niels Quack, Christophe Galland, Igor Aharonovich, M. Wolff, C. Schuck, Neil Sinclair, Marko Lončar, Tin Komljenovic, David Weld, Shayan Mookherjee, Sonia Buckley, Marina Radulaski, Stephan Reitzenstein, Benjamin Pingault, Bartholomeus Machielse Debsuvra Mukhopadhyay, Alexey Akimov, Aleksei Zheltikov, Girish S. Agarwal, Kartik Srinivasan, Juanjuan Lu, Hong X. Tang, Wentao Jiang, Timothy P. McKenna, Amir H. Safavi-Naeini, Stephan Steinhilber, Ali W. Elshaari, Val Zwiller, Paul S. Davids, Nicholas Martinez, Michael Gehl, John Chiaverini, Karan K. Mehta, Jacqueline Romero, Navin B. Lingaraju, Andrew M. Weiner, Daniel Peace, Robert Ceand Mirko Lobino, Eleni

- Diamanti, Luis Trigo Vidarte, and Ryan M. Camac. Roadmap on integrated quantum photonics. *arXiv:2102.03323*, 2021.
- [MSW⁺08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Phys. Rev. Lett.*, 100:200502, May 2008.
- [MvdS99] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*. Springer, 1st edition, 1999.
- [MZM⁺20] K. K. Mehta, C. Zhang, M. Malinowski, T.-L. Nguyen, M. Stadler, and J. P. Home. Integrated optical multi-ion quantum logic. *Nature*, 586:533–537, 2020.
- [NBUM⁺19] Crystal Noel, Maya Berlin-Udi, Clemens Matthiesen, Jessica Yu, Yi Zhou, Vincenzo Lordi, and Hartmut Häffner. Electric-field noise from thermally activated fluctuators in a surface ion trap. *Phys. Rev. A*, 99:063427, Jun 2019.
- [NCH⁺15] T. L. Nicholson, S. L. Campbell, R. B. Hutson, G. E. Marti, B. J. Bloom, R. L. McNally, W. Zhang, M. D. Barrett, M. S. Safronova, G. F. Strouse, W. L. Tew, and J. Ye. Systematic evaluation of an atomic clock at 2×10^{-18} uncertainty. *Nat. Commun.*, 6:6896, 2015.
- [NCP⁺20] Yunseong Nam, Jwo-Sy Chen, Neal C Pienti, Kenneth Wright, Conor Delaney, Dmitri Maslov, Kenneth R Brown, Stewart Allen, Jason M Amini, Joel Apisdorf, et al. Ground-state energy estimation of the water molecule on a trapped-ion quantum computer. *npj Quantum Inf.*, 6:33, 2020.
- [NH12] A. H. Nizamani and W. K. Hensinger. Optimum electrode configurations for fast ion separation in microfabricated surface ion traps. *App. Phys. B*, 106:327 – 338, 2012.
- [Nic06] T. L. Nicholson. *A new record in atomic clock performance*. PhD thesis, University of Colorado, 2006.
- [NMW⁺12] T. L. Nicholson, M. J. Martin, J. R. Williams, B. J. Bloom, M. Bishof, M. D. Swallows, S. L. Campbell, and J. Ye. Comparison of two independent Sr optical clocks with 1×10^{-17} stability at 10^3 s. *Phys. Rev. Lett.*, 109:230801, Dec 2012.
- [NP73] Laurence W. Nagel and D. O. Pederson. SPICE (simulation program with integrated circuit emphasis). Technical Report UCB/ERL M382, EECS Department, University of California, Berkeley, Apr 1973.
- [NRL⁺00] H. C. Nägerl, C. Roos, D. Leibfried, H. Rohde, G. Thalhammer, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Investigating a qubit candidate: Spectroscopy on the $S_{1/2}$ to $D_{5/2}$ transition of a trapped calcium ion in a linear Paul trap. *Phys. Rev. A*, 61:023405, Jan 2000.
- [NSSA⁺20] R. J. Niffenegger, J. Stuart, C. Sorace-Agaskar, D. Kharas, S. Bramhavar, C. D. Bruzewicz, W. Loh, R. T. Maxson, R. McConnell, D. Reens, G. N.

- West, J. M. Sage, and J. Chiaverini. Integrated multi-wavelength control of an ion qubit. *Nature*, 586:538–542, 2020.
