Laser-free operations in a mixed-species trapped ion ${\bf processor}$

by

Hannah Marie Knaack

B.S., Harvey Mudd College, 2017M.S., University of Colorado Boulder, 2019

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2024

Committee Members:
Daniel H. Slichter, Chair
Dietrich Leibfried
Scott Glancy
Heather J. Lewandowski
J. Mathias Weber

Knaack, Hannah Marie (Ph.D., Physics)

Laser-free operations in a mixed-species trapped ion processor

Thesis directed by Dr. Daniel H. Slichter

Trapped ions are a leading platform for quantum computing, but, like all other candidate platforms, the techniques that will enable scaling to larger systems are still an area of active research. Laser-free control operations for ion-based qubits offer many advantages for scaling. The microwave and radiofrequency currents used to perform these operations are cheaper and more straightforward to generate, control, and integrate with trapped ions than lasers. They offer power-efficient parallelization of gate operations relative to laser-based gates, and are intrinsically free from the photon scattering, laser phase noise, and beam pointing errors that are the primary sources of infidelity in laser-based control. Additionally, operating with multiple co-trapped ion species (e.g., a 'data' and a 'helper' species) offers many advantages for computing, as well as for exploring fundamental physics and precision measurement by extending quantum control techniques to more exotic charged particles. This thesis demonstrates the first use of multiple ion species in a trap designed for laser-free qubit control and entangling operations, and describes a theoretical framework and early experimental results for mixed-species quantum logic operations based on magnetic field gradients. Laser-free mixed-species quantum information processing offers many potential advantages for scaling to larger systems while avoiding some common error sources and challenges intrinsic to laser-based quantum information with trapped ions.

Dedication

For my mom, who taught me how to make things, and my dad, who taught me how to fix them.

Acknowledgements

It has been incredible to be able to do my PhD in the Ion Storage group where there is so much history 1 and such a spirit of kindness and collaboration. When I met Dietrich Leibfried at a SQUINT conference I attended as an undergrad I had no idea that I would eventually get to work with him on a daily basis. His passion and love for physics is a constant inspiration. Andrew Wilson did a lot of my early training, especially working with the doublers, and I am thankful to have gotten to learn from the doubler master himself. I will never forgive JILA for taking him from the Ion Storage Group! 2 Even though he left the group fairly early in my PhD, he has always been available to dispense practical wisdom in his delightful accent when I needed it the most, for which I am eternally grateful. Daniel Slichter's high standards in every area have made me a better physicist and I have learned a tremendous amount about grit and perseverance on this project. Our theory collaborators in the group led by Manny Knill and Scott Glancy always helped us push our work to be as rigorous as possible, and I learned a great deal from them.

Daniel Slichter and David Allcock did a great deal of the early planning and building on what became this experiment. Raghu Srinivas and Shaun Burd were the senior students on the Magtrap when I joined, and while I wish I had had more time in which to absorb their knowledge, I am grateful for this experiment they built and maintained, and also sorry for the amount of their work I ended up tearing down. Raghu handled the majority of my initial training on the project and I

¹Both metaphorically and also like, literally, still there in the basement of NIST.

²Just kidding. Maybe.

think he might be the single kindest and most patient person on this planet. He also introduced me to Phoebe Bridgers before she was cool, and continues to have impeccable taste in music and books. Raghu, I'm sorry for that time I made you cry by being so mad at math.

Loz Stephenson joined me as a postdoc less than a month before the COVID-19 lockdowns, so it took me a long time to get to really know him, but he turned out to be perhaps the second most patient person on the planet. Almost every time I was deep in the weeds on this project, he was right there with me with brilliant ideas and calm encouragement. He also did a huge amount of the thankless kind of work that makes a lab a nicer place to work—things like cable management and code documentation—so Loz, this is me thanking you for that (and also the rainbow LED keyboard that helped me write this thesis up so quickly). His enthusiasm and love for mountain sports and life in Colorado was delightful and contagious, and I hope he can move back into the mountains someday. Loz, I understand why you had to leave, but I hated being here without you. Alejandra Collopy, the other postdoc on this project, saved me from working completely alone during the COVID occupancy restrictions, for which I was grateful, and also served as a source of inspiration and motivation. Several postdocs who were not on my project were also my good friends and informal advisors: Adam Brandt, Jules Stuart, Panyu Hou, and Nathan Lysne all provided advice on technical matters, career decisions, and more. I am grateful that I got to work with all of you and wish you the best of luck in your future careers! To the next generation of quantum graduate students in the group: I wish you happy and well-behaved Stingers and best of luck on your projects. They will probably not be fast PhDs, knowing this group, but we can hope. Justin Niedermeyer has been my most constant companion in that cold, lonely lab and I have been very grateful for his company and the many times he has run in to save various pieces of equipment from power, chilled water, or HVAC outages. He is a great friend and a lab hero.

I have been lucky throughout my PhD to have so many lovely and talented colleagues, many of whom I got the chance to cite in this work. It's been lovely to live in Boulder where there is such a high density of quantum physicists.³ But for those of you who I wasn't able to cram a citation in

³A real conversation between Shawn Geller and I once went like this: "What fraction of the population would you

for (hello condensed matter friends), still know that it has been enriching to be around so much different physics.

I also need to thank my climbing and skiing partners, Aaron Young, Hope Whitelock, Charlie McLemore, Laurent Stephenson, and Theo Lukin Yelin for dragging me up and down rocks of various sizes up and down the Front Range (and sometimes Utah and Wyoming!).

Aaron Young was of course my partner in more than just climbing and skiing for 90% of my time in Boulder, and I am very grateful to him for all those years of love and support (even though he insists on working with one more electron than God intended). I loved getting to explore Colorado and grow up with you, and I'm very excited to see what you do next.

Hope Whitelock has been my best friend for many years now. She is always there when I need to vent, or to go up and down mountains without talking at all. She's pushed me to be radder, and also provided lots of space to be absolutely terrible at my chosen hobbies. She deserves at least six PhDs for the amount of garbage she has had to deal with during her graduate career.

I want to thank my long-term roommates Will Milner and Shawn Geller for all the late-night physics discussions, making sure my annual Halloween party was well attended with both physicists and whoever Shawn happened to meet at the Trident that week, and for helping me eat all that bread I kept baking. A special shoutout to Shawn for always tolerating my intolerance of paradoxes.

Susanna Todaro and Raghu Srinivas were my trans-Atlantic technical and emotional support team, and their help and support have been invaluable especially in this last year. I am so grateful to have met and gotten to work with both of you.

I need to thank the Women and Gender Minorities in Physics group at CU, with which I have been heavily involved during my time here. The work that that group does is so important and so special. I need to thank Iona Binnie and Joanna Lis especially for taking it over when Hope and I needed to retire, but also everyone who comes to the meetings, especially the angry ones. That space has helped keep me sane and grounded, and I hope it will continue to help others.

My family has also been so supportive of my decision to spend all of my 20's in school, and I estimate are physicists? Like 70-80%?" "Yeah I think that sounds about right."

am glad they all were able to come visit me in Boulder at various times. My sister Carolyn was able to come live with me this last winter and spring and her shining presence and constant love helped me through an extremely difficult time in my life. I can't wait to get completely sunburnt and devoured by mosquitos on our annual backpacking trip every year for the rest of our lives. My dad spent my entire childhood teaching me science and engineering, so I guess a lot of this is his fault. My mom used to give me math problems while we were grocery shopping, so she's not free from blame either.

Special thanks to Grant Mitts at UCLA, fellow Stinger sufferer, for reviewing my appendix on Stinger operation. Also to Katie David and Lindsay Sonderhouse from the Quantum Networks experiment for their feedback on my indium trap mounting appendix, and especially to Katie for taking the process photos I was able to include. I also need to thank everyone who read drafts of various chapters: Susanna Todaro, Raghu Srinivas, Dietrich Leibfried, and Daniel Slichter. Their comments and suggestions vastly improved this thesis and any remaining confusions and errors are mine alone.

Thanks to Tina Bowers and Daniel Slichter for motivating me during this final push. I will watch the future of this project with great interest.

Contents

Chapter

1	Intr	roduction	1
	1.1	Mixed-species ion crystals for quantum information	2
	1.2	Laser-free control and entanglement	4
	1.3	The all-electronic qubit	5
	1.4	This work	7
2	Qua	antum computing with trapped ions	8
	2.1	Rf ion traps	11
	2.2	Surface electrode traps	17
	2.3	Ion loading and transport	19
	2.4	Trapped ion qubits	21
	2.5	Motional and qubit state control	21
	2.6	Entangling gates	24
		2.6.1 The Cirac-Zoller gate	24
		2.6.2 Geometric phase gates	26
3	Lase	er-free quantum information with trapped ions	28

	3.1	Laser-	free spin-motion coupling		31
		3.1.1	Static gradient spin-motion coupling		33
		3.1.2	Oscillating magnetic field gradient spin-motion coupling		33
	3.2	Groun	nd state cooling using laser-free spin-motion coupling		35
	3.3	Laser-	free single ion addressing		36
	3.4	Laser-	free gate implementations		38
		3.4.1	Static gradient		38
		3.4.2	Near-qubit frequency oscillating gradient		38
		3.4.3	Near-motional frequency oscillating gradient		39
4	Mix	ed spe	ecies trapped ion quantum information	2	41
	4.1	Motion	nal modes for mixed-species ion crystals		42
	4.2	Sympa	athetic cooling		46
	4.3	Quant	cum logic		47
		4.3.1	Projective state preparation		48
		4.3.2	Quantum logic readout		49
	4.4	Mixed	l-species entanglement		50
5	Mag	gnesiuı	m & calcium ions	ļ	52
	5.1	Magne	esium		52
		5.1.1	Ion production		52
		5.1.2	Micromotion compensation		54
		5.1.3	Axial stray field compensation		55
		5.1.4	Cooling		55
		5.1.5	State preparation & manipulation		57
		5.1.6	Readout		59
	5.2	Calciu	<u>ım</u>		61
		5.2.1	Ion production		63

		5.2.2	Micromotion compensation	63
		5.2.3	Cooling	65
		5.2.4	State preparation and manipulation	68
		5.2.5	Readout	72
	5.3	Trap st	tability & loading for mixed species crystals	72
6	App	aratus		74
	6.1	Vacuur	m chamber	75
		6.1.1	Re-entrant top viewport	75
		6.1.2	Optical access viewport modifications	76
		6.1.3	Pillbox modifications	76
		6.1.4	Bias-able mesh	76
		6.1.5	Ovens	78
		6.1.6	Vacuum pumps	79
	6.2	Cryost	at	79
	6.3	High-fi	eld coils	81
	6.4	Shim c	oils	83
	6.5	Imagin	ng system	83
	6.6	Laser s	systems	85
		6.6.1	Magnesium	85
		6.6.2	Calcium	89
		6.6.3	Beam pointing	93
	6.7	Trap .		93
	6.8	Trap r	f generation	97
		6.8.1	Rf modulation for parametric micromotion compensation	98
	6.9	dc volt	age generation	98
	6 10	Micron	vovog	101

	6.11	Control system		101
7	Effe	cts of stray potentials	on mixed-species loading and tran	sport 104
	7.1	Trap geometry		106
	7.2	Historical loading and ch	harging in this trap	106
	7.3	Calcium-only operation		109
		7.3.1 Loading		
		7.3.2 Transport		
		7.3.3 Shuttle loading .		
	7.4	Mixed-species operation		
		7.4.1 Effects of magnes	sium loading on transport	
		7.4.2 Effects of magnes	sium loading on merging	
		7.4.3 Merge-free loadin	ng attempts	124
	7.5	Strategies for managing	and mitigating stray fields	126
		7.5.1 Velocity-selective	e loading of magnesium	126
			ld mitigation strategies	
	7.6	Summary of stray field e	$_{ m effects}$	129
		v		
8	Lase	er-free quantum logic		130
	8.1	Projective state preparate	tion	131
	8.2	Projective state readout		133
		8.2.1 Zeeman microway	ves	134
	8.3	Advantages of laser-free	quantum logic	134
	8.4	Mixed-species laser-free	entanglement for quantum logic readout	137
	8.5	Preliminary steps		137
		8.5.1 Sympathetic cool	ling	138
		8.5.2 Parametric mode	e-mode coupling	138
		8.5.3 Mode-mode coup	oling for quantum logic	145

		8.5.4	Faster magnesium-magnesium near-motional frequency oscillating magnetic	
			field gradient gates	145
		8.5.5	Calcium-calcium near-motional frequency oscillating magnetic field gradient	
			gate	149
9	Mix	æd-spe	ecies laser-free entangling gates	152
	9.1	Single	-species near-motional frequency oscillating magnetic field gradient gate	153
	9.2	Exten	ding the gate to multiple ion species	154
		9.2.1	Independent motional modes with added coupling	155
		9.2.2	Most complicated: non-independent motional modes with added coupling $\ \ .$	160
	9.3	Simula	ations	162
		9.3.1	Independent motional modes, with and without coupling	162
		9.3.2	Non-independent motional modes, no coupling	162
	9.4	Altern	nate method: PHRAP-style twisting	164
10	Con	clusio	ns & future directions	166
	10.1	Sugge	stions for improvements to this apparatus	166
			stions for future generations of laser-free mixed-species experiments	
\mathbf{R}_{0}	efere	nces		169
A	pper	ndix		
A	Reb	ouild		188
	A.1	Trap		188
		A.1.1	Scatter spot testing	190
	A.2	Pillbo	x	192
		A.2.1	Trap rf	192
		A.2.2	Filterboard	193

	A.3	Top viewport	. 193
	A.4	Homemade viewports	. 196
В	Stin	nger closed-cycle cryostat upgrade & operation	199
	B.1	Old cryostat	. 199
	B.2	Stinger cryostat	. 200
		B.2.1 Transitioning from the flow cryostat to the Stinger	. 201
		B.2.2 Cleaning cycles	. 202
		B.2.3 Supelco filter manifold	. 206
		B.2.4 Miscellaneous plumbing upgrades	. 210
		B.2.5 Automatic clog catcher	. 211
		B.2.6 Daily operation	. 211
		B.2.7 Cooling power measurements	. 215
\mathbf{C}	Two	o methods for mounting microfabricated ion traps with indium solder	216
	C.1	Indium soldering using NanoFoil	. 218
	C.2	Indium soldering using a hotplate	. 223
	C.3	Use in experiments	. 227
D	Qua	artz crystal thickness monitors	229
	D.1	Background	. 229
	D.2	Board design	. 230
		D.2.1 Passive	. 231
		D.2.2 Active	. 231
		B.2.2 1001/0	
	D.3	Soldering instructions & tips	. 233
	D.3		
	D.3	Soldering instructions & tips	. 233

			xiv	
	D.5	Guidelines for use	234	
	D.6	Implementation in experiments	235	
\mathbf{E}	VEC	CSEL source for beryllium photoionization	237	
	E.1	Introduction	237	
	E.2	Laser setup	240	
		E.2.1 940 nm VECSEL	242	
		E.2.2 Doubling to 470 nm and 235 nm	243	
	E.3	Photoionization testing	243	
	E.4	Doppler spectroscopy	244	
	E.5	Conclusion	245	

Tables

Table

4.1	Motional mode frequencies and excursions for example parameters in our trap (z	
	lenotes the trap axis)	

Figures

Figure

1.1	The quantum charge-coupled device (QCCD) $\ \ldots \ \ldots \ \ldots \ \ldots \ \ldots$	3
1.2	The all-electronic qubit	6
2.1	Quadrupole mass analyzer versus linear rf Paul traps	12
2.2	Cirac-Zoller vs. Mølmer-Sørensen gate schemes	25
3.1	Far-field vs. near-field magnetic field gradients	32
4.1	Laser-based projective state preparation	49
4.2	Quantum logic readout	50
5.1	Magnesium level diagram at 212.8 G	53
5.2	Effects of axial stray fields	56
5.3	Magnesium state preparation histogram	58
5.4	Magnesium microwave Rabi flops	60
5.5	Calcium level diagram at 212.8 G	62
5.6	Micromotion sideband ratio for large shim scan	64
5.7	Parametric micromotion compensation process	66

	xvii
5.8	Calcium EIT cooling
5.9	Calcium cooled using microwave sidebands
5.10	Calcium state preparation histogram
5.11	Calcium microwave control of Zeeman qubit
5.12	Calcium flopping on quadrupole transition
6.1	Top flange CAD comparison
6.2	Trap pillbox
6.3	Apparatus
6.4	Stinger plumbing diagram
6.5	Imaging system upgrade
6.6	Magnesium laser systems
6.7	Calcium laser table optical layout
6.8	Trap and laser beam layout
6.9	Trap table optics
6.10	Trap photos
6.11	New trap heating rates
6.12	Trap rf modulation circuit
6.13	Microwave electronics schematic
7.1	Trap aspect ratio
7.2	Large, deep 'bathtub' well for loading
7.3	Radial mode drifts following calcium loading
7.4	Effect of axial stray fields on splitting
7.5	Sensitivity of splitting to axial stray fields
7.6	Stray axial potential from loading magnesium in the center well
7.7	Visualization of different merge strategies
7.8	Stray axial potential from loading magnesium in the load zone

A.4	Trap rf, cracked and reconnected
A.5	Pillbox apertures and top window
A.6	Trap mesh upgrade
A.7	Vacuum chamber upgrades
A.8	Homemade vacuum viewports
B.1	Custom Stinger bayonet to retrofit ST-400
B.2	Stinger and accessories
В.3	Supelco manifold
B.4	Clog catcher
B.5	Stinger timeline
B.6	Cooling power measurements
C.1	Epoxy-mounted trap tests
C.2	3D printed indium soldering jig
C.3	NanoFoil trap mounting process
C.4	Hotplate soldered traps
C.5	Hotplate soldered traps, close up
C.6	Indium-mounted traps
D.1	QCTM circuit diagrams
D.2	QCTMs installed in experiments
D.3	Plated QCTM
D.4	QCTM measurements
E.1	Energy levels relevant to QIP with beryllium ions
E.2	VECSEL and doubling cavities
E.3	940 nm VECSEL tuning
E.4	Doubling cavities to convert 940 nm light to 235 nm light

E.5	Measured output powers and efficiencies of frequency-doubling stages used in the
	940 nm VECSEL system
E.6	Beryllium photoionization spectroscopy
E.7	Beryllium photoionization test chamber

1

Introduction

The control of quantum systems for quantum computing, sensing, and simulation is a broad and active field of research. Quantum computing entails using two-level quantum systems as qubits in a quantum information processor, analogous to the bits in a conventional computer. It can then leverage the unique features of quantum mechanics to perform calculations which would be challenging or even impossible with a conventional computer [1, 2]. We are currently in an age of quantum computing which has been called the 'noisy intermediate-scale quantum' (NISQ) era [3] because the current systems tend to lack the number of qubits and quality of control necessary to perform large or complex computations, but lots of exciting work is currently underway to address these deficiencies. Since the first demonstrations of trapped ion entangling gates at the turn of the century, error rates have decreased by orders of magnitude for both single- and two-qubit gates, and qubit numbers have reached double digits in trapped ions (and triple digits for neutral atoms).

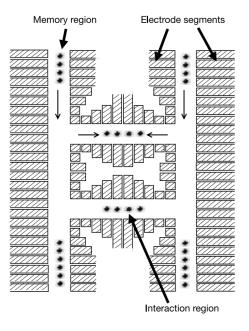
Progress in quantum information and sensing over the past three decades can be broadly split into two main branches: improving the quality and complexity of control for established systems as the numbers of qubits increases, and extending existing control techniques to new, more exotic quantum systems such as molecules, radioactive ion species, and subatomic particles. Trapped ions are one of the leading platforms in both branches due to the possibility to trap and control single particles up to large ensembles using the suite of techniques developed over the past seventy years since the initial proposal of the ion trap [4]. Researchers have demonstrated high-quality entanglement of ion pairs [5–9], control of many ions in a single system [10–13], as well as control of exotic particles and molecules [14–18]. Trapped ions have also formed a basis for both current and proposed investigations for fundamental physics [19–22]. This thesis focuses on two particular strategies: using multiple trapped ion species within a single processor, and performing as much control as possible with microwave and radiofrequency fields and field gradients rather than lasers. Both of these approaches have already enabled significant advances and are well-positioned to continue to do so.

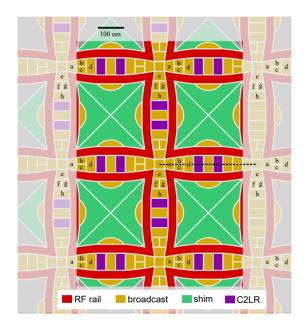
1.1 Mixed-species ion crystals for quantum information

Using multiple trapped ion species within a single processor was first proposed in 1998 [23] for the purpose of sympathetically cooling computational ions. In that initial proposal, which was expanded upon in [24], the idea of a 'quantum charge-coupled device' (QCCD) was also put forward. This is an architecture for a trapped ion computer involving different zones for different purposes, with ions being transported between zones depending on the desired operations. Long linear chains of ions have many modes which need to be tracked and cooled, and suffer from motional mode crowding, crosstalk errors, and slow entangling gates. The QCCD offers a path towards larger trapped ion systems while avoiding these difficulties inherent to long chains. An illustration of the system as initially conceived is shown in Fig. 1.1, showing the different zones for different purposes, next to a recent realization from Quantinuum that arguably demonstrates many of the properties of a QCCD.

Sympathetic cooling with a 'helper' ion between quantum logic gates allows longer sequences to be performed without needing to stop and cool data ions, which would disrupt their quantum states, and allows reinitialization of ion motion after rearrangements in a QCCD-like system [26, 27]. The addition of mode-mode coupling even enables cooling of modes with poor coolant ion participation

Figure 1.1: The quantum charge-coupled device (QCCD) describes a vision for scaling trapped ion computing using a zoned architecture. The figure on the left is from one of the early QCCD proposals [24], while the figure on the right is from Quantinuum's most recent trap [25]. Both images show a grid-based multi-zone trap that can hold and control multiple small strings of ions and move them arbitrarily through different trap zones.





[28–30]. Many groups and companies pursuing large-scale trapped ion quantum systems choose to use multiple ion species for all of these reasons [25, 31–37]. The second species can also be used as a 'memory' ion which is not disturbed by operations on the computational ions [24], another important step towards larger and more complex quantum processors.

In 2005, the demonstration of quantum logic spectroscopy (QLS) [38] provided another motivation for using multiple ion species in a given system. QLS uses the shared motion of two co-trapped charged particles to enable use of an ion species which does not have an accessible cycling transition for cooling or fluorescence detection - notably ²⁷Al⁺ in the record-breaking Mg⁺-Al⁺ clock [39], though the technique has since been extended to molecules and highly-charged ions as well [40–42]. While additional ion species increases experimental complexity, the benefits often outweigh the costs. And as researchers seek to expand the variety of particles that can be controlled at the quantum level, quantum logic could prove a powerful tool for target particles that lack easily addressable optical transitions, as long as we can couple an ion's spin to its motion.

1.2 Laser-free control and entanglement

The second focus of this thesis is a technique that was also proposed in 1998 [23]. This is 'laser-free' entanglement, where magnetic field gradients, which can be static or oscillate at microwave and radio frequencies, are used to perform gates instead of lasers. Successful entanglement with this strategy was first demonstrated in 2011 [43] (more than a decade after the first laser-based ion-ion entangling gate [44]), and recently reached comparable fidelities and speeds with laser gates [8, 9, 45]. Laser-free entanglement offers a lot of potential benefits: one of the primary advantages is the elimination of photon-scattering and other laser-based errors, which are the dominant error source in most previous high-fidelity laser-based entangling demonstrations. There are practical advantages as well: microwave hardware is cheap relative to lasers, phase control and stability are vastly easier to achieve, and integration of electrodes to drive qubit interactions into traps is based on established technology, especially when compared to the challenges of integrated photonics. Microwave currents within an electrode can drive gates in parallel on any charged particles with a magnetic moment

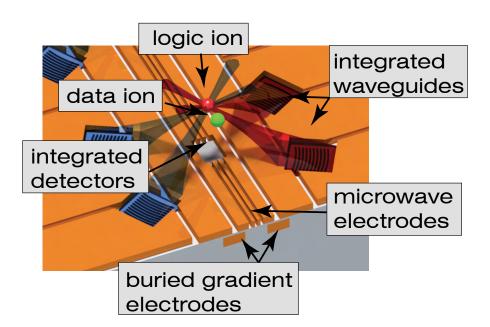
above the electrode, allowing for very power-efficient parallel gates. Compare this to laser-based parallel gates, which require either many tightly focused beams or very high power sheet beams, with more optical power required to perform additional gates. And while it may not be as obvious as individual ion addressing using tightly focused laser beams, individual ions and pairs can still be addressed with a combination of local electric and magnetic fields and global microwave magnetic field gradients [8, 46–49].

Laser-free control and entanglement could also help extend quantum control to more unusual charged particles, such as protons or exotic molecules. In this vein, ongoing work at UCLA aims to manipulate polar molecules with electric rather than microwave field gradients, pioneering a new style of laser-free interactions [50]. Laser-free quantum logic, the proposal of which is a focus of this thesis, could prove a powerful tool for both ordinary trapped ions and more exotic charged particles.

1.3 The all-electronic qubit

This thesis fits into the broader project in our lab of developing an 'all-electronic' qubit. This means a qubit species that requires only microwave and rf control fields. A second 'helper' species would have lasers for cooling and fluorescence detection, tasks which require relatively little optical power compared to that needed for laser-based entangling gates. By using integrated photonics as well as integrated superconducting single-photon detectors, both technologies which have been separately demonstrated [51–54], an ion trap could be developed with no need for the optical access viewports currently used to supply cooling and control light, as well as collect fluorescence signals, that currently occupy much of the space on any trapped ion vacuum chamber. This scheme would require only electronic wires and optical fibers entering the chamber, a huge step towards miniaturizing trapped ion systems while also reducing sensitivity to vibrations and other external perturbations. The vision as presented in Fig. 1.2 shows a combination of integrated trap technologies that would allow for such a system. Each of these have been demonstrated separately and future work in the lab will focus on systems which integrate multiple of these technologies into a single trap chip.

Figure 1.2: The all-electronic qubit would require only electrical input and output lines for qubit species control. The helper species would still require laser light but this could be supplied via integrated photonics and detected using integrated detectors, eliminating the need for free-space optics and the space-consuming windows and imaging systems they entail. Figure from Daniel Slichter.



1.4 This work

This work focuses on laser-free operations in a mixed-species trapped ion system with the aim of having one laser-free species. The apparatus described in this thesis has previously used only magnesium ions, which make good qubits but require lasers in the ultraviolet. We added calcium as a helper species due to its demonstrated success with integrated photonics [51, 53] and reasonable mass match. We will begin this thesis with a broad introduction to ion traps and how trapped ions can be controlled at a level that renders them useful for quantum information processing. Then we will review the principles of laser-free entanglement and control including the motivation and recent milestones, followed by the arguments in favor of using multiple ion species for larger scale processors. We then describe the apparatus designed to perform the experiments described in the final chapters. These proposed techniques of laser-free quantum logic and laser-free mixed species entanglement each constitute a chapter detailing theoretical underpinnings, followed by preliminary experimental results or simulations demonstrating feasibility. Significant results presented include trapping, transport, and cooling of mixed-species ion crystals in this system for the first time; fast, coherent mode-mode coupling of one- and two-ion crystals; and a theoretical description of a mixed-species laser-free entangling gate using parametric coupling to increase the speed of the entangling interaction. Finally, a series of technical appendices describe the various optical and electrical engineering projects undertaken during this thesis, the details of our vacuum chamber rebuild, how to mount ion traps using indium, and recommendations for installation and operation of a Stinger closed-cycle cryostat.

Quantum computing with trapped ions

Classical computers operate by performing logic operations between two-level systems, known as bits. Quantum computers perform quantum logic operations on quantum bits, or qubits, and the quantumness of these bits opens up new avenues for interesting and complex computations. A bit is always either 0 or 1, while the quantum nature of a qubit means it can be in a superposition of 0 and 1. Qubits can also become entangled with each other. These two properties of entanglement and superposition underlie the quantum algorithms that define tasks where quantum computers could dramatically outperform their classical counterparts [55, 56].

A qubit can be formed using any quantum two-level system. Popular physical implementations of qubits include trapped ions and neutral atoms, where two electronic levels of the atomic structure can be selected to form the qubit states, and superconducting qubits, where the energy states of a small superconducting circuit form the quantum levels. For a trapped ion or neutral atom qubit, the two levels are often labeled as 'up' and 'down' despite (usually) not being a single spin. In a trapped ion, there are several types of qubit that may be used, categorized by the frequency of the transition splitting the qubit states. Optical qubits occur when the two chosen levels are

¹While some are specific and refer to this as a 'pseudospin', in this work I will call it a spin for simplicity.

separated by a frequency in the optical regime. Hyperfine qubits, where the levels are chosen out of the ground-state² hyperfine manifold of an ion, are typically split by microwave frequencies and usually manipulated using Raman transitions. There can also be Zeeman qubits, where the frequency splitting is often in the radiofrequency band. Because these qubits are not necessarily actual two-level systems, we can also choose to encode multiple qubits within a single ion, or to shift the encoding of a given qubit for specific purposes. This last technique is the basis of the *omg* scheme, discussed in more detail in Chapter 3.

While other technologies have recently made rapid advances in qubit number [57, 58], gate fidelities [59–63], and complexity of operations [57, 58, 62, 64–66], trapped ions remain one of the leading platforms for quantum information processing [8, 9, 67].

When Benioff, Manin, and later Feynman proposed what would become known as quantum computing in the early 1980's [1, 68, 69], the ability to laser cool and trap single ions—tools that would enable the use of those trapped ions as qubits—had already been demonstrated [70–74].

In 1995 Cirac and Zoller proposed using trapped ions as a platform for quantum information processing, based on a scheme for entangling trapped ion qubits using their shared motion as an information bus [75]. However, this Cirac-Zoller gate requires perfect ground-state cooling and individual addressing of the qubit ions, both of which can be technically challenging. When Mølmer and Sørensen proposed their eponymous gate in 1999 [76, 77], it relaxed the stringent requirements on ion temperature and addressing and became the dominant method for entangling trapped ions. In fact, it was successfully demonstrated in 2000 [44], before the Cirac-Zoller gate which was not demonstrated until 2003 [78].

By the time DiVincenzo laid out his criteria for the ideal physical implementation of a quantum computer in 2000 [79], trapped ions had already fulfilled many of the criteria, and have only continued to improve in the intervening years:

• A scalable physical system with well-characterized qubits: Atomic qubits, such as neutral atoms and trapped ions, are always identical which makes them very well

²Or metastable state.

characterized. Scalability is still a work in progress but control of up to 56 trapped ion qubits has been demonstrated [13] and there are blueprints in place to extend systems to more [23, 24].

- The ability to initialize the state of the qubits to a simple fiducial state: Trapped ion qubits can be prepared in a specific state with extremely high fidelity, typically by optical pumping. Combined state prep and measurement errors of 10⁻⁴ have been achieved in trapped ions [80, 81].
- Long relevant decoherence times, much longer than the gate operation time:

 Trapped ion qubits can have exceptionally long coherence times, recently surpassing an hour [82]. Typical gate times are on the order of tens to hundreds of microseconds, so this criterion is well-satisfied.
- A 'universal' set of quantum gates: 'Universal' here means capable of performing universal computation. A universal set of gates can be made from arbitrary single-qubit rotations, which can be achieved with extremely high fidelity by either laser or microwave control [83], and a two-qubit entangling gate [2]. Bell-state fidelities in trapped ions surpassed 99.9% fidelity in 2016 [5, 6] for laser-based entanglement, and more recently for laser-free entanglement [8, 9].
- A qubit-specific measurement capability: Trapped ion qubits are typically read out using fluorescence detection, where a closed cycling transition scatters photons only if the ion is in one of the two qubit states. Readout fidelity can be pushed even higher using quantum non-demolition techniques—ones that do not destroy the quantum information, allowing for repeated readout of the same qubit and errors at or below the 10⁻⁴ level [84, 85].

There are two additional bonus criteria for connecting physically separated quantum systems (also known as quantum networking):

- The ability to interconvert stationary and flying qubits
- The ability faithfully to transmit flying qubits between specified locations

both of which have been satisfied by various networking experiments [86–89].

In the nearly 30 years since Cirac and Zoller's initial proposal, experiments with trapped ions have proven their capability as a platform for quantum computation. This chapter will begin with the basic principles of ion trapping, then progress into the control of ions' internal and motional states at the quantum level, finishing with a discussion of trapped ion qubits and different trapped-ion entangling gates.

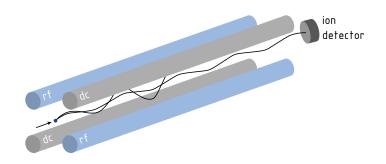
2.1 Rf ion traps

Because ions carry an electric charge, they can be trapped by means of their interaction with electric and magnetic fields. There are two main varieties of ion trap: Penning and Paul traps. Penning traps can contain large numbers of ions by taking advantage of the cyclotron motion of charged particles in strong magnetic fields, while linear Paul traps use rapidly oscillating electric fields applied to linear electrodes to confine ions at the minima of the oscillating field.³ We will focus on these linear Paul traps in this work, where one or several ions are confined along a one-dimensional line of minimal rf electric field.

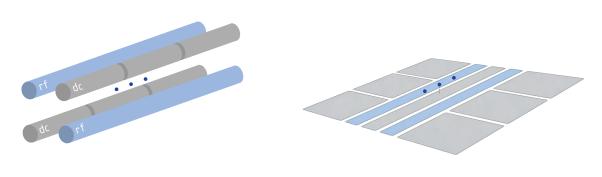
Although the first linear Paul trap was not demonstrated until the early 1990's [90], linear Paul traps operate using the same principles as the quadrupole mass analyzers developed for mass spectrometry in the mid-20th century. The similarities can be seen in Fig. 2.1. These mass analyzers work by applying varying amplitudes and frequencies of radiofrequency (rf) voltage to two of the four rods that form the device (the other two are grounded). Only ions with certain charge-to-mass ratios can travel successfully along the analyzer without entering an unstable orbit, and the chosen charge-to-mass ratio can be adjusted by adjusting the rf and dc voltages. Similarly, in modern Paul traps care has to be taken to provide the correct rf and dc voltages to stably trap the desired ion

³While Paul traps need not be linear, this work will focus only on the linear variety.

Figure 2.1: Cartoons showing a quadrupole mass analyzer, with both a stable and unstable ion trajectory; a 3D linear Paul trap with segmented electrodes to provide axial confinement; and a surface-electrode trap. In all three figures, rf electrodes are blue while dc electrodes are gray.



quadrupole mass analyzer



3D Paul trap

surface electrode trap

species. Following first [91] and then [23], we will examine how these quadrupole mass filters work and then extend that to the trapping of ions.

For a particle to be trapped in a given dimension, it must experience a restoring force such that the always pushes it back towards the origin. The simplest restoring force is a linear one, F = -kx, where x is the position of our particle from the origin, and k is the effective spring constant.

For a particle with charge e in an external electric field (still in one dimension for now), this condition can be fulfilled by a potential

$$\Phi_x = \frac{\xi}{2}x^2 \tag{2.1}$$

where $\xi = k/e$ is a geometric constant with units of electric potential over length squared. The electric field along x is then

$$E_x = -\frac{d\Phi_x}{dx} = -\xi x. \tag{2.2}$$

For a particle with charge e in this parabolic potential, it experiences the force

$$F_x = -e\xi x = -kx \tag{2.3}$$

and so it is trapped in this dimension.

Extending this to three dimensions gives

$$\Phi = \frac{\Phi_0}{2R^2} (\alpha x^2 + \beta y^2 + \gamma z^2)$$
 (2.4)

where R is a characteristic distance to the electrodes creating the potential, and α , β , and γ are constants. If we had no constraints on these coefficients, we could easily define a potential that trapped in all three directions simultaneously, but Laplace's equation requires $\alpha + \beta + \gamma = 0$. The consequence that a charged particle cannot be stably trapped by purely static fields is known as Earnshaw's theorem [92].