- [NTHS93] P.-A. Nicati, K. Toyama, S. Huang, and H. J. Shaw. Temperature effects in a Brillouin fiber ring laser. *Opt. Lett.*, 18(24):2123–2125, Dec 1993.
- [NYD89] W. Nagourney, N. Yu, and H. Dehmelt. Laser spectroscopy of a single barium ion using “shelving”. In A. De Marchi, editor, *Frequency Standards and Metrology*, pages 312–318. Springer, Berlin, Heidelberg, 1989.
- [OBK⁺18] N. T. Otterstrom, R. O. Behunin, E. A. Kittlaus, Z. Wang, and P. T. Rakich. A silicon Brillouin laser. *Science*, 360(6393):1113–1116, 2018.
- [OM01] Mário H. Oliveira and José A. Miranda. Biot-Savart-like law in electrostatics. *Eur. J. of Phys.*, 22(1):31–38, jan 2001.
- [Pau90] Wolfgang Paul. Electromagnetic traps for charged and neutral particles. *Rev. Mod. Phys.*, 62:531–540, Jul 1990.
- [PDF⁺21] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. Demonstration of the trapped-ion quantum ccd computer architecture. *Nature*, 592:209–213, 2021.
- [PH17] Darrell W. Pepper and Juan C. Heinrich. *The Finite Element Method*. CRC Press, 3rd edition, 2017.
- [PHK⁺18] G Pagano, P W Hess, H B Kaplan, W L Tan, P Richerme, P Becker, A Kyprianidis, J Zhang, E Birkelbaw, M R Hernandez, Y Wu, and C Monroe. Cryogenic trapped-ion system for large scale quantum simulation. *Quantum Sci. and Tech.*, 4(1):014004, oct 2018.
- [PMGP⁺15] M. Palmero, S. Martínez-Garaot, U. G. Poschinger, A. Ruschhaupt, and J. G. Muga. Fast separation of two trapped ions. *New J. Phys.*, 17:093031, 2015.
- [Pob07] Frank Pobell. *Matter and Methods at Low Temperatures*. Springer, 3rd edition, 2007.
- [Pre18] J. Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- [Qi21] L. Qi. Optimizing ion-shuttling operations in trapped-ion quantum computers. B.S. thesis, Massachusetts Institute of Technology, 2021.
- [RBKD⁺02] M.A. Rowe, A. Ben-Kish, B. DeMarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W.M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D.J. Wineland. Transport of quantum states and separation of ions in a dual rf ion trap. *Quantum Inf. Comp.*, 2:257, 2002.
- [RHS⁺20] Zak David Romaszko, Seokjun Hong, Martin Siegele, Reuben Kahan Puddy, Foni Raphaël Lebrun-Gallagher, Sebastian Weidt, and Winfried Karl Hensinger. Engineering of microfabricated ion traps and integration of advanced on-chip features. *Nat. Rev. Phys.*, 2:285–299, 2020.

- [RKL⁺17] T. Ruster, H. Kaufmann, M. A. Luda, V. Kaushal, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger. Entanglement-based dc magnetometry with separated ions. *Phys. Rev. X*, 7:031050, Sep 2017.
- [RLB⁺06] R. Reichle, D. Leibfried, R. B. Blakestad, J. Britton, J. D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Transport dynamics of single ions in segmented microstructured Paul trap arrays. *Fortschr. Phys.*, 54(8-10):666–685, 2006.
- [Roo00] Christian Roos. *Controlling the quantum state of trapped ions*. PhD thesis, Universität Innsbruck, 2000.
- [RWK⁺14] T. Ruster, C. Warschburger, H. Kaufmann, C. T. Schmiegelow, A. Walther, M. Hettrich, A. Pfister, V. Kaushal, F. Schmidt-Kaler, and U. G. Poschinger. Experimental realization of fast ion separation in segmented Paul traps. *Phys. Rev. A*, 90:033410, Sep 2014.
- [SABK⁺18] C. Sorace-Agaskar, S. Bramhavar, D. Kharas, W. Loh, P. W. Juodawlakis, J. Chiaverini, and J. M. Sage. Multi-layer integrated photonics from the ultraviolet to the infrared. *Proc. of SPIE*, 10510:105100D, 2018.
- [Sag06] J. Sage. *Optical Production of Ultracold Polar Molecules*. PhD thesis, Yale University, 2006.