For a quadrupole mass analyzer with its axis along z, we select $\gamma=0$ i.e. no confinement along the axis of the analyzer, and then $\alpha=-\beta=1$. This gives the electric fields along all three directions

$$E_x = -\frac{1}{R^2}x,\tag{2.5}$$

$$E_y = \frac{1}{R^2} y,\tag{2.6}$$

and

$$E_z = 0. (2.7)$$

The effects of these fields are none in the z direction, a restoring force in the x direction, i.e. stable oscillatory motion, and an anti-confining force in the y direction, inducing runaway acceleration. This means that motion in that direction is unbounded, and will increase until the particle crashes into an electrode. But by setting our potentials to oscillate, we can alternate between stable confinement in x and y, creating an effective confining potential in both directions (if we choose our parameters carefully) that will allow the motion to remain stable along the length of the mass filter.

To see this, we start by applying an oscillating voltage $V_{rf} = U + V \cos(\Omega_{rf}t)$ to the rf electrodes. Following [91], this creates a potential along the center axis of the trap with the form

$$\Phi(x, y, t) = \frac{U + V \cos(\Omega_{rf}t)}{2} \left(1 + \frac{x^2 - y^2}{R^2}\right)$$
 (2.8)

and the equations of motion in our radial directions are

$$\ddot{x} + \frac{e}{mR^2}(U + V\cos(\Omega_{rf}t))x = 0$$
(2.9)

and

$$\ddot{y} + \frac{e}{mR^2}(U + V\cos(\Omega_{rf}t))y = 0. \tag{2.10}$$

If we rewrite these in terms of the unitless parameters

$$a = \frac{4eU}{mR^2\Omega_{rf}^2},\tag{2.11}$$

$$q = \frac{2eV}{mR^2\Omega_{rf}^2},\tag{2.12}$$

and

$$\tau = \frac{\Omega_{rf}t}{2} \tag{2.13}$$

we recover the Mathieu equations

$$\frac{d^2x}{d\tau^2} + (a + 2e\cos(2\tau))x = 0 (2.14)$$

and

$$\frac{d^2y}{d\tau^2} - (a + 2e\cos(2\tau))y = 0 (2.15)$$

which have well-known solutions. The parameters a and q are known as the Mathieu stability parameters, where a corresponds to the static part of the potential and q to the dynamic part: a two-dimensional plot of these parameters is known as a stability plot. The parameters are mass-dependent which is how these mass analyzers worked - by carefully selecting U, V, and Ω_{rf} only a certain range of masses will successfully pass through the analyzer. Charged particles whose parameters fall outside of the stability region will enter unstable orbits and leave. In real ion traps, typically $a \ll q$, the static potential is much less than the oscillating potential.

With $a \ll q < 1$, the solution to first order in q is

$$x(t) = x_0 \cos(\omega_r t + \phi) \left[1 + \frac{q}{2} \cos(\Omega_{rf} t) \right]$$
 (2.16)

where

$$\omega_r = \frac{q\Omega_{rf}}{2\sqrt{2}} \tag{2.17}$$

and x_0 and ϕ depend on the initial conditions. For a stably trapped ion, ω_r is smaller than Ω_{rf} so this looks like a slow oscillation at some characteristic frequency that depends on the ion and trap parameters (what we call secular motion) modified by a small, fast oscillation at the trap rf frequency (what we call micromotion). If we neglect the micromotion, we can approximate the ion as a harmonic oscillator in a pseudopotential

$$\Phi_p = \frac{1}{2}m\omega_r^2(x^2 + y^2). \tag{2.18}$$

The degeneracy between the radial secular frequencies can be broken by applying a small static potential U to the trap electrodes [23, 93].⁴ In order to trap ions in three dimensions, a static, harmonic potential along the trap axis can be applied to generate confinement along z as well (though this necessarily generates anti-confinement in x and y, per Laplace's equation). Typically, this axial confinement is designed to be weaker than the radial confinement so the ions arrange themselves into linear chains rather than zigzagging or otherwise bulging in the radial directions. In the pseudopotential approximation (where we average the force of the rf over a full period, approximating the effect of the periodic field into an effective potential [94, 95]), and adding this harmonic confinement along z, we can create the desired three-dimensional harmonic potential

$$\Phi = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2), \tag{2.19}$$

where $\omega_x = \omega_y = \omega_r$ and $\omega_z = (2e\kappa U/mz_0^2)^{1/2}$, with U the electrostatic axial potential applied to the endcap electrodes, κ a geometric factor, and z_0 related to the axial well parameters.

For linear chains of multiple ions in a harmonic trap, the normal modes of the chain can be calculated from the normal modes of a single ion in the same trap [23, 96]. If the masses of the ion species in the chain differ, this is still true but the radial pseudopotential is mass-dependent and the math becomes trickier as the ion ordering starts to affect mode structure [27]. A discussion of the normal modes of mixed-species ion chains can be found in Chapter 4.

⁴For us this is done with our rf bias, as discussed in Chapter 6.

When micromotion is neglected, a trapped ion can be treated as a three-dimensional quantum harmonic oscillator (QHO). The secular frequencies mentioned above form the normal modes of this oscillator, and we can write the Hamiltonian for a single mode of the QHO (neglecting the zero-point energy)

$$\hat{H}_{OHO} = \hbar \omega \hat{n} \tag{2.20}$$

where \hbar is the reduced Planck constant and \hat{n} is the number operator for the mode, $\hat{n} = \hat{a}^{\dagger}\hat{a}$, with \hat{a}^{\dagger} and a the raising and lowering operators of a quantum harmonic oscillator. Ion motional states are typically described in terms of the Fock 'number state' basis $|n\rangle$, where n is the number of motional quanta in that mode. We cool the ion's motional modes near their ground state for many of the applications described in this work, i.e. to a motional state approximating $|0\rangle$.

In general neglecting micromotion is valid because the oscillations are small and fast relative to the secular motion, but if an ion is pushed away from the axis of the trap (and thus away from the null of the rf quadrupole), what is called 'excess' micromotion can be quite large and can negatively affect control of the ions [97].⁵ This excess micromotion can be minimized, as discussed in Chapter 4.

2.2 Surface electrode traps

While lots of interesting work is actively underway in terms of new and interesting 3D architectures for ion traps [98–101], surface electrode traps are considered a promising technology for scaling trapped ion systems [102, 103]. Rather than the traditional four-rod ion traps described above that must be carefully assembled and aligned by hand, and have limited flexibility in terms of trap design, a surface electrode trap 'unfolds' the electrode structure into a single surface, as seen in Fig. 2.1. Placing the electrodes onto a single surface allows scientists to take advantage of existing microfabrication techniques, meaning that large quantities of high-precision, identical trap chips with small and intricate features can be manufactured at once. This could provide an advantage in

⁵Though it can also be intentionally induced, to effect individual addressing [46].

scaling these types of systems as many identical unit cells could be created and connected. The flexibility and sophistication of these fabrication techniques allow for unique and complicated trap geometries beyond the purely linear. Placing all the electrodes onto one surface also provides improved optical access. In our one-off, proof-of-principle style experiments which typically use unique trap chips with different form factors and wiring requirements in every experiment, we still require most of our wiring and trap assembly to be painstakingly done by hand, but the newest generation of experiments in the group has started using a modular system where traps are mounted to uniform 'pucks' with electrical quick-connects that allows more rapid iteration of prototype traps [104]. In the industry sector, utilizing traps with small identical building blocks can allow for parallel wiring and operations as companies work to address the difficulties of controlling larger numbers of ions [25, 105]. Another compelling alternate architecture for larger-scale surface traps is the micro Penning trap developed at ETH Zurich [106].

Surface traps also have a lot of potential for integrating different technologies into the trap chips themselves using advanced microfabrication techniques. Integrated technologies such as microwave antennae and resonators [43, 107], DACs [108], waveguides and grating couplers [52, 53, 109], optics [110–115], detectors [54, 116–118], and active elements for light modulation [119] have all been demonstrated and future work in this group and others will strive to combine multiple technologies within a single trap chip, or multiple chips within a single system [120].

Despite their many advantages, there are certainly drawbacks to using surface electrode traps, such as the shallower trap depths and larger anharmonicities in the trapping potentials compared to 3D traps [102]. Additionally, ions are often confined much closer to electrode surfaces than in traditional Paul traps - tens of microns rather than hundreds. Given that anomalous heating, where ions near surfaces experience heating whose source is not understood, increases as $1/d^4$ where d is the ion-electrode distance, the effects are worsened as ions get closer to surfaces. Having laser beams very close to the trap surfaces can cause laser-induced charging of the trap surface as well [121] Cryogenic operation reduces anomalous heating [122, 123] and increases ion lifetime, so many surface

⁶Though of course this is possibly in 3D traps as well, especially those built using stacked wafers.

traps are used at cryogenic temperatures, though cryogenic operation can introduce new challenges (see Appendix B for details of our cryostat). While it is possible to operate in a surface trap at room temperature, ion lifetime and heating rates may create challenging experimental conditions.

There are a multitude of intermediate 'wafer'-style traps [124–126] that offer some benefits of both surface and four-rod traps, though these can suffer from difficulties in alignment and construction. Surface electrode traps remain a leading candidate for scaling trapped ion quantum processors.

2.3 Ion loading and transport

There are several established techniques for producing ions and loading them into the trap. Perhaps the most common is oven loading: a thermal beam of neutral atoms of the desired species is produced by heating up a sample of metal, and this beam interacts with either an electron beam or resonant photoionization light in the trapping region. Cooling light is typically applied as the ions are generated and ions that are cooled to low enough energies inside the trapping zone can become trapped. Because the neutral atoms produced by an oven have energies following a thermal distribution with a mean energy much higher than most surface trap depths, typically orders of magnitude more ions and electrons are created for each ion that is successfully loaded. These extraneous charged particles can adhere to trap surfaces, leading to large and variable stray fields, as discussed in Chapter 7.

A second popular technique is ablation loading, wherein a high peak power pulsed laser (typically a pulsed Nd:YAG) is focused onto a solid sample containing the metal of choice [127]. This technique has the advantage of often not requiring separate photoionization lasers as the ablation directly produces ions. However, these ions tend to have high velocities and ablation loading into surface traps is often difficult, though not impossible [128]. Another technique which is gaining popularity is MOT loading, where neutral atoms of the desired species are kept near the trap in a magneto-optical trap (MOT) and neutral flux is produced from the MOT rather than directly

⁷See Appendix E for some investigations of the thermal properties of the neutral flux produced by beryllium ovens.

from an oven [129]. This has the advantage of starting off which much cooler neutral atoms, leading to increased loading rates especially for shallow traps, at the cost of a fair amount of technical overhead.

Once ions have been loaded into the trap, they will often need to be moved out of the loading zone and away from the stray fields created by the loading process. The initial proposals for the QCCD involved a zoned architecture with separate zones for loading ions, performing entangling operations, and storing memory qubits [23, 24]. As larger-scale ion traps become more and more of a reality, most groups have stuck to this zoned architecture technique [130, 131] and it has even recently spread into the neutral atom field [58]. A key preliminary operation for these large-scale traps is fast, stable ion transport even through the potentially complicated pseudopotentials produced by novel trap architectures with multiple zones separated by junctions. Ions can be transported along a linear trap by gradually moving the center of the well providing the axial confinement. Transport was first demonstrated in 2002 [132], followed by successful coherent transport through a junction with one ion [133], and recently with a mixed-species crystal [134]. Coherent transport has also recently been demonstrated in a multi-layer array of trapping sites [135]. Early transport efforts were performed on timescales much slower than the oscillatory period of the ions in the moving well (adiabatically) to avoid motional excitation that would need to be cooled away, but coherent transport can be sped up without inducing motional excitation by either carefully timing the length of the transport [136] or optimizing transport velocity profiles [137]. All transport performed in this thesis is adiabatic.

For loading of mixed-species crystals into a surface trap, it is often necessary to separately load and then merge the two ion species. Splitting and merging of wells is a crucial step for operation of a QCCD-style system, and was first demonstrated for a pair of ions in 2002 [132]. See Chapter 7 for a detailed discussion of loading, transport, splitting, and merging as it pertains to the mixed-species operations described in this thesis.

2.4 Trapped ion qubits

At the beginning of this chapter, we introduced the idea of the trapped ion qubit, formed by selecting two electronic levels of the ion and assigning them the labels $|\uparrow\rangle$ and $|\downarrow\rangle$. Regardless of the exact levels chosen, the Hamiltonian describing the qubit is the same:

$$\hat{H}_S = \frac{\hbar\omega_0}{2}\hat{\sigma}_z \tag{2.21}$$

where $\hat{\sigma}_z$ is the Pauli operator for the chosen spin and ω_0 is the qubit frequency.

Trapped ions make nice, well-isolated qubits because the Coulomb force keeps the spins, which form the qubit, much farther apart than the dipole-dipole interaction distance. This same Coulomb repulsion strongly couples the ion motion, and the shared motion can be used as an information bus if we can independently control and couple the ions' spins and motion.

2.5 Motional and qubit state control

In order to couple the spin to the motion, we need what is called a 'state-dependent' or 'spin-dependent' force.⁸ This can be produced by a field gradient, because different spin states will move in opposite directions in the gradient - the force they experience is thus spin-dependent.

To see how such a force can be generated using an external oscillating field, we will at first limit ourselves to two internal states and one motional mode. The Hamiltonian for the ion qubit plus single motional mode (sometimes called the 'bare ion' Hamiltonian) is

$$\hat{H} = \frac{\hbar\omega_0}{2}\hat{\sigma}_z + \hbar\omega\hat{n},\tag{2.22}$$

where we have just summed Eq. 2.20 and Eq. 2.21 from above.

To couple the motion and the spin, we need to now consider a Hamiltonian coupling the ion's internal states to its motion (such a Hamiltonian can be generated, for example, by a laser⁹), in the

⁸These two terms are used interchangeably.

⁹These same interactions can of course be driven without lasers, see Chapter 3.

interaction picture with respect to the bare ion Hamiltonian above, in the Lamb-Dicke regime, and after making the rotating-wave approximation ¹⁰:

$$\hat{H}_I = \frac{\hbar\Omega}{2}\hat{\sigma}_+ [1 + i\eta(\hat{a}e^{-i\omega t} + \hat{a}^\dagger e^{i\omega t})]e^{i(\varphi - \delta t)} + H.c.$$
 (2.23)

where Ω is the Rabi frequency,¹¹ $\hat{\sigma}_{\pm}$ are the spin state raising and lowering operators, δ is the detuning of the oscillating field from the qubit frequency ω_0 , η is the Lamb-Dicke parameter $\eta = |\mathbf{k}|\cos(\theta)z_0$ describing the ratio of the ion's ground-state wavepacket extent $z_0 = \sqrt{\hbar/2m\omega}$ to the wavevector $\mathbf{k} = 2\pi/\lambda \hat{k}$ of the interacting field projected along the mode of interest where θ is the angle between them, and φ is an arbitrary phase. Operating in the Lamb-Dicke regime means that $\eta \sqrt{\langle (\hat{a} + \hat{a}^{\dagger})^2 \rangle}$ is sufficiently small ($\ll 1$) that higher-order interactions are suppressed. For laser-based interactions, typical values of η are on the order of 0.01-0.1 while for laser-free interactions they are much smaller, see Chapter 3.1 for more discussion of why.

Following [93] and [138], this simplified interaction Hamiltonian has three resonances: the carrier, red sideband, and blue sideband interactions. These resonances occur when the detuning δ has values of 0 and $\pm \omega$. To see these resonances separately, a second rotating wave approximation must be done assuming we are only near one resonance at a time [138]. We will start with $\delta = 0$, the carrier transition. This Hamiltonian drives spin flips of the qubit at the Rabi frequency Ω :

$$\hat{H}_{car} = \frac{\hbar\Omega}{2} (\hat{\sigma}_{+} e^{i\varphi} + \hat{\sigma}_{-} e^{-i\varphi}). \tag{2.24}$$

The next resonance, $\delta = +\omega$ is called the blue sideband, because it is tuned blue of the carrier by an amount corresponding to the motional frequency of the near-resonant mode:

$$\hat{H}_{bsb} = \frac{\hbar\Omega}{2} \eta (\hat{\sigma}_{+} \hat{a}^{\dagger} e^{i\varphi} - \hat{\sigma}_{-} \hat{a} e^{-i\varphi}). \tag{2.25}$$

¹⁰For the full details and justification of all the assumptions and approximations involved to get to this point, I refer the reader to [23] and [93].

¹¹Here, using the definition in [23].

This drives oscillations between the $|\downarrow\rangle|n\rangle$ and $|\uparrow\rangle|n+1\rangle$ states with a frequency that now depends on both the Lamb-Dicke parameter as well as the motional mode occupation:

$$\Omega_{n,n+1} = \sqrt{n+1}\eta\Omega. \tag{2.26}$$

The final resonance, $\delta = -\omega$ is called the red sideband, as it is tuned red of the carrier:

$$\hat{H}_{rsb} = \frac{\hbar\Omega}{2} \eta (\hat{\sigma}_{+} \hat{a} e^{i\varphi} + \hat{\sigma}_{-} \hat{a}^{\dagger} e^{-i\varphi}). \tag{2.27}$$

This drives oscillations between the $|\downarrow\rangle|n\rangle$ and $|\uparrow\rangle|n-1\rangle$ states with frequency

$$\Omega_{n,n-1} = \sqrt{n\eta}\Omega. \tag{2.28}$$

The fact that the spin flip Rabi frequency is different depending on the motional state can be used to extract estimates of the Fock state distribution of certain states of motion, i.e. squeezed or displaced states. This analysis, and many examples of its use in trapped ion experiments, are covered in detail in [139].

One thing to keep in mind is that the sideband speed is reduced from the carrier speed by the Lamb-Dicke parameter for plane wave fields. This makes it impractical to drive sidebands using free-space microwave and rf radiation, where the wavelength is 3-4 orders of magnitude longer than for laser light, and thus the Lamb-Dicke parameter 3-4 orders of magnitude smaller. However, the larger field gradients that can be achieved by working in the near-field of a conductor carrying microwave and rf currents can give an increased sideband Rabi frequency for a fixed drive power, as is discussed in more detail in Chapter 3. Red and blue sidebands can be used to create non-classical states of motion [140, 141], perform cooling of motional modes very near to their ground state [142, 143], and also to drive entangling gates between ions. It is this last function that is the focus of the final section of this chapter.

2.6 Entangling gates

Quantum computing relies on the ability to create high-fidelity entanglement between qubits [79]. Trapped ions were the first platform in which a concrete proposal for generating such entanglement was proposed, by Cirac and Zoller in 1995 [75]. The next decade saw many new proposals for different types of trapped-ion gate, the most popular of which I will review here, roughly following [138]. Since trapped ion gates almost always use the shared motion to transmit information between ions, we often talk about gates being done 'on' a mode, or refer to a 'gate mode' or 'bus mode.' A gate can be performed on any motional mode, though there are many technical considerations that make certain modes better or worse (or faster or slower!) for different kinds of gates. Some of these considerations as they relate to this thesis are covered in Chapter 9.

2.6.1 The Cirac-Zoller gate

The Cirac-Zoller (CZ) scheme involves driving red sidebands on two different ions to enact a controlled-NOT gate via the shared motion of an ion chain. A cartoon in Fig. 2.2 illustrates the procedure. A chain of ions is first cooled into its ground state. A red sideband π pulse is driven on one ion in the chain, mapping any excitation in the ion's internal state into the motion of the ion chain. Then a red sideband 2π pulse is driven on a second ion, passing through a third auxiliary state so as not to flip the qubit, which causes it to acquire a phase if the motion of the chain was excited and the second ion was in its electronic ground state when the pulse was applied. A second red sideband π pulse on the first ion restores its internal state. Thus a phase shift is created in a second ion depending on the internal state of the first ion, generating entanglement.

However, high-fidelity entanglement using this scheme requires perfect ground-state cooling of the gate mode, as well as individual addressing of the ions, which can be technically challenging (although it can also be viewed as an advantage when it is desirable to perform gates on only certain ions out of a long chain, as demonstrated by [144]). The Cirac-Zoller gate was not physically implemented until 2003 [78] by which time a new gate scheme had entered the scene.

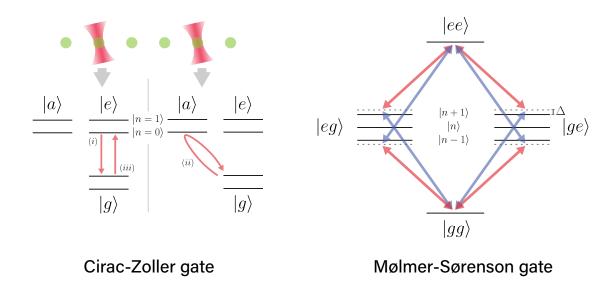


Figure 2.2: Schematics showing the basic operation of the Cirac-Zoller gate scheme, left, and Mølmer-Sørensen gate scheme, right. The Cirac-Zoller scheme requires individual addressing of the ions to be entangled, and uses the shared motion of the ion chain to add a phase to one ion dependent on the phase of a second ion. The states $|g\rangle$ and $|e\rangle$ form the qubit, while $|a\rangle$ is sometimes called the auxiliary state. Because it requires discrete π pulses to be successful, it requires ground-state cooling of the bus mode. The Mølmer-Sørensen scheme drives transitions on both ions simultaneously via a virtual state that is detuned by an amount Δ from the motional states, as shown. Because both sidebands must be driven simultaneously for a successful transition, the exact motional state does not matter (if in the Lamb-Dicke regime) and the $|eg\rangle$ and $|gg\rangle$ states are never populated. This gate does not require individual addressing or ground-state cooling, making it technically more accessible.

2.6.2 Geometric phase gates

While the CZ gate uses red sideband pulses to perform a phase shift on one ion conditioned on the shared motional occupation of the chain, these new gates involve generating differential phase shifts depending on the ions' internal states. These can be visualized as driving enclosed loops (or other exciting shapes, so long as they return to the starting point [145]) in phase space. This new class of gates is appealing from a technical standpoint as it does not require individual addressing or perfect ground-state cooling. There are two variants which I will discuss separately.

2.6.2.1 The Mølmer-Sørensen gate $(\hat{\sigma}_{\phi} - \hat{\sigma}_{\phi})$

In 1999, Klaus Mølmer and Anders Sørensen proposed a new method for entangling trapped ions that relaxed requirements on ion temperature and did not require individual addressing [76, 77, 146]. The next year, it was experimentally realized at NIST on chains of two and four ions [44].

The Mølmer-Sørensen (MS) gate involves simultaneously applying a red and blue sideband to a pair of ions, slightly detuned from resonance with the motional mode acting as the information bus. The Hamiltonian for simultaneous red and blue sidebands on two ions (i.e. two copies each of \hat{H}_{rsb} and \hat{H}_{bsb} from above, with subscripts 1 and 2 denoting the ion number), detuned symmetrically from the qubit state by Δ , looks like

$$\hat{H}_{MS} = \frac{\hbar\Omega}{2} \eta (\hat{\sigma}_{x1} \pm \hat{\sigma}_{x2}) (\hat{a}e^{i\Delta t} + \hat{a}^{\dagger}e^{-i\Delta t})$$
(2.29)

The gate scheme is often visualized as in Fig. 2.2: both sidebands are detuned by Δ from resonance such that they can jointly drive transitions between $|gg\rangle$ (both qubits in their ground state) and $|ee\rangle$ (both ions in the excited state), but single sidebands cannot be driven. The gate is thus insensitive to the exact motional state because every transition both adds and subtracts a quantum of motion. It is a popular gate scheme and has been used to perform many of the highest-fidelity gate operations on trapped ions to date [6, 9]. The gate does not require perfect ground-state cooling of the motional mode because at the end of the gate the motion and the spin

are (ideally) completely disentangled. It only requires that the ion be cooled within the Lamb-Dicke limit, i.e. that the ion's motion does not cause it to sample a large fraction of one wavelength of the radiation driving the gate.

2.6.2.2 $\hat{\sigma}_z \hat{\sigma}_z$ gates

The $\hat{\sigma}_z\hat{\sigma}_z$ gate has many similarities with the MS gate: it was proposed around the same time [147], it also involves global illumination with a pair of laser beams, was first demonstrated at NIST [148], and was first used to perform a high-fidelity gate in 2016 [5]. It has also been used in some of the fastest trapped-ion gates performed to date [145]. The difference is the basis in which the interaction acts as well as the mechanism which creates the interaction. These gates commute with $\hat{\sigma}_z$ meaning that they can easily be dynamically de The $\hat{\sigma}_z\hat{\sigma}_z$ gate involves creating a spin-dependent force such that ions in different internal states experience physical displacements in opposite directions. In [148] this spin-dependent force was generated by means of a walking wave created by two lasers with a difference frequency near the gate mode.

The gate scheme presented in [149] is capable of generating both $\hat{\sigma}_{\phi}\hat{\sigma}_{\phi}$ and $\hat{\sigma}_{z}\hat{\sigma}_{z}$ interactions, but the demonstrations presented in [8], performed using the apparatus described in this thesis, focused on $\hat{\sigma}_{z}\hat{\sigma}_{z}$ gates. Details of how this and other laser-free gates can be performed are presented in the following chapter, and the extension of the scheme to multiple ion species is the primary subject of Chapter 9.

Laser-free quantum information with trapped ions

One of the biggest challenges for scaling any atom-based quantum computing system will be delivering the quantity and quality of laser light needed for state preparation, readout, and control. Trapped ion qubits have an advantage over neutral atom qubits in that lasers are not needed for trapping but only for cooling, detection, and control. Of these three tasks, cooling and detection require only a few tens of microwatts per ion per zone for typical systems, but laser-based entangling gates often require watt-level power. The photon scattering errors which often dominate the error budgets of many laser-based entangling interactions can be suppressed by detuning the lasers farther, at the cost of higher laser power [150], but laser-free gates are completely free of these errors and any fundamental limits on their fidelity are orders of magnitude below the threshold for fault-tolerance.

Integrated photonics offer a very promising technique for delivery of laser light to many sites in large-scale traps, and the few-micron beam waists offered by on-chip grating couplers [51, 52] reduce the power needed to reach the desired intensities. But transmission losses and stray charging

on dielectric surfaces in such systems have proven to be stubborn technical challenges [52, 53, 109, 151]. Especially as we progress towards genuine QCCD-like systems with many parallel gate zones, quantity and quality of light delivery will continue to be an important focus of technological progress. Outsourcing as much control as possible to integrated microwave electrodes rather than lasers could substantially improve prospects for scaling of trapped ion systems. Additionally, our ultimate goal of a mixed-species system with one laser-free species could allow the use of ion species whose wavelengths are not well-suited to integrated photonics to still be used in integrated photonic systems.

In their 1998 review of experimental issues in trapped ion quantum information, Wineland et al. state that spontaneous emission errors could be mitigated by driving hyperfine or Zeeman transitions with radiofrequency and microwave radiation, as these have effectively no scattering [23]. However, they continue, these sidebands would be very slow compared to laser-based sidebands unless large magnetic field gradients could be created (they estimate 29 kT/m to achieve a 1 MHz sideband Rabi rate for an ion with a magnetic dipole moment equal to the Bohr magneton and $z_0 \sim 10$ nm), but say that strong gradients could be achieved by placing a very small electrode very close to the ions (both the size of the electrode and the distance to the ions must be much less than the wavelength of the radiation to satisfy this condition). They also note that individual addressing could pose a challenge due to the long wavelength of such radiation [23]. In 2001, Mintert and Wunderlich pointed out that the introduction of a large static magnetic field gradient could mitigate the two main drawbacks of microwave radiation, namely the smaller Lamb-Dicke parameters and lack of focusing for individual addressing [152]. Applying one large static gradient across the trap zone breaks the degeneracy of the qubit frequencies for identical ions, allowing individual addressing, and combining the large static gradient with weaker microwave drives allows faster microwave sidebands without needing each microwave itself to generate a strong gradient [152]. Another strategy for driving interactions with static magnetic field gradients was presented by Leibfried et al. in Ref. [153] which involves transporting ions through an inhomogeneous magnetic field landscape created by either current-carrying wires or small permanent magnets, where the frequency of the

interaction would be controlled by the ion's transport speed, though interactions in this style have not yet been implemented.

In 2008, Ospelkaus et al. proposed the use of a microwave antenna integrated into the then-new technology of surface-electrode ion traps to apply strong oscillating magnetic field gradients very close to the ions [154], and this technique was realized experimentally in 2011 [43]. Laser-free gates have since made steady progress in both gate fidelity and speed [48, 155–157] with recent work proving that laser-free gates can be competitive with laser-based gates in terms of speed and fidelity [8, 9, 45, 158]. Generating microwave and radiofrequency tones is technically easier and cheaper than generating high-power, phase-stable laser light, and delivery to the trap via coaxial cables and microfabricated electrodes is simpler and more robust than aligning free-space optics. Once applied to a trap electrode, the microwave and radiofrequency tones can drive gates in any ion pair above the electrode with no additional power required even as the system gets larger, enabling efficient parallel gates. In contrast, to perform parallel laser-based entangling gates, either multiple focused beams or a higher-power 'sheet' beam that can illuminate many ions at once are required—power cannot be as easily distributed to and recycled between multiple ions with lasers as with currents in an electrode. Additionally, since these gates are performed on a given motional mode, individual ions and pairs over a given electrode can be selectively entangled by adjusting their confining wells to push modes in or out of resonance with the gate driving fields [48, 154].

The previous chapter demonstrated the importance of spin-motion coupling in trapped ion entangling gates. This chapter will focus on generating these interactions without using lasers. We will examine a few non-entanglement applications, then examine laser-free entanglement of trapped ions, detailing the various schemes in use. We will finish by looking at past work using the apparatus described in this thesis, focusing primarily on the gate presented in [149] and [8], to lay the groundwork for extending that gate to multiple ion species, the theory of which is the focus of Chapter 9.

3.1 Laser-free spin-motion coupling

The fundamental difficulty that needs to be overcome is the intrinsically small Lamb-Dicke parameter of microwave and rf fields, which determines the ratio of the sideband Rabi rate to the carrier Rabi rate. We will now quantify what Lamb-Dicke parameters could be expected using far-field microwaves, using numbers relevant to this work.

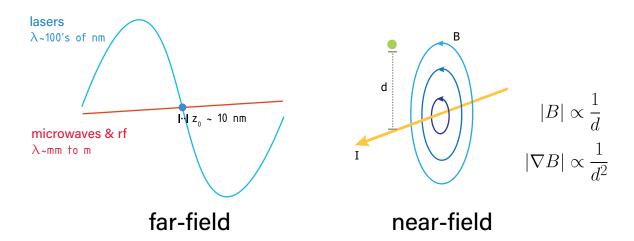
The Lamb-Dicke parameter η is

$$\eta = \frac{2\pi z_0}{\lambda} \tag{3.1}$$

where z_0 is the spatial extent of the ground-state wavefunctions $z_0 = \sqrt{\hbar/2m\omega_r}$. As in the previous chapter, m is our ion mass and ω_r is the motional frequency, and now λ is the wavelength of the radiation driving our transition. For the species and trap parameters used in this work, $z_0 \sim 5$ nm. For calcium, the Zeeman qubit splitting is 596 MHz, giving a wavelength of 50 cm, and the magnesium stretch qubit is 1.326 GHz for a wavelength of about 23 cm. This puts the Lamb-Dicke factors in the 10^{-8} range, meaning that we cannot effectively drive the transition using plane waves. For laser-based interactions, wavelengths in the hundreds of nm lead to Lamb-Dicke parameters ~ 0.1 . This occupies a sweet spot where the Lamb-Dicke criterion $\eta \sqrt{\langle (\hat{a} + \hat{a}^{\dagger})^2 \rangle} \ll 1$ is typically satisfied (i.e. the ion does not sample so much of the field gradient that non-linearities must be accounted for) but the interaction strength is not unduly hampered.

In 1998 Wineland et al. noted that while the wavelength of the radiation needed to drive a transition cannot be changed, a strong gradient can be generated by using small electrodes (much shorter than the wavelength of the radiation) very close to the ions [23]. Either static or oscillating magnetic field gradients can be used, and we will discuss the advantages and disadvantages of both schemes. In either case, the Biot-Savart law tells us that for a long, straight current-carrying wire the generated magnetic field goes as 1/d and the field gradient as $1/d^2$, where d is the distance between the ion and the wire [159].

Figure 3.1: Left: in the far field, microwaves have a much gentler field gradient than lasers, leading to very small Lamb-Dicke parameters and weak spin-motion coupling. Right: by going to the near-field, we can uncouple the gradient strength from the radiation wavelength, allowing much stronger magnetic field gradients.



3.1.1 Static gradient spin-motion coupling

In 2001, Mintert and Wunderlich stated that a static magnetic field gradient of a few hundred T/m [152], which had already been experimentally achieved using microfabricated conductors [160], could increase the ratio of sideband Rabi rate to carrier Rabi rate (what they call an effective Lamb-Dicke parameter $\tilde{\eta}$) up to ~ 0.01 . This meant that laser-free sidebands could be driven at speeds which made them a competitive technology for trapped ion entangling gates. Because the spin-motion coupling is mediated by the strong gradient, the GHz-frequency tones needed to drive motional sidebands on hyperfine qubits do not need to be especially powerful, which is useful as large GHz currents are technically demanding to generate. The static gradient can also be oriented so as to shift each qubit to a different frequency, allowing straightforward individual addressing of ions. However, this increases the technical and bookkeeping overhead because each qubit frequency must be tracked and generated separately. It also requires using magnetic field-dependent qubits, which intrinsically have increased sensitivity to magnetic field noise.

3.1.2 Oscillating magnetic field gradient spin-motion coupling

In 2008, Ospelkaus et al. proposed using the then-new technology of surface electrode traps to place trap-integrated conductors very near to ions [154]. They estimated that the strong gradients made possible by these integrated electrodes could allow microwave-driven gates to reach competitive speeds with laser-driven gates, and noted that both $\hat{\sigma}_z\hat{\sigma}_z$ and $\hat{\sigma}_\phi\hat{\sigma}_\phi$ interactions can be created.

Following [154] and [161], we introduce an oscillating magnetic field B(r,t), dependent on time and ion position r, creating an interaction Hamiltonian

$$\hat{H}_I(t) = -\sum_{n=1}^{N} \boldsymbol{\mu} \cdot \boldsymbol{B}(\boldsymbol{r}_n) \cos(\omega t - \varphi)$$
(3.2)

where μ is the ion's magnetic dipole moment, the index n denotes each ion up to the total N, and ω and φ are the frequency and phase of the magnetic field. In this section, z will be the axis defined by our quantization magnetic field, and x will be the out-of-plane radial mode. The third axis, y,

corresponds to the axis along which the control currents flow, and thus there is negligible magnetic field along this direction.¹

In the interaction picture with respect to the bare ion Hamiltonian, after making a rotating wave approximation, 2 this gives the Hamiltonian for the nth ion:

$$\hat{H}_{n} = -\frac{\hbar}{2} \Omega^{x} \hat{\sigma}_{+,n} e^{-i((\omega - \omega_{0})t + \varphi)}$$

$$-\frac{\hbar}{2} \Omega^{z} \hat{\sigma}_{z,n} e^{-i(\omega t + \varphi)}$$

$$-\sum_{j} \frac{\hbar}{2} \Omega^{x}_{j,n} \hat{\sigma}_{+,n} e^{-i((\omega - \omega_{0})t + \varphi)} (e^{-i\omega_{j}t} \hat{a}_{j} + e^{i\omega_{j}} \hat{a}_{j}^{\dagger})$$

$$-\sum_{j} \frac{\hbar}{2} \Omega^{z}_{j,n} \hat{\sigma}_{z,n} (e^{-i\omega_{j}t} \hat{a}_{j} + e^{i\omega_{j}} \hat{a}_{j}^{\dagger}) + H.c.$$

$$(3.3)$$

where j indexes the motional modes and ω_j are their respective frequencies. This is the same Hamiltonian as Eq. 1 from [154], but with the terms rearranged so we can look at them one by one. First, though, we will also define all of these new Rabi frequencies: for our carrier-like terms

$$\Omega^x = \frac{1}{\hbar} B_x \mu_{x\uparrow\downarrow} \tag{3.4}$$

and

$$\Omega^z = \frac{1}{2\hbar} B_z (\mu_{z\uparrow\uparrow} - \mu_{z\downarrow\downarrow}) \tag{3.5}$$

where $B_{x(z)}$ is the magnetic field projected along x(z) and $\mu_{x\uparrow\downarrow} = \langle \uparrow | \mu_x | \downarrow \rangle$, $\mu_{z\uparrow\uparrow} = \langle \uparrow | \mu_z | \uparrow \rangle$, etc. are the matrix elements of the ion's magnetic dipole moment; and for our sideband-like terms

$$\Omega_{j,n}^x = \frac{b_{j,n}q_0}{\hbar} B_x \prime \mu_{x\uparrow\downarrow},\tag{3.6}$$

and

$$\Omega_{j,n}^{z} = \frac{b_{j,n}q_0^{j}}{2\hbar} B_{z} \prime (\mu_{z\uparrow\uparrow} - \mu_{z\downarrow\downarrow})$$
(3.7)

¹In our particular trap the quantization field is actually 22.5 degrees off from perpendicular to the trap axis, but the physics holds after the requisite projections are made.