- [SAKY⁺19] Cheryl Sorace-Agaskar, Dave Kharas, Siva Yegnanarayanan, Ryan Maxson, Gavin Neal West, William Loh, Suraj Bramhavar, Rajeev J Ram, John Chiaverini, Jeremy Sage, and Paul Juodawlakis. Versatile silicon nitride and alumina integrated photonic platforms for the ultraviolet to short-wave infrared. *IEEE J. Sel. Top. Quantum Electron.*, 25(5):1–15, 2019.
- [SAMI⁺11] Y. G. Shee, M. H. Al-Mansoori, A. Ismail, S. Hitam, and M. A. Mahdi. Multiwavelength Brillouin-erbium fiber laser with double-Brillouin-frequency spacing. *Opt. Express*, 19(3):1699–1706, Jan 2011.
- [ŠB02] Marek Šašura and Vladimír Bužek. Cold trapped ions as quantum information processors. *J. Mod. Opt.*, 49(10):1593–1647, 2002.
- [SBJ⁺98] R. Storz, C. Braxmaier, K. Jäck, O. Pradl, and S. Schiller. Ultrahigh long-term dimensional stability of a sapphire cryogenic optical resonator. *Opt. Lett.*, 23(13):1031–1033, Jul 1998.
- [SBK⁺21] R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T. C. Allcock, and D. H. Slichter. High-fidelity laser-free universal control of two trapped ion qubits. *arXiv:2102.12533*, 2021.
- [SBT⁺18] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas. Fast quantum logic gates with trapped-ion qubits. *Nature*, 555(7694):75, 2018.
- [SBTL15] V. M. Schäfer, C. J. Ballance, C. J. Tock, and D. M. Lucas. Optical injection and spectral filtering of high-power ultraviolet laser diodes. *Opt. Lett.*, 40(18):4265–4268, Sep 2015.

- [Sch04] Thomas Schneider. *Nonlinear Optics in Telecommunications*. Springer-Verlag Berlin Heidelberg, 1st edition, 2004.
- [SCR⁺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, and D. J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Phys. Rev. Lett.*, 96:253003, 2006.
- [SDF⁺11] E. A. Salim, J. DeNatale, D. M. Farkas, K. M. Hudek, S. E. McBride, J. Michalchuk, R. Mihailovich, and D. Z. Anderson. Compact, microchip-based systems for practical applications of ultracold atoms. *Quantum Inf. Process.*, 10:975–994, 2011.
- [SGS⁺18] J. A. Sedlacek, A. Greene, J. Stuart, R. McConnell, C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. Distance scaling of electric-field noise in a surface-electrode ion trap. *Phys. Rev. A*, 97:020302, Feb 2018.
- [Sho94] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, Nov 1994.
- [SHO⁺06] D. Stick, W. K. Hensinger, S. Olmschenk, M. J. Madsen, K. Schwab, and C. Monroe. Ion trap in a semiconductor chip. *Nat. Phys.*, 2:36–39, 2006.
- [SIS⁺21] W. J. Setzer, M. Ivory, O. Slobodyan, J. W. Van Der Wall, L. P. Parazzoli, D. Stick, M. Gehl, M. Blain, R. R. Kay, and H. McGuinness. Fluorescence detection of a trapped ion with a monolithically integrated single-photon-counting avalanche diode. *arXiv:2105.01235*, 2021.
- [SKC12] J. M. Sage, A. J. Kerman, and J. Chiaverini. Loading of a surface-electrode ion trap from a remote, precooled source. *Phys. Rev. A*, 86:013417, Jul 2012.
- [SKHR⁺03] Ferdinand Schmidt-Kaler, Hartmut Häffner, Mark Riebe, Stephan Gulde, Gavin P. T. Lancaster, Thomas Deuschle, Christoph Becher, Christian F. Roos, Jürgen Eschner, and Rainer Blatt. Realization of the Cirac-Zoller controlled-not quantum gate. *Nature*, 422:408–411, 2003.
- [SM00] Anders Sørensen and Klaus Mølmer. Entanglement and quantum computation with ions in thermal motion. *Phys. Rev. A*, 62:022311, Jul 2000.
- [SNKW⁺13] A. Safavi-Naini, E. Kim, P. F. Weck, P. Rabl, and H. R. Sadeghpour. Influence of monolayer contamination on electric-field-noise heating in ion traps. *Phys. Rev. A*, 87:023421, Feb 2013.
- [SNN⁺20] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance. High-rate, high-fidelity entanglement of qubits across an elementary quantum network. *Phys. Rev. Lett.*, 124:110501, Mar 2020.