²For a lovely and rigorous mathematical description of all the intermediate steps, I refer the reader to [161].

where $B_{x(z)}$ denotes the gradient of the magnetic field projected onto x(z), and $b_{j,n}q_0^j$ is the participation of ion n in mode j (q_0^j is the same as z_0 in Chapter 2, defined now for mode j).

Looking at the first term in Eq. 3.3, we can see that this has a resonance at $\omega - \omega_0 = 0$ and thus drives carrier transitions with a rate proportional to the magnetic field along that axis when the frequency of the magnetic field is near the qubit frequency. The second term modulates the energy of the qubit splitting at the frequency of the magnetic field. Both of these carrier terms are proportional to the magnetic field strength at the ion's position.

The second pair of terms describe sideband interactions, as we can see because they contain both spin operators $\hat{\sigma}_i$ and motional operators \hat{a} and \hat{a}^{\dagger} . The first one has resonances at $\omega = \omega_0 \pm \omega_j$, i.e. motional sidebands about the qubit frequency. This is the term that can be used to drive $\hat{\sigma}_{\phi}\hat{\sigma}_{\phi}$ interactions when two sidebands are applied [76, 77]. The other term has resonances at $\omega = \pm \omega_j$, and can be used to drive $\hat{\sigma}_z\hat{\sigma}_z$ interactions. Both of these terms are proportional to the strength of the magnetic field gradient. In the context of trying to perform an entangling gate, it is not desirable to drive off-resonant carrier transitions, so typically the magnetic field is nulled to minimize the Rabi frequencies described in Eq. 3.4 and Eq. 3.5. This can be done passively, for example by using a single specially-shaped resonator [107, 154, 156, 162] or actively using multiple current-carrying electrodes [8, 43, 154]. The static gradient case can be thought of as a special case of this interaction with $\omega = 0$, and we can see why we said earlier that the field gradient acts as an effective Lamb-Dicke parameter: now that is the factor dictating the speed of the sidebands relative to the carrier. Now that we have shown that we can drive sidebands using microwaves, let us look at some applications of this capability.

3.2 Ground state cooling using laser-free spin-motion coupling

It is important to note that I did not call this section 'laser-free cooling.' Cooling requires dissipation, and rf and microwave transitions do not dissipate energy. However, we can perform an analog of resolved sideband cooling using our laser-free motional sidebands and a low-power

repump beam.³ Compared to the pair of high-power 280 nm Raman beams historically used for ground-state cooling in this apparatus [139, 163], it reduces the requirements for laser power and control.

Magnetic field gradient sideband cooling had been performed previously with both oscillating [43] and static gradients [164, 165], but the oscillating gradient scheme in [166] reaches a lower temperature ($\bar{n} = 0.09$) in an order of magnitude less time (2.5 ms) than the static gradient examples. In [166] they used magnesium ions, and we have since extended this technique to calcium, as reported in Chapter 4.

The basic protocol is as follows:

- Cool ion(s) to the Doppler limit, either directly or sympathetically
- Prepare the ion in its bright state
- Perform a laser-free motion-subtracting sideband pulse
- Perform a repumping step to return the ion to the bright state
- Repeat steps 3-4 until desired temperature is reached

Demonstration of this cooling protocol on calcium is presented in Chapter 5. Ground state cooling is helpful to reduce certain types of gate errors, and crucial for the quantum logic operations described in Chapter 8.

3.3 Laser-free single ion addressing

Having explained how laser-free sidebands can be generated with sufficient strength to drive useful interactions, there is still an open question of how these interactions, which are generated with global fields, can be used to address individual ions and pairs in a larger system. In laser-based systems, individual addressing is achieved using tightly focused beams, but microwaves cannot be focused in the same way. A static magnetic field gradient can be oriented along the ion crystal axis,

³Perhaps we can call it 'reduced-laser sideband cooling.'

splitting each qubit in frequency space and allowing straightforward individual addressing [152, 167]. This has been successfully used with extremely low crosstalk [47]. However, the frequency splitting that allows each ion to be individually addressed also means that the drive fields for each ion must have different frequencies, creating substantial technical overhead and a bandwidth crowding problem for larger systems.

Another proposed strategy for individual addressing with microwaves is to have different trap zones, such that fields can be individually nulled within the different regions [154]. This has also been successfully demonstrated [168] but requires that the ions be physically moved between zones and thus adds some transport overhead to any algorithm.

In 2013, Warring et al. described four different methods for individual addressing with microwaves based on twisting an ion crystal such that one ion in a pair leaves the rf and/or oscillating magnetic field null, and taking advantage of the different field environments between the two ions in various ways [46]. The method we use is based on this crystal-twisting technique, although thanks to some special features of our gate we do not have to twist the ion crystal in and out of the null but can leave it twisted. This method was demonstrated in [8] and [163].

First the ion crystal is twisted so that each ion will experience a different magnetic field from the gradient once it is applied. Then a global $\pi/2$ pulse is applied, and we apply current to just one of our three microwave electrodes such that the ions begin precessing at slightly different rates due to their differential ac Zeeman shift. A second global $\pi/2$ pulse can then be timed such that the ions have acquired a π phase difference.

This technique was used in conjunction with the gates, described below, to create a high-fidelity entangled singlet state of two ions, something that can only be done using individual addressing. We have also reproduced this individual addressing scheme in calcium ions, presented in Chapter 8. An alternative strategy for laser-free individual addressing is presented in [49], based on localized electric fields.

3.4 Laser-free gate implementations

I stated above that these laser-free sideband interactions could be used to drive different styles of entangling gates but I did not detail how. Here we will look at the experimental demonstrations of different laser-free gates. These gates can be divided into static and oscillating gradients, and the oscillating gradients can be further split by the frequency of oscillation. Most of these gates utilize the Mølmer-Sørenson gate scheme described in Chapter 2 where blue and red sidebands are simultaneously applied.

3.4.1 Static gradient

When performing a gate using the static gradient scheme, the gradient is often created using permanent magnets. This has the advantage of not requiring strong drive currents to create a strong gradient, but permanent magnets are not tunable once in vacuum and are vulnerable to temperature fluctuations affecting field gradients. This takes advantage of Mintert & Wunderlich's initial insight that the microwave sidebands themselves do not need to have a strong gradient, which is advantageous since generating high power GHz currents can be difficult. A gate using the static gradient scheme was first demonstrated using a $\hat{\sigma}_z\hat{\sigma}_z$ -style interaction in 2012 [169], with a fidelity of 0.64(5) achieved in 8 ms, and a higher-fidelity demonstration followed in 2016 [48] that reached a fidelity of 0.985(12) with a gate duration of 2.7 ms. Recent work demonstrated an order-of-magnitude speed up in static magnetic field gradient gates, achieving a fidelity of 0.98(+2,-3) in 313 μ s using a gradient of only 19 T/m [158].

3.4.2 Near-qubit frequency oscillating gradient

The first demonstrated laser-free entanglement used a pair of magnetic field gradients oscillating near the qubit frequency to perform a $\hat{\sigma}_{\phi}\hat{\sigma}_{\phi}$ -style gate and achieved a fidelity of 76(3)% in 400 μ s. A high-fidelity (99.7(1)%) demonstration followed in 2016 [155] with a gate duration of 3.25 ms. In 2019, Hahn *et al.* reported a faster gate of 808 μ s, achieving a fidelity of 98.2(1.2)% [156]

and Zarontonello *et al.* demonstrated a higher-fidelity gate of 99.7(1)% in just under 3 ms [157]. Recently, a fast and high-fidelity demonstration by Oxford Ionics was presented at DAMOP 2024, claiming a lower bound on their fidelity of 99.97% in 120 μ s [9].

3.4.3 Near-motional frequency oscillating gradient

Generating high-power fields at the GHz frequencies needed to drive many hyperfine qubits can be more technically challenging than at few-MHz frequencies. As frequencies get higher and wavelengths get shorter, components such as mixers and amplifiers get more expensive, skin depths get smaller and thus heat dissipation gets larger, and induced return currents increase, reducing the amount of gradient achieved for a given applied current. However, an oscillating magnetic field gradient is often accompanied by an oscillating electric field, and when near-resonant with a motional frequency this can cause enough coherent motional excitation to rapidly eject an ion from the trap.

The high-fidelity laser-free gates reported in [8] used an alternate approach, presented in [149], that combines a strong near-motional frequency oscillating magnetic field gradient with a pair of microwave sidebands in order to drive either $\hat{\sigma}_z\hat{\sigma}_z$ or $\hat{\sigma}_\varphi\hat{\sigma}_\varphi$ interactions without lasers. Typically, gate detunings are on the order of a few kHz. Given the gradient strengths necessary to drive a laser-free entangling gate, this is not enough detuning to avoid exciting the motional mode if the gradient is oscillating near the motional frequency. By moving the gradient a few megahertz away from the motional frequency and bridging that extra detuning using a pair of microwave sidebands detuned symmetrically about the qubit frequency, this strategy can combine the advantages of both techniques above: as in the static gradient case, the microwave power needed is lower; but just as in the near-qubit frequency oscillating gradient case, our qubits can stay at the same frequency which simplifies control. One way to understand the gate is that in the interaction picture with respect to the oscillating gradient, it looks like the static gradient gate with all motional frequencies shifted by the gradient frequency. The theoretical underpinnings and experimental details of this gate are presented in great detail in [149] and [8] respectively, so here we will just provide a brief overview of

the details relevant to this work, using the same notation used in those papers.

In the lab frame, the Hamiltonian for the gradient oscillating at ω_g and the two microwave tones with detuning δ from the qubit frequency is

$$H(t)_{lab} = 2\hbar\Omega_{\mu}\hat{S}_x[\cos((\omega_0 + \delta)t) + \cos((\omega_0 - \delta)t)] + 2\hbar\Omega_g\hat{S}_z(\hat{a} + \hat{a}^{\dagger})$$
(3.8)

where \hat{S}_i is the multi-ion spin operator $\hat{S}_i = \sum_i \hat{\sigma}_i$, and the motional frequency ω has gained the subscript r to differentiate it from the qubit frequency ω_0 . Ω_{μ} is the Rabi rate for the microwave sidebands, identical to Ω_x from Section 3.1.2 but relabeled now to match the notation in [163], while Ω_g is the gradient Rabi rate, identical to $\Omega_{j,n}^z$ from Section 3.1.2 and relabeled for the same reason.

In the ion frame, the gate interaction is [149]

$$H(t) = 2\hbar\Omega_{\mu}\cos(\delta t)\hat{S}_{x} + 2\hbar\Omega_{q}\cos(\omega_{q}t)\hat{S}_{z}[\hat{a}e^{-i\omega_{r}t} + \hat{a}^{\dagger}e^{i\omega_{r}t}]. \tag{3.9}$$

From here, we transform into the bichromatic interaction picture (i.e. the interaction picture with respect to the red and blue microwave sidebands) and perform a Jacobi-Anger expansion to arrive at

$$\hat{H}_{I} = 2\hbar\Omega_{g}\cos(\omega_{g}t)[\hat{a}e^{-i\omega_{r}t} + \hat{a}^{\dagger}e^{i\omega_{r}t}](\hat{S}_{z}\left[J_{0}(4\Omega_{\mu}/\delta) + 2\sum_{n=1}^{\infty}J_{2n}(4\Omega_{\mu}/\delta)\cos(2n\delta t)\right] + 2\hat{S}_{y}\sum_{n=1}^{\infty}J_{2n-1}(4\Omega_{\mu}/\delta)\sin([2n-1]\delta t)).$$
(3.10)

If we now pick $2n\delta \approx |\omega_r - \omega_g|$ and n = 1, we make a rotating wave approximation such that this simplifies into

$$\hat{H}_I \approx \hbar \Omega_g J_2(\frac{4\Omega_\mu}{\delta}) \hat{S}_z[\hat{a}e^{i\Delta t} + \hat{a}^\dagger e^{-i\Delta t}]$$
(3.11)

where Δ is our gate detuning such that $\Delta = 2n\delta - (\omega_r - \omega_g)$. This now looks like a geometric phase gate where we have a Bessel function modifying our Rabi frequency.

4

Mixed species trapped ion quantum information

In the earliest schemes envisioning how a large-scale trapped-ion quantum computer might work, scientists proposed using different ion species for different tasks [23, 24]. Oftentimes a single ion species may not have every property desirable in a system, so two species may be used in conjunction. Additionally, having separate species eliminates crosstalk errors between ions, important for long-term memory and mid-circuit readout, and sympathetic cooling allows longer sequences of operations to be performed without direct cooling, which scrambles quantum information. Finally, the technique of quantum logic spectroscopy makes available ion species which do not have an accessible cycling transition by using a second species of ion to prepare and readout the state [38, 84]. An alternative method for realizing some of these desirable features is the *omg* architecture, which involves encoding qubits in the different optical, metastable, and ground states of a single ion or atom species. As this offers some benefits of mixed-species operation without the attending technical overhead, many groups have recently implemented this scheme both within and outside of the trapped ion community [170–174].

The benefits of mixed-species operation are compelling enough that many groups pursue it in spite of the difficulties [25, 33–37]. Especially as the toolbox of quantum control extends its reach to more exotic species of ions, molecules, and even subatomic particles that may not be amenable to direct control, the ability to manipulate one species by means of another will continue to be essential for not only quantum information but also investigations of fundamental physics [21, 22].

Adding a second ion species can provide many advantages, although there are complications to consider. We will begin this chapter with a discussion of the character of motional modes of mixed-species trapped ion crystals, followed by a description of techniques for mixed-species quantum information. I will focus primarily on the two most relevant techniques for our plan to create a mixed-species quantum computer with one laser-free species: sympathetic cooling and quantum logic spectroscopy. I will also briefly discuss generation of entanglement between different ion species.

4.1 Motional modes for mixed-species ion crystals

A major consideration when using mixed-species ion crystals for quantum information is the effect on the motional modes of having mismatched ion masses. We will focus here on the simplest case as it is also most relevant to our planned experiments: a two-ion crystal with a single ion of each species.

In Chapter 2, we looked at a single ion in a harmonic potential with degenerate radial modes. Now we will, following [26] and [27], break our radial mode degeneracy and look at motional modes for ion crystals with different ion masses. Because we are now breaking our radial mode degeneracy, we cannot set U = 0 and thus the stability parameter a is also non-zero. The full potential including the axial confinement is

$$\Phi(x, y, z, t) = \frac{V}{2} \cos(\Omega_{rf} t) \left(1 + \frac{x^2 - y^2}{R^2}\right) + \kappa U\left(\frac{z^2 - \alpha x^2 - (1 - \alpha)y^2}{z_0^2}\right)$$
(4.1)

where U is the static potential, κ is a geometric factor, z_0 is related to the axial trap size, and

 α describes how asymmetric the radial confinement is (i.e. the degree to which the radial mode degeneracy has been broken).

The axial trap frequency for a single ion is the same as in Chapter 2,

$$\omega_z = \sqrt{\frac{2e\kappa U}{mz_0^2}}. (4.2)$$

What we previously defined as our radial frequency

$$\omega_r = \frac{q\Omega_{rf}}{2\sqrt{2}} \tag{4.3}$$

(where we remember that q is species-dependent) becomes just a parameter in the definition of our non-degenerate radial mode frequencies

$$\omega_x = \sqrt{\omega_r^2 - \alpha \omega_z^2} \tag{4.4}$$

and

$$\omega_y = \sqrt{\omega_r^2 - (1 - \alpha)\omega_z^2}. (4.5)$$

Because these motional modes frequencies are dependent on the ion mass, but all other trap parameters stay constant, by measuring the motional mode frequencies of a single ion in a given well one can calculate the frequencies and participations for a different ion (or multiple different ions) in that same well. This is especially important for experiments such as the gates described in Chapter 9 where care must be taken to avoid heating any motional modes with the powerful fields that drive our gate interactions.

For a crystal of two ions, there will be six normal modes: an in-phase (the ions move in the same direction at any given time) and an out-of-phase (the ions move in opposite directions) mode along each trap axis. If both ions have the same mass, the motion of each ion in each mode will be perfectly balanced because they are identical coupled harmonic oscillators. When the ion masses

differ, however, it gets more complicated. Along the axial trap direction, the Coulomb force between the ions couples the motion quite strongly so even for relative large mass differences, the mode participation is relatively balanced. But for the radial motional modes, the motion becomes almost completely decoupled for masses that differ by more than a factor of 4 [27].

To see this, we can continue to follow [27] and [26] to write the total oscillation along a given axis for both ions

$$x_1(t) = b_1 z_i \sin(\omega_i t + \phi_i) + b_2 z_o \cos(\omega_o t + \phi_o)$$

$$(4.6)$$

and

$$x_2(t) = \sqrt{\frac{m_1}{m_2}} (b_2 z_i \sin(\omega_i t + \phi_i) + b_1 z_o \cos(\omega_o t + \phi_o)).$$
 (4.7)

Here, $b_{1,2}$ are the normalized mode participations such that $b_1^2 + b_2^2 = 1$; $z_{i,o}$, $\omega_{i,o}$, $\phi_{i,o}$ are the amplitudes, frequencies, and phases of the in-phase and out-of-phase modes respectively. To calculate the mode participations and frequencies, the interactions between the ions and their trapping potential can be solved to lowest order and [27] parameterizes the results in terms of the reduced mass $\mu = m_2/m_1$, $\epsilon = \omega_r/\omega_z$, and

$$a = \sqrt{\epsilon^4(\mu^2 - 1)^2 - 2\epsilon^2(\mu - 1)^2\mu(1 + \mu) + \mu^2[1 + (\mu - 1)\mu]}$$
(4.8)

to give

$$\omega_{i,z} = \sqrt{\frac{1 + \mu - \sqrt{1 - \mu + \mu^2}}{\mu}} \omega_z, \tag{4.9}$$

$$\omega_{o,z} = \sqrt{\frac{1 + \mu + \sqrt{1 - \mu + \mu^2}}{\mu}} \omega_z,$$
(4.10)

$$b_{1,z}^2 = \frac{1 - \mu + \sqrt{1 - \mu + \mu^2}}{2\sqrt{1 - \mu + \mu^2}},\tag{4.11}$$

$$\omega_{i,x,y} = \sqrt{-\frac{\mu + \mu^2 - \epsilon^2 (1 + \mu^2) - a}{2\mu^2}} \omega_z, \tag{4.12}$$

$$\omega_{o,x,y} = \sqrt{-\frac{\mu + \mu^2 - \epsilon^2(1 + \mu^2) + a}{2\mu^2}} \omega_z, \tag{4.13}$$

and

$$b_{1,x,y}^2 = \frac{1 - \mu + \sqrt{1 - \mu + \mu^2}}{2\sqrt{1 - \mu + \mu^2}}.$$
(4.14)

This is all fairly overwhelming to look at, but the reality is that these are linear equations depending on just the ion masses and the measured motional modes of one of those ions in the well you wish to use.

For some typical parameters in our trap, the results of these calculations can look like Table 4.1.

Parameter	Mode 1	Mode 2	Mode 3	Mode 4	Mode 5	Mode 6
Mode Freq. (MHz)	2.252	3.422	4.111	4.201	6.153	6.876
$x_{Mg} \text{ (nm)}$	0	0.790	0	0	5.702	0
$y_{Mg} \text{ (nm)}$	0	0	0	0.628	0	5.4
$z_{Mg} \text{ (nm)}$	5.065	0	5.927	0	0	0
$x_{Ca} \text{ (nm)}$	0	-6.043	0	0	0.466	0
$y_{Ca} \text{ (nm)}$	0	0	0	-5.461	0	0.388
$z_{Ca} \text{ (nm)}$	6.331	0	-2.963	0	0	0

Table 4.1: Motional mode frequencies and excursions for example parameters in our trap (z denotes the trap axis).

To match this to the parameters above, each column represents a mode and each row is the excursion of each ion (labeled 1 or 2) along the given direction as a part of that mode. The excursion is the product of the parameters $z_{i,o}$ and $b_{1,2}$ from above. The modes are listed in order of ascending frequency, and you can determine if they are radial or axial by looking at the direction along which the ions travel as part of that mode, and in-phase versus out-of-phase can be determined by comparing the sign of the excursion between the ion species - opposite signs mean the ions are moving in opposite directions, i.e. they are out-of-phase. The primary thing to take away from this

table is that our radial mode participations are vastly unequal, which will have consequences for the rest of the operations presented in this chapter.

4.2 Sympathetic cooling

Some of the earliest proposals for using multiple ion species in a given processor planned to use one species as a 'coolant' helper ion. In fact, sympathetic cooling of one ion species using another predates proposals for trapped ion quantum information [175]. Since quantum information is stored in the spin states of a given ion, and cooling necessarily involves dissipative transitions between different spin states, direct cooling of an ion cannot be performed without destroying any quantum information it may be storing. In addition, heating of ions involved in quantum computations causes decoherence and gate errors [23], and so cannot be ignored. In 2000, Kielpinski et al. expanded in detail the initial proposal from [23] of using sympathetic cooling for quantum logic with trapped ions: because the motion of ions in the same well is coupled so strongly via the Coulomb force, if one ion is directly cooled, their shared motion is cooled as well [26]. The first demonstration of sympathetic cooling near the ground state of multi-ion chain involved two calcium ions, with only one being directly illuminated [176], but the technique was soon extended to multiple species [177], and later molecules [178] and highly-charged ions [179].

One important consideration when selecting a second ion species to use as a coolant is the mass match. In order for the non-coolant ion to become cooled, it has to participate in the mode that is being cooled, and in order for a mode to be cooled, the coolant ion has to participate in it. In the axial direction the Coulomb force helps maintain strong coupling, but for radial modes participation can be very uneven. Sympathetic cooling is most efficient when masses are best-matched.

Sympathetic cooling allows for continuous cooling during long computations without transport, crosstalk, or risk of quantum information being destroyed because the cooling lasers for the helper species are far off-resonant from transitions in the qubit species. Certain encodings using the omg architecture allow for sympathetic cooling, but in general care must be taken to store any information (whether it is part of an active computation or stored in memory) in the metastable

state before cooling [170]. The *omg* architecture claims mixed species-like operation with only one set of lasers, but our ultimate goal is the same - with the exception of our photoionization laser, we are building a suite of control tools to enable mixed-species operation with only one set of lasers. Additionally, our scheme does not require careful bookkeeping of which ions in a given chain are in which qubit encoding, and is free from the potential error source that is constantly needing to coherently translate information back and forth between encodings.

Another strategy for cooling during longer sequences of operations is the same-species exchange cooling [180] developed at GTRI. This involves preparing a bank of pre-cooled refrigerant ions which are then transported to wherever they are needed. Both this technique and *omg* can provide faster cooling due to the perfect mass match of a single-species system, but have other forms of technical overhead. While these strategies offer a lot of advantages for single-species experiments, truly mixed species sympathetic cooling can be simpler to actually execute.

One major drawback of sympathetic cooling historically has been the inability to cool modes in which the coolant ion does not strongly participate, e.g. the 'breathing' mode of a three ion crystal. However, recent work has demonstrated how heat can be swapped out of poorly cooled modes using either quantum logic [29] or parametric coupling [30]. These techniques allow all modes of a mixed-species crystal to be cooled with only one set of lasers, an important step towards processors with one laser-free species. Indeed, these techniques can extend sympathetic cooling capabilities to systems where charge-to-mass ratios or mode participations cannot be matched, such as larger molecules, highly charged ions, and macroscopic charged particles.

4.3 Quantum logic

However, helper ions can be used for much more than just cooling. Other early work involving multiple species of ion comes from the arena of trapped ion clocks. In order to perform high precision spectroscopy for atomic timekeeping, a narrow clock transition is desired. But for an ion to be used in an atomic clock, an accessible transition for Doppler cooling and fluorescence detection is also necessary [38]. For example, mercury and aluminum ions both have very narrow clock transitions

suitable for high-precision spectroscopy, but they also have Doppler cooling transitions in the far ultraviolet.

In 2005, spectroscopy using quantum logic was proposed and demonstrated [38], and was quickly used to realize high-fidelity readout of the clock transition in aluminum ions, which lack an accessible cycling transition [84]. The technique requires co-trapping one ion (the 'data' ion) with inaccessible cooling and readout transitions with a second, so-called 'helper' ion that can more easily be addressed. The data ion can then be cooled using the sympathetic cooling technique described above, and the state of the data ion can be projectively prepared and read out using the helper ion. This technique has enabled the most accurate atomic clock to date [39], as well as extending quantum control to molecules [40, 41] and highly-charged ions [42].

4.3.1 Projective state preparation

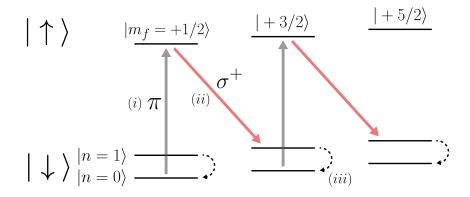
Projective state preparation is performed via a series of sidebands driven on the shared motion of the two ions. I will describe the original laser-based technique here, and the proposed laser-free variation will be presented in Chapter 8.

Often, qubit state preparation in trapped ions is performed via optical pumping, where a laser with carefully selected frequency and polarization drives population selectively into one state, typically to one side of a ground-state hyperfine manifold. For quantum logic state preparation, we need to perform a sort of optical pumping with no cycling transition. This can be done with a linearly polarized carrier and a σ^- -polarized red sideband and the experimental procedure as presented in [38] is described below.

First the bus mode to be used must be cooled to the ground state. A π -polarized carrier π pulse on the data ion drives population into the excited state, and then a σ^- -polarized red sideband π pulse, also on the data ion, maps that internal state excitation into a motional excitation on the next state up the manifold, as shown in Fig. 4.2. Cooling on the helper ion gets rid of the motional excitation, collapsing population into the ground state of motion and making the sequence irreversible. These steps are repeated until the population has been driven to the far right side of

the manifold and re-cooled to the ground state.

Figure 4.1: Figure showing the laser-based projective state preparation of ²⁷Al⁺ from [38]. Grey and red arrows denote the carrier and red sidebands. Population is moved to one side of the ground-state hyperfine manifold one step at a time using successive carrier (i) and red sideband pulses (ii) on the data ion with specifically chosen polarization, interleaved with ground-state cooling on the helper ion (iii).



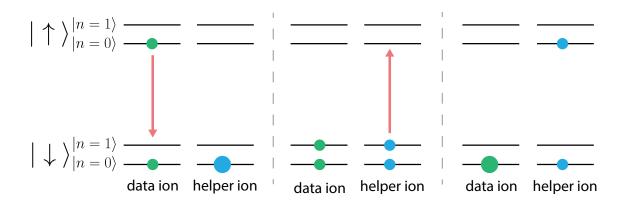
4.3.2 Quantum logic readout

The readout process operates on a similar set of principles, this time using a red sideband pulse on both ions. First, assume the data ion is in some superposition of its internal states, the helper ion is prepared in the electronic ground state, and the bus mode has been cooled to the motional ground state. A red sideband π pulse on the data ion will map the internal state superposition into a motional state superposition in the shared motion. A red sideband π pulse on the helper ion will map the superposition in the shared motion into a superposition in the internal states of the helper ion, as shown in Fig. 4.2, where it can be read out using standard fluorescence detection techniques.

While these techniques allow the use of a data ion whose cycling transition is not easily addressable or even nonexistent, the capability to drive some sort of motional sideband on the data ion is still required. Because performing quantum logic on a mixed species ion crystal with laser-free sidebands is a primary goal of this project, we will save a detailed discussion of the procedure for

Chapter 8.

Figure 4.2: Figure showing quantum logic readout as described in [38]. Both ions must be cooled to their motional ground state in preparation, and the helper ion is prepared in its electronic ground state. A red sideband π pulse on the data ion maps the data ion's internal state into the shared motion, then a red sideband π pulse on the helper ion maps the motional superposition into the internal state of the helper ion.



4.4 Mixed-species entanglement

In order to transfer information into a 'memory' helper ion for storage or a 'communication' helper ion for transmission through a photonic interconnect, it is desirable to be able to generate entanglement between ions of different species. A mixed-species entangling gate can be used to improve quantum logic readout by relaxing the requirements on ground-state cooling [38, 181]. Creating high-fidelity entanglement requires careful calibration of many different parameters, and entangling between two species nearly doubles the complexity of such an interaction. The first mixed-species entangling gates were not performed until 2015, when Ballance et al. generated entanglement between two isotopes of calcium with 99.8(6)% fidelity [182] and Tan et al. entangled a magnesium and a beryllium ion with 97.9(1)% fidelity [183].

It was only recently that mixed-species laser-based gates reached competitive fidelities with same-species laser-based gates [184], and that result was achieved by choosing ion species such that both species could be driven with a single pair of laser beams, a strategy which reduces experiment complexity but also severely limits the ion species that can be used.

If we return to the description of entangling gates from Chapter 2, we can see that in order to drive a gate on two ions of different species we will need to be able to drive red and blue sidebands on the same motional mode for both species. These sidebands will need to be carefully balanced and phase-stabilized such that interactions on each ion species are coherently driven at the same rate. For a detailed discussion of laser-based mixed-species entangling gates, I will refer the reader to [185].

For laser-based gates, stabilizing the phase of different laser beams of different wavelengths can be technically challenging. In theory, laser-free interactions should make phase stabilization, and thus mixed-species entanglement, vastly easier. But there is one major issue which needs to be overcome. All laser-based mixed-species entangling gates performed to date have been done on an axial motional mode. For mixed-species gates this is especially important as better ion participation in the shared mode of motion means faster gates. This poses an issue for laser-free gates in the apparatus described in this thesis, however, which can act only on radial modes of motion due to our trap geometry. Performing a laser-free mixed-species entangling gate is another primary goal of the project described in this thesis, and our plans for overcoming the difficulties involved are laid out in detail in Chapter 9. For now, we turn to a description of our chosen ion species and the methods we use to trap and control them.

5

Magnesium & calcium ions

The previous chapter motivated the use of multiple ion species within a single quantum system. This chapter describes the two species chosen for this system, ²⁵Mg⁺ and ⁴⁰Ca⁺, and the methods used to generate, trap, and control ions of both species.

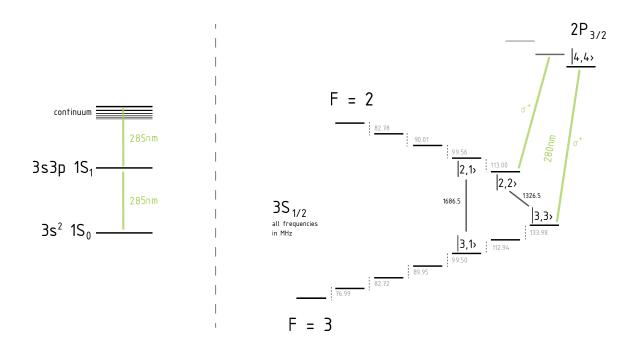
5.1 Magnesium

All past work performed using this apparatus was done with 25 Mg⁺ ions [139, 163]. The magnetic field of 212.8 G was chosen to give a magnetic field-insensitive 'clock' qubit at 1.686 GHz, but we can also use the magnetic field-sensitive 1.325 GHz 'stretch' qubit at the far end of the manifold. The stretch qubit is prepared via optical pumping, and preparation of other states is performed by means of microwave π pulses between hyperfine levels.

5.1.1 Ion production

A beam of neutral magnesium is produced by running current through a pillow-style oven filled with shavings of magnesium metal. For magnesium, we run 6.5 A through the oven and

Figure 5.1: The levels of neutral magnesium relevant for photoionization on the left, and level structure of $^{25}{\rm Mg}^+$ at 212.8 G on the right. Figure modified from Laurent Stephenson.



the typical loading temperature is 260 C as measured by the on-oven thermocouple.¹ Magnesium photoionization is a two-photon process where a single 285 nm photon resonantly drives the $3s^2$ $^1S_0 \leftrightarrow 3s3p$ 1P_1 transition, and a second photon, at either 280 or 285 nm, excites the electron into the continuum as shown in Fig. 5.1 [186].

5.1.2 Micromotion compensation

Once a magnesium ion has been trapped, we can apply electric 'shim' fields to compensate for any stray fields and push the ion towards the null of the rf field, where micromotion is minimized. We refer to this process as 'micromotion compensation.' Coarse compensation of the in-plane micromotion can be done by monitoring fluorescence as the field shims are adjusted. Because an ion experiencing large-amplitude micromotion spends a substantial portion of its time Doppler shifted out of resonance with the cooling light, it appears dimmer to the PMT. This method of compensation can determine the correct in-plane shim to within roughly 100 V/m, as shown in [163].

For fine adjustment of the in-plane and out-of-plane shim fields, we used to drive a microwave micromotion sideband on the clock qubit as described in [163]. The closer the ion was to the rf null, the more difficult it was to drive this sideband, so micromotion could be compensated by adjusting stray field shims to minimize the speed of this micromotion flop. We later added the method of parametric micromotion excitation [187]. This technique required some modifications to the trap rf drive, described in Chapter 6, as well as the transition from the Scan1D Artiq framework to NDScan to easily run multi-dimensional scans. However, it allows relatively rapid and precise calibration of stray fields while simultaneously measuring both radial mode frequencies.

The basic principle is as follows: a weak modulation is applied to the trap rf over a range of frequencies near the ion's motional modes. At the same time, one of the field shims is scanned, producing a two-dimensional plot with the ion fluorescence providing the third variable. If the ion is near the rf null, this modulation has a negligible effect on the ion fluorescence even when it

¹More details on the oven construction and thermocouples can be found in Chapter 6.

is near-resonant with a motional mode simply because being at the rf null protects the ion from modulations on the rf. However, if the ion is pushed out of the null by a stray field with projection along that mode, it will experience very strong disturbances as the modulation crosses the motional resonance, manifesting as decreased fluorescence. In this way we can zero the projection of the stray field along each of the radial modes, and when both modes experience minimal excitation at the same point the shims are correctly calibrated. In this way, by doing one quick two-dimensional scan we can quickly determine the in-plane and out-of-plane field shims to minimize micromotion and measure the frequencies of the radial modes.

5.1.3 Axial stray field compensation

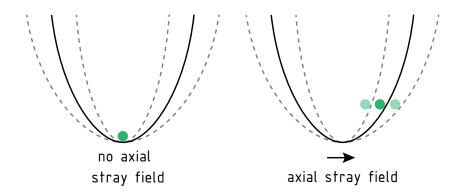
In the above section only in-plane and out-of-plane radial shims are mentioned as those are the only directions in which stray fields push the ion off the rf null and thus the only shims relevant for micromotion compensation. The axial shim moves the ion back and forth along the rf null, and so we cannot detect stray fields in this direction with the above method. Indeed, before we began trying to transport in earnest, we very rarely calibrated the axial shim. However, it becomes a key player in Chapter 7 and so we will discuss the calibration here.