- [SNRWS11] A. Safavi-Naini, P. Rabl, P. F. Weck, and H. R. Sadeghpour. Microscopic model of electric-field-noise heating in ion traps. *Phys. Rev. A*, 84:023412, Aug 2011.

- [SPB⁺19] J. Stuart, R. Panock, C.D. Bruzewicz, J.A. Sedlacek, R. McConnell, I.L. Chuang, J.M. Sage, and J. Chiaverini. Chip-integrated voltage sources for control of trapped ions. *Phys. Rev. Appl.*, 11:024010, Feb 2019.
- [SPM⁺10] Kilian Singer, Ulrich Poschinger, Michael Murphy, Peter Ivanov, Frank Ziesel, Tommaso Calarco, and Ferdinand Schmidt-Kaler. Colloquium: Trapped ions as quantum bits: Essential numerical tools. *Rev. Mod. Phys.*, 82:2609–2632, Sep 2010.
- [SSJ17] U. I. Safronova, M. S. Safronova, and W. R. Johnson. Forbidden $M1$ and $E2$ transitions in monovalent atoms and ions. *Phys. Rev. A*, 95:042507, Apr 2017.
- [SSK12] C. T. Schmiegelow and F. Schmidt-Kaler. Light with orbital angular momentum interacting with trapped ions. *Eur. Phys. J. D*, 66:157, 2012.
- [SSL⁺18] J. A. Sedlacek, J. Stuart, W. Loh, R. McConnell, C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. Method for determination of technical noise contributions to ion motional heating. *J. Appl. Phys.*, 124(21):214904, 2018.
- [SSS⁺18] J. A. Sedlacek, J. Stuart, D. H. Slichter, C. D. Bruzewicz, R. McConnell, J. M. Sage, and J. Chiaverini. Evidence for multiple mechanisms underlying surface electric-field noise in ion traps. *Phys. Rev. A*, 98:063430, Dec 2018.
- [STB⁺11] D. V. Strekalov, R. J. Thompson, L. M. Baumgartel, I. S. Grudinin, and N. Yu. Temperature measurement and stabilization in a birefringent whispering gallery mode resonator. *Opt. Express*, 19(15):14495–14501, Jul 2011.
- [SVB⁺14] G. Shu, G. Vittorini, A. Buikema, C. S. Nichols, C. Volin, D. Stick, and Kenneth R. Brown. Heating rates and ion-motion control in a Y-junction surface-electrode trap. *Phys. Rev. A*, 89:062308, Jun 2014.
- [SZE91] S. P. Smith, F. Zarinetchi, and S. Ezekiel. Narrow-linewidth stimulated Brillouin fiber laser and applications. *Opt. Lett.*, 16(6):393–395, Mar 1991.
- [TGLW09] T. G. Tiecke, S. D. Gensemer, A. Ludewig, and J. T. M. Walraven. High-flux two-dimensional magneto-optical-trap source for cold lithium atoms. *Phys. Rev. A*, 80:013409, Jul 2009.
- [TKK⁺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, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Phys. Rev. A*, 61:063418, 2000.
- [TLF⁺12] K. H. Tow, Y. Léguillon, S. Fresnel, P. Besnard, L. Brilland, D. Méchin, D. Trégoat, J. Troles, and P. Toupin. Linewidth-narrowing and intensity noise reduction of the 2nd order stokes component of a low threshold Brillouin laser made of Ge₁₀As₂₂Se₆₈ chalcogenide fiber. *Opt. Express*, 20(26):B104–B109, Dec 2012.
- [TMA⁺12] U. Tanaka, K. Masuda, Y. Akimoto, K. Koda, Y. Ibaraki, and S. Urabe. Micromotion compensation in a surface electrode trap by parametric excitation of trapped ions. *Appl. Phys. B*, 107:907–912, 2012.

- [Tod20] S. L. Todaro. *Scalable state detection and fast transport of trapped-ion qubits for quantum computing*. PhD thesis, University of Colorado, 2020.
- [TVM⁺21] S. L. Todaro, V. B. Verma, K. C. McCormick, D. T. C. Allcock, R. P. Mirin, D. J. Wineland, S. W. Nam, A. C. Wilson, D. Leibfried, and D. H. Slichter. State readout of a trapped ion qubit using a trap-integrated superconducting photon detector. *Phys. Rev. Lett.*, 126:010501, Jan 2021.
- [Vah03] K. Vahala. Optical microcavities. *Nature*, 424:839–846, 2003.