Axial confinement is provided by the dc electrodes where each well is designed with the center in a specific location. A stray field along the trap axis pushes the ion out of the bottom of this well. This is not detectable if the well shape is not changed so long as the field is not powerful enough to push the ion completely out of the expected position. But if the well strength is adjusted, an ion in the center of the well will remain stationary, whereas an ion that has been pushed out will not. We calibrate this shim then by measuring deviations in ion position along the trap axis by horizontally scanning the detection beam, at different axial well strengths.

5.1.4 Cooling

Magnesium ions are Doppler cooled by the blue Doppler detuned (BDD) and blue Doppler (BD) beams, whose naming conventions and generation are described in Chapter 6. The BDD is

Figure 5.2: A stray field along the trap axis does not cause micromotion, but it does push the ion out of the bottom of the intended well. This effect is only noticeable when the well strength is being changed, at which point the ion's horizontal position will vary by an amount proportional to its displacement from the well center. By adjusting the axial shim field, we can minimize the change in ion position with respect to well strength and thus null the axial stray field.



tuned 370 MHz red of the ${}^2S_{1/2} \leftrightarrow {}^2P_{3/2}$ cycling transition and has a power of 20 μ W at the trap, which gives a Doppler temperature of nbar $\sim 2.[163]$ The BD beam can perform near-resonant cooling when tuned ~ 10 MHz red of the transition.

Historically there was a pair of perpendicular Raman beams at 280 nm for ground-state cooling in magnesium-only experiments, but we chose not to reinstall these beams when reconfiguring the optics around the trap chamber. While the beam launchers and mirrors are in place, the dichroics necessary to combine the 280 nm light with the calcium light in those ports were removed due to paranoia about beam quality and scattering. In 2024 we gave away the 559 to 280 nm doubler, which is going to make it very hard to get the Raman ground state cooling back.

Fortunately ground-state cooling and sideband thermometry can be performed using the 5 MHz magnetic field gradient in concert with motional microwave sidebands and an optical repump, as described in Chapter 3.1. [163].

5.1.5 State preparation & manipulation

State preparation into the $|F=3,m_f=3\rangle$ bright state is performed with a combination of optical and microwave pumping. The 280 nm beams for Doppler cooling and state readout are all σ^+ -polarized, meaning that Doppler cooling also optically pumps population towards the right side of the manifold.² The quality of this state preparation depends on both the purity of the polarization and its alignment with the magnetic field, so there are small shim coils that can be used to adjust the field alignment to optimize state prep (see Chapter 6).

When especially high-quality state preparation is required, a combination of microwave and optical pumping can be used to maximize population in $|3,3\rangle$ by clearing it out of $|2,2\rangle$ and $|3,2\rangle$ as follows [163]:

• π pulse $|2,2\rangle \leftrightarrow |3,3\rangle$: all previously prepared population now in $|2,2\rangle$, some leakage population possibly in $|3,3\rangle$

²A very un-fun fact that we discovered is that, if the drive for the green shift AOM decides to be 150 MHz instead of 364.5 MHz, and in desperation to trap anything we flip the polarization of the Doppler beams, we can actually also drive the cyling transition on the other side of the manifold!

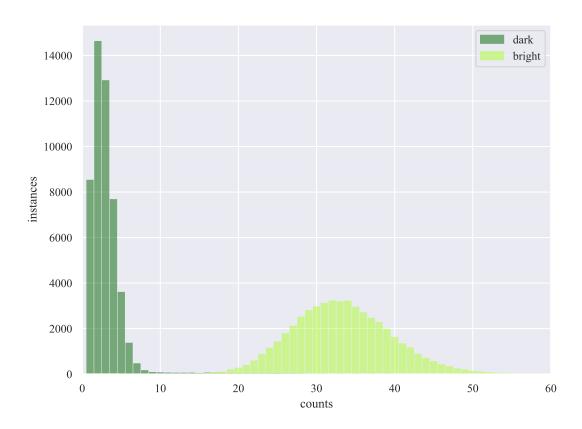


Figure 5.3: High statistics histogram showing bright and dark state preparation for magnesium.

- π pulse $|2,2\rangle \leftrightarrow |3,2\rangle$: all previously prepared population now in $|3,2\rangle$, some leakage population possibly in $|3,3\rangle$ and $|2,2\rangle$
- Pulse RD repump laser: population in $|2,2\rangle$ now mostly in $|3,2\rangle$ (some in $|3,3\rangle$)
- π pulse $|2,2\rangle \leftrightarrow |3,3\rangle$: population left in $|3,3\rangle$ now in $|2,2\rangle$ where it can be repumped
- Repeat steps 3 and 4 five times
- π pulse $|2,2\rangle \leftrightarrow |3,2\rangle$
- π pulse $|2,2\rangle \leftrightarrow |3,3\rangle$

The state preparation and measurement, when maximally optimized as it was for the randomized benchmarking experiments reported in [188],³ has an error of 1.8% for bright states (due primarily to repumping during detection) and 0.2% for dark states.

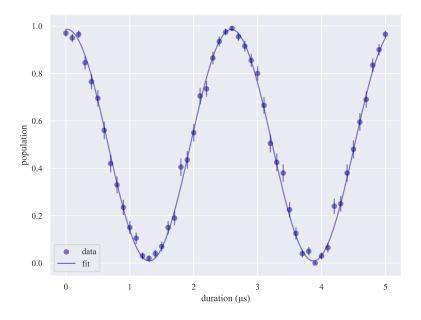
After the $|3,3\rangle$ state is prepared, experiments can be performed either on the field-sensitive stretch qubit, in which case the qubit is already in the bright state, or population can be coherently transferred anywhere within the ground state manifold using microwaves.⁴ For example, if we wish to perform an experiment on the field-insensitive clock qubit, we pulse the carrier microwaves to take us from $|3,3\rangle$ to $|2,2\rangle$, then $|2,2\rangle$ to $|3,1\rangle$. This is the 'bright' state for the clock qubit. Each microwave carrier pi pulse takes a few microseconds, depending on the transition, and the error per pulse is below 10^{-4} when properly calibrated.

5.1.6 Readout

Fluorescent state readout is performed using the blue Doppler (BD) laser beam, tuned nearresonant with the ${}^2S_{1/2} \leftrightarrow {}^2P_{3/2}$ transition. For high-fidelity state readout of the stretch qubit, population in $|2,2\rangle$ is shelved via a series of microwave pi pulses to the $|2,-1\rangle$ state.

³Publication in preparation.

⁴Though in practice, we never go further out than the $|2,-1\rangle$ shelf.



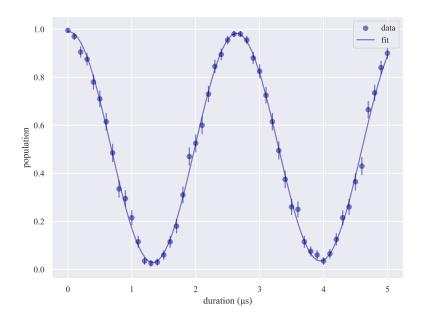


Figure 5.4: Microwave control of the 596 MHz Zeeman qubit between the $S_{1/2}$ ground state levels of calcium 40. Top: a detuning scan around the nominal qubit frequency. Bottom: a microwave flop between the qubit states demonstrating high contrast and fast π times.

In order to read out populations in the clock qubit, some care must be taken as the $|2,1\rangle \leftrightarrow$ $|3,2\rangle$ and $|3,1\rangle \leftrightarrow |2,2\rangle$ transitions are nearly degenerate (~ 100 kHz separation). This near-degeneracy can be compensated for by performing a $\pi/2$ pulse on the $|3,1\rangle \leftrightarrow |2,2\rangle$ transition, waiting 6-7 μ s, then performing a second $\pi/2$ pulse. This allows off-resonantly excited population from $|2,1\rangle$ to be returned to $|2,1\rangle$ while population that was in $|3,1\rangle$ is successfully transferred into $|2,2\rangle$. Typical bright counts for a magnesium ion in the trap used in this thesis are 225 kilocounts per second and the dark counts are ~ 5 kilocounts per second, depending on how tightly the imaging system is stopped down.

5.2 Calcium

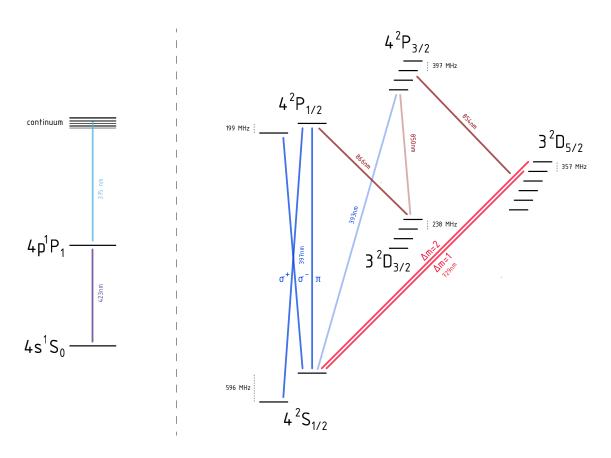
Our chosen helper ion is ⁴⁰Ca⁺. Calcium was selected because of its readily available lasers and relatively good mass match with magnesium, which is important for sympathetic cooling. However, the magnetic field dictated by the magnesium clock qubit is not one where calcium is typically run. ⁴³Ca⁺ has a zero-field clock qubit, and so often fields of just a few Gauss are used to break the degeneracy between hyperfine levels. There is a second clock qubit at 146 G, used in [161, 189], and a higher-field one at 288 G [189]. Though we use an isotope with no nuclear spin, and thus no hyperfine structure, the laser wavelengths and thus much of the general control scheme are similar - we just have fewer levels to worry about. The magnetic field using in this thesis of 212.8 G lies awkwardly between the two fields used at Oxford, and much of the initial strategy was similar to their high-field scheme. However, we later found that some of the strategies used at ETH Zurich, where they use calcium at 119 G (a magnetic field which provides a clock qubit in beryllium), ended up working better [190].⁵

The initial planning for this calcium laser system was done by David Allcock and preliminary setup was done by Shaun Burd, but Laurent Stephenson and I made extensive renovations to both the plans and existing optics, especially the 397 nm beamline.⁶ The 729 nm beamline and lock was

⁵The main differences are using two σ -polarized beams, rather than a σ and a π , and shelving with 729 nm light rather than 393 nm light.

⁶And at the time of this writing, Christina Bowers is doing yet another overhaul to switch from IntraAction to

Figure 5.5: The levels of neutral calcium relevant for photoionization on the left, and the level structure of 40 Ca $^+$ at 212.8 G on the right. Calcium is not typically used at this magnetic field, so there are some correspondingly atypical features of the laser system. Figure modified from Laurent Stephenson.



initially set up by Alejandra Collopy and the lock was later remodeled by Laurent Stephenson for improved stability.

5.2.1 Ion production

Neutral calcium is produced using a pillow-style oven run at a current of 7 A. The typical loading temperature as measured by the oven-mounted thermocouple is around 270 C. This loading temperature was selected as being slightly below the threshold where we begin to see a neutral fluorescence signal.

Photoionization of calcium is a two-photon process that requires two separate lasers: the first step $4s^1S_0$ to $4p^11$ is resonantly driven by a 423 nm laser [191], and excitation to the continuum is performed by a 375 nm laser. We typically use 150 μ W of 375 nm light and 120 μ W of 423 nm light, and both of these beams are focused to a waist of $\sim 15~\mu$ m. Though we have found that having at least 50 μ W of light in the far-detuned 397 cooling beam (analogous to the BDD beam in magnesium) leads to higher loading rates, we did successfully trap calcium ions for a year and a half before installation of this beam so it is not strictly necessary.

5.2.2 Micromotion compensation

Micromotion compensation of calcium ions is performed identically to that of magnesium ions. However, we initially had a great deal of difficulty finding compensation values that worked well for calcium, we now suspect due to the very different charge landscapes that occur when loading the different ion species. This is discussed in more detail in Chapter 7.

The parametric micromotion compensation technique is most sensitive to in-plane micromotion and works best when you are already near the rf null. If compensation is extremely far off, the scans become garbled and difficult to interpret. When we first began working with calcium, the compensation was so far off that we spent nearly a week going in circles alternately compensating the in-plane micromotion using the parametric excitation and the out-of-plane micromotion using Brimrose AOMs.

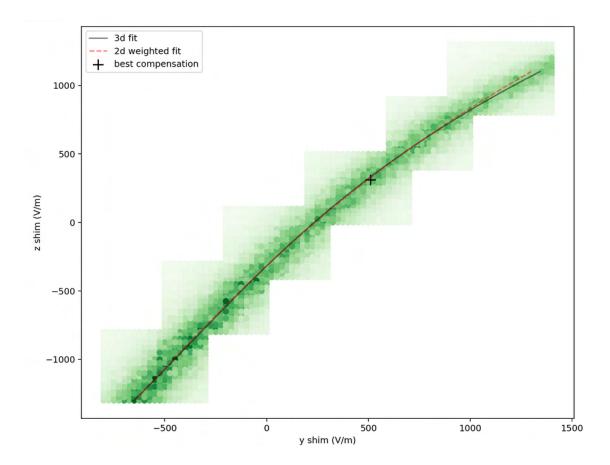


Figure 5.6: In order to compensate the micromotion from scratch, we needed to find the point where the in-plane and out-of-plane micromotion minima coincided. This plot shows a composite scan of the ratio of the carrier to the micromotion sideband as both in-plane (y) and out-of-plane (z) shims are scanned over a large range (note that the trap axes used in these scans are different from those in Chapter 2). Darker green represents a larger carrier:sideband ratio, i.e. lower micromotion. This demonstrates the linear relationship between in-plane and out-of-plane shims where in-plane micromotion is minimized.

the micromotion flop method described for magnesium.

In order to accurately determine compensation starting from zero, this is what we found to be useful. Firstly, perform a two-dimensional scan of ion fluorescence as a function of both in-plane and out-of-plane shims. Ion loss is very common when scanning shim fields over large ranges for a poorly-compensated ion, so it is easiest to do these scans in relatively small chunks. Because in-plane and out-of-plane micromotion are not truly independent, we find a relation as shown in Fig. 5.6. Then, we step along those pairs of in-plane and out-of-plane shim values, checking the out-of-plane micromotion using the flop method. There is only one true micromotion null, and so there is one minimum of out-of-plane micromotion along the curve of minimal in-plane micromotion. Fortunately we only needed to complete this process once as the calcium stray fields are relatively stable, but I include it here as a cautionary tale.

5.2.3 Cooling

Doppler cooling of the calcium ions is performed by both the high-power, far-detuned 397 nm beam as well as the near-detuned σ -polarized beam pair. The far-detuned light is tuned 440 MHz below the 397 nm σ^- beam for far-off-resonant Doppler cooling. The σ -polarized beam pair can be detuned using the common 'shift' AOM to perform near-resonant Doppler cooling.

Resolved sideband cooling can also be performed on the 729 nm quadrupole transition. Sideband cooling works very well but it slow - in order to perform rapid cooling of many motional modes in parallel, we implemented electromagnetically-induced transparency (EIT) cooling as well. This utilizes both the 729 and the weak 397 nm π polarized beam. Though EIT does not reach as low of a base temperature as sideband cooling, it can simultaneously cool multiple modes very quickly, and then the modes of interested can be further cooled with sideband cooling. In this way the total amount of time that must be dedicated to ground-state cooling can be reduced. Results of EIT cooling on the modes of a single calcium ion are shown in Fig. 5.8.

Finally, we have one last method for cooling the calcium ions which is the microwave sideband cooling described in [163], here implemented in calcium for the first time (to my knowledge). The

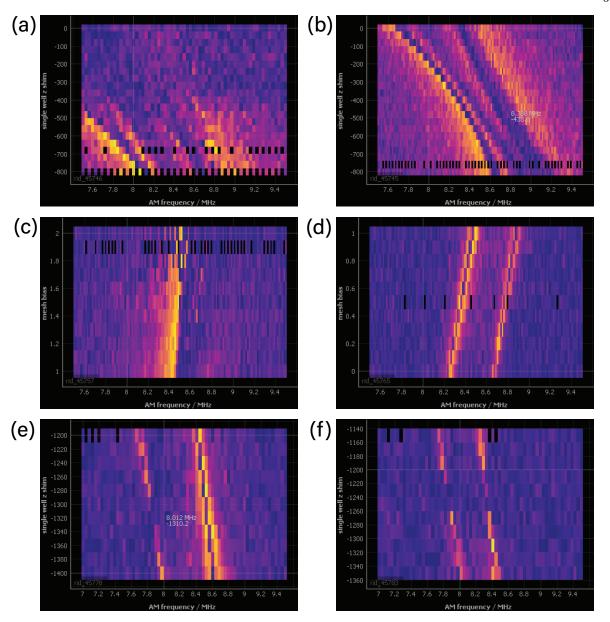
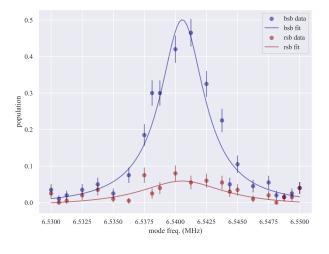
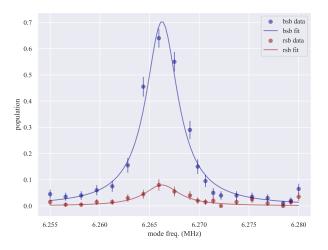


Figure 5.7: A series of parametric micromotion scans showing the process of compensating from far off. Blue regions are where the ion fluorescence is not disturbed by the applied modulation, where yellow indicates regions of high disturbance. We alternate between scanning the out-of-plane (z) shim and the mesh voltage as they act in the same way: 1 V applied to the mesh is equivalent to -250 V/m of applied z shim. Using the mesh can be very helpful when compensating very large shims in multiple directions to reduce the fields that need to be supplied by the dc electrodes. (a) The scan is garbled, no clear mode signatures visible. The stronger signal at the bottom of the screen suggests that we need to move in that direction. (b) The modes are visible but they are not clear or separated - possibly very far off OR being driven too hard. (c) There is a cleaner signal but it does not look like two clean, separate modes. (d) There are now two clear, separate modes, though neither minima is within the scan range. (e) There is now one minimum within the scan range. (f) Now both minima are centered in the scan range and at the same shim, which means the shim not being scanned is at the correct value.





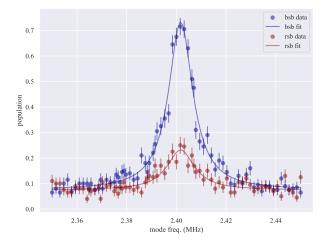


Figure 5.8: Red and blue sideband measurements for a single calcium ion after cooling both radials (top) and the axial mode (bottom) using just EIT cooling.

method is similar to that used for magnesium except that the repumping portion of the cooling is performed using 397 nm light. We drive a motion-subtracting sideband on the mode to be cooled using the strong 5 MHz gradient and a weak microwave sideband, then pulse the σ -polarized 397 beams to dissipate energy and return the population to $|\uparrow\rangle$.

5.2.4 State preparation and manipulation

Calcium 40 has no nuclear spin and thus no hyperfine structure, so we only need to calculate the Zeeman splitting of its levels. The magnetic field sensitivity of the $S_{1/2}$ is 2.8 MHz per G, so at 212.8 G the ground state Zeeman qubit has a splitting of 596 MHz. The $P_{1/2}$ levels are split by 199 MHz, so in order to drive what is typically a single cycling transition at low field, we need two tones separated by nearly 800 MHz as shown in Fig. 5.5. This is similar to what was done at ETH [190], although they had π -polarized beams to connecting both pairs of $S_{1/2}$ and $P_{1/2}$ levels where we only have one. There is typically have \sim 10-20 μ W in each beam, depending on whether we are performing Doppler cooling or detection, in a beam with a 15 μ m waist.

The splitting of the ground-state qubit means that state preparation into either the $|\uparrow\rangle$ or $|\downarrow\rangle$ qubit states is very simple: turn off the beam pumping out of the state you wish to prepare. From here, the qubit can be manipulated by means of a carrier microwave at 596 MHz (described in Chapter 6), and we can also drive microwave motional sidebands using a combination of a strong oscillating magnetic field gradient and a weak microwave tone, as described in Chapter 3.

The $D_{3/2}$ and $D_{5/2}$ levels are also split further than in typical calcium ion experiments. Rather than using multiple frequencies and polarizations of the 866 nm repump light as was done at ETH [190], we use an excess of power (500 μ W in a 36 μ m beam waist) to power-broaden the transition enough to clear out all the $D_{3/2}$ levels. A similar strategy is employed with the 854 nm repump light that clears out the $D_{5/2}$ shelf, where we use 440 μ W in a 32 μ m beam waist.

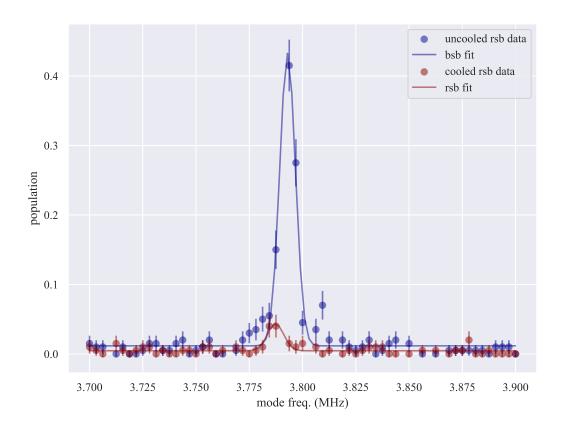


Figure 5.9: Measurement of the red sideband of a radial mode of a single ion of calcium without microwave sideband cooling (blue) and with microwave sideband cooling (red). Both measurements are performed with the red microwave sideband to eliminate possible miscalibrations between the red and blue sidebands due to the large AC Zeeman shifts that can be generated by the microwave sidebands, so we cannot strictly extract an \bar{n} from this measurement.

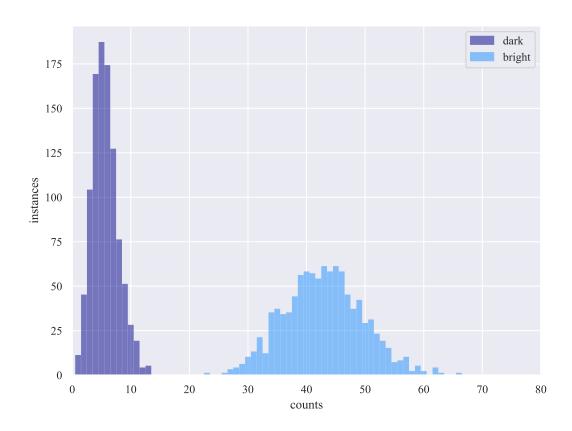
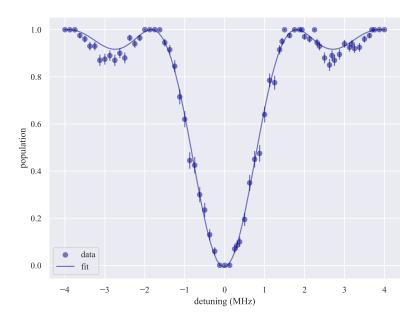


Figure 5.10: Histogram showing the bright and dark state preparation for calcium.



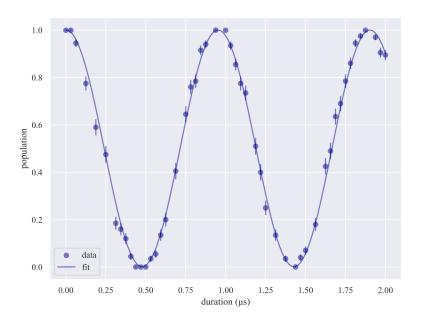


Figure 5.11: Microwave control of the 596 MHz Zeeman qubit between the $S_{1/2}$ ground state levels of calcium 40. Left, we see a detuning scan around the nominal qubit frequency and right, a microwave flop between the qubit states demonstrating high contrast and fast π times.

5.2.5 Readout

In order to readout the qubit, we need to first shelve the dark state into the $D_{5/2}$ manifold using the 729 nm laser. We can then turn on both of the 397 σ beams to drive the cycling transition and collect fluorescence. Typical bright counts for a calcium ion in the trap are 220 kilocounts per second. Because the ion goes dark within just a few photon scatters without the 866 repump, we can perform background-subtracted detection by comparing the levels of 397 nm scatter with the 866 on versus off. Typical background levels are 20 kilocounts per second.

5.3 Trap stability & loading for mixed species crystals

One of the biggest difficulties of operating a single trap with multiple ion species is the compromises which must be made between the different stability parameters for each species. Because the masses are 25 amu and 40 amu, nearly a 1:2 ratio, the range in which both ion species are stably trapped is much narrower than for a single species.

We have found that while both species are happy at a wide variety of parameters⁷ once trapped, the initial trapping conditions are much more difficult. See Chapter 7 for (much) more discussion of actual mixed species loading parameters in this trap.

⁷Like, decades of rf levels. Very wide.

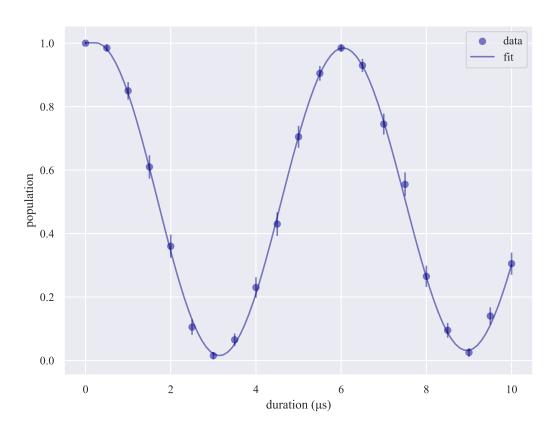


Figure 5.12: Flopping on the 729 nm quadrupole transition used for shelving and sideband cooling.

6

Apparatus

By now we have hopefully motivated the two main techniques driving the work in this thesis. In this chapter we will describe the apparatus that was designed and built to execute these experiments. This system has a few unique features dictated by the work we hope to do. Because the strength of the laser-free interactions increases as the ions get closer to the trap, we work very close to our trap surface - our ion height is 30 μ m, whereas even in other surface traps¹ people prefer to work at heights closer to 100 μ m [35, 161]. This is because problems such as anomalous heating get worse with proximity to surfaces: anomalous heating scales as $1/d^4$, where our gate interaction strength only scales as $1/d^2$. In order to reduce the heating rates and thus improve ion lifetimes, we operate the experiment at cryogenic temperatures. Cryogenic operation has been proven to drastically reduce heating rates [122], as well as improving ion lifetime by reducing background gas through cryopumping.

Because the apparatus as it existed in 2020 was thoroughly covered in [139] and [163], I will focus primarily on the various technical upgrades that have been made since and will gloss over topics that are covered in much more detail in those documents. Many of these upgrades were necessary

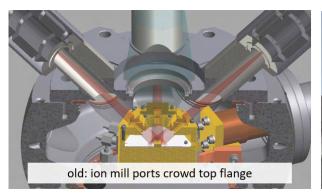
¹Even when doing laser-free gates, as in [161]!

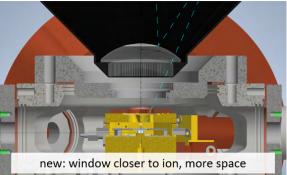
to expand the project to use a second ion species, but we chose to make additional improvements when we discovered a cracked rf connection in January 2022 that necessitated breaking vacuum. A thorough comparison of the old vacuum chamber and pillbox to the new one can be found in Appendix A. The original system was designed and built by David Allcock, Daniel Slichter, Shaun Burd, and Raghavendra Srinivas. The modifications described in this chapter were performed by Alejandra Collopy, Laurent Stephenson, Daniel Slichter, and myself.

6.1 Vacuum chamber

The rf breakage discovered in January 2022 offered the opportunity² to make several upgrades to the vacuum system described in [139, 163].

Figure 6.1: Cutaway CAD models of the vacuum chamber showcasing the difference between our old and new top flanges. On the left, the argon ion guns and the amount of space required above the trap to provide adequate line of sight. On the right, the new top flange is re-entrant, which allows the achromatic objective to get close enough to the trap. Not needing access for the ion guns anymore freed us up to install a new mesh much closer to the trap surface, which can also be seen in this rendering.





6.1.1 Re-entrant top viewport

One of the nicest upgrades in terms of mixed-species operation was a re-entrant top viewport on the octagon which allowed us to exchange the existing imaging system for an achromatic

²Whether I wanted it or not.

Schwarzchild objective (more details below). Our old top viewport also featured four argon ion guns, which were intended for milling the trap surface if the heating rates were higher than desired. However, [192] discovered that ion milling of trap surfaces can actually increase heating rates for gold trap surfaces at cryogenic temperatures, and so the guns were never fired. These ion guns occupied a lot of space above the trap chamber, and removing them increased the space available for our imaging system, as can be seen in Fig. 6.1.

6.1.2 Optical access viewport modifications

The original viewports were coated for 280 and 285 nm light, and had stray reflections of 2-5% at calcium wavelengths. We chose to replace the viewport through which the 397 nm σ beams for cooling and detection enter the chamber with a homemade viewport made with an Edmunds UV-VIS coat 25 mm diameter, 2 mm thick window (see Appendix A for more details on the assembly of this viewport) in order to reduce stray 397 light both on the table and in the chamber.

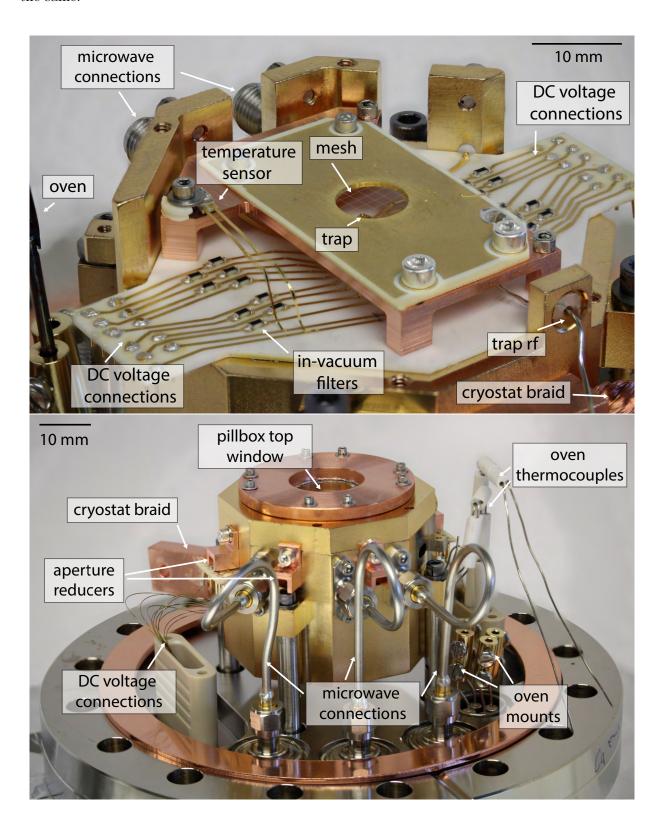
6.1.3 Pillbox modifications

This system was not initially intended to operate at liquid helium temperatures, and so does not have any heat shielding around the pillbox. We chose to add a top window (Edmunds UV coat 25 mm diameter, 2 mm thick, mounted with an indium gasket to prevent cracking during temperature cycling) as well as reducers on all of our optical access apertures which decrease our line-of-sight to room temperature by 70%. The goal of this was to reduce the gasload inside our pillbox from blackbody radiation; however it also reduced conductance out of our pillbox (see Appendix B for more details on why this may be an issue).

6.1.4 Bias-able mesh

Previous iterations of this experiment utilized a grounded mesh 16 mm above the trap surface, which served to shield the trap from any potential charging on the top viewport. The new mesh we installed is 4 mm above the trap surface and wired such that a bias voltage can be applied.

Figure 6.2: Annotated photograph of the trap pillbox with the lid off and with it on, from different angles. Note that it is in different stages of assembly in the different photos, so not all elements are the same.



Without this capability, we would occasionally reach a situation where we had such extensive stray charges built up on our trap surface that we could no longer compensate them enough to trap ions. This mesh allows us to put a voltage 'lid' on the trap which prevents that from happening, and has proven extremely useful for trapping calcium in particular. We can apply bias voltages of ± 10 V (although in practice applying a negative bias to the mesh just sucks any ions up out of the trap), and each volt applied on the mesh corresponds to 250 V/m of effective negative z shim.

6.1.5 Ovens

When we opened the chamber we needed to make a new calcium oven as calcium oxidizes rapidly when exposed to air, and we chose to also make a new magnesium oven as the current one had been in service for a long time. We took the opportunity to attach thermocouples (Accu-Glass 112114 Type K) to both ovens as is done in the Oxford ion trapping group [193]. These thermocouples were easy to install and have been invaluable for debugging our loading, especially when we had an out-of-vacuum short that was preventing one oven from heating up. They also allow us to do simple 'bang-bang' temperature control during loading, to keep our oven temperatures as low as possible while still loading (previously, oven current was left on full blast and so ovens could get quite hot during repeated load attempts³). This helps both to prolong oven lifetime and to minimize local heating, which could potentially than cryopumped gases. The ovens are otherwise the same design previously used: a thin (1.65 mm OD, 1.39 mm ID) stainless steel (AISI 316L) tube, 10 mm long with a 0.66 mm hole drilled in one side is pinched closed on one end, filled with shavings of the metal of interest (i.e. calcium (enriched to 20% Ca-43 from ORNL⁴) or magnesium (natural abundance, Alfa Aesar)), pinched closed on the other end and then gap welded onto thick (1.6 mm diameter) stainless steel (AISI 316) wires. When current flows through the oven, the relative difference in conductivity between the thicker wire and the hollow tube means that most of the resistive heating occurs in the pillow where it also heats the metal shavings inside. The hole drilled

³But we don't know how hot because we didn't have on-oven thermocouples!

⁴In case we ever wanted a helper ion with a hyperfine qubit.

into the front side of each pillow allows a beam of hot neutral flux to escape the oven and travel over to the trapping region where it can be ionized and trapped.

6.1.6 Vacuum pumps

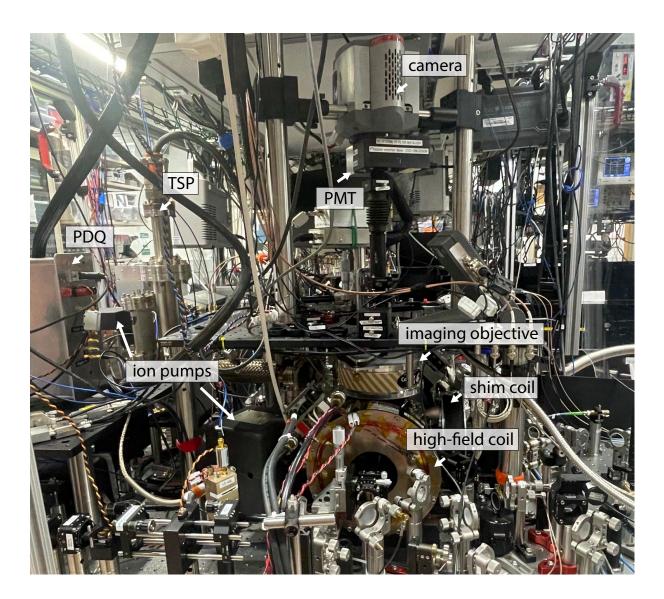
The old system had a single Agilent VacIon 20 L/s ion pump, a non-evaporable getter (NEG), and a titanium sublimation pump (TSP) with a cryoshroud to maintain good vacuum pressure. The TSP portion was separated from the main chamber by an all-metal valve, to allow replacement of TSP filaments without venting the main chamber, as well as for paranoia against leaks caused by the high temperatures involved when firing the TSP. However, the cryoshroud around the TSP occupied a substantial amount of optical table space and reduced conductance between the TSP and the main experiment chamber. We