- [VCLB06] K. Vant, J. Chiaverini, W. Lybarger, and D. J. Berkeland. Photoionization of strontium for trapped-ion quantum information processing. *arXiv:0607055v1*, 2006.
- [vMMG⁺20] Martin W. van Mourik, Esteban A. Martinez, Lukas Gerster, Pavel Hrmo, Thomas Monz, Philipp Schindler, and Rainer Blatt. Coherent rotations of qubits within a surface ion-trap quantum computer. *Phys. Rev. A*, 102:022611, Aug 2020.
- [WAF⁺13] Kenneth Wright, Jason M. Amini, Daniel L. Faircloth, Curtis Volin, S. Charles Doret, Harley Hayden, C.-S. Pai, David W. Landgren, Douglas Denison, Tyler Killian, Richard E. Slusher, and Alexa W. Harter. Reliable transport through a microfabricated X-junction surface-electrode ion trap. *New J. Phys.*, 15(3):033004, 2013.
- [WAS⁺14] W. Weng, J. D. Anstie, T. M. Stace, G. Campbell, F. N. Baynes, and A. N. Luiten. Nano-kelvin thermometry and temperature control: Beyond the thermal noise limit. *Phys. Rev. Lett.*, 112:160801, Apr 2014.
- [WBD⁺19] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pisenti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducore, A. Blinov, S. M. Kreikemeier, V. Chaplina, M. Keesan, C. Monroe, and J. Kim. Benchmarking an 11-qubit quantum computer. *Nat. Commun.*, 10:5464, 2019.
- [Wes08] J. H. Wesenberg. Electrostatics of surface-electrode ion traps. *Phys. Rev. A*, 78:063410, Dec 2008.
- [WHLL⁺11] Shannon X Wang, Guang Hao Low, Nathan S Lachenmyer, Yufei Ge, Peter F Herskind, and Isaac L Chuang. Laser-induced charging of microfabricated ion traps. *J. Appl. Phys.*, 110(10):104901, 2011.
- [WI79] D. J. Wineland and Wayne M. Itano. Laser cooling of atoms. *Phys. Rev. A*, 20:1521–1540, Oct 1979.
- [WI81] D.J. Wineland and Wayne M. Itano. Spectroscopy of a single Mg^+ ion. *Phys. Lett. A*, 82(2):75–78, 1981.
- [WJE⁺20] Yong Wan, Robert Jördens, Stephen D. Erickson, Jenny J. Wu, Ryan Bowler, Ting Rei Tan, Pan-Yu Hou, David J. Wineland, Andrew C. Wilson, and Dietrich Leibfried. Ion transport and reordering in a 2D trap array. *Adv. Quantum Technol.*, 3(11):2000028, 2020.

- [WLG⁺10] S. X. Wang, J. Labaziewicz, Y. Ge, R. Shewmon, and I. L. Chuang. Demonstration of a quantum logic gate in a cryogenic surface-electrode ion trap. *Phys. Rev. A*, 81:062332, Jun 2010.
- [WLK⁺19] Gavin N. West, William Loh, Dave Kharas, Cheryl Sorace-Agaskar, Karan K. Mehta, Jeremy Sage, John Chiaverini, and Rajeev J. Ram. Low-loss integrated photonics for the blue and ultraviolet regime. *APL Photonics*, 4, 2019.
- [WLL⁺11] S. X. Wang, G. H. Low, N. S. Lachenmyer, Y. Ge, P. F. Herskind, and I. L. Chuang. Laser-induced charging of microfabricated ion traps. *J. Appl. Phys.*, 110:104901, 2011.
- [WM02] Guy K. White and Philip J. Meeson. *Experimental Techniques in Low-Temperature Physics*. Oxford University Press, 4th edition, 2002.
- [WMI⁺98] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *J. Res. Natl. Inst. Stand. Technol.*, 103:259–328, 1998.
- [WUZ⁺17] Ye Wang, Mark Um, Junhua Zhang, Shuoming An, Ming Lyu, Jing-Ning Zhang, L-M Duan, Dahyun Yum, and Kihwan Kim. Single-qubit quantum memory exceeding ten-minute coherence time. *Nat. Photonics*, 11(10):646, 2017.
- [WZR⁺12] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger. Controlling fast transport of cold trapped ions. *Phys. Rev. Lett.*, 109:080501, Aug 2012.
- [YCIB99] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist. Visible lasers with subhertz linewidths. *Phys. Rev. Lett.*, 82:3799–3802, May 1999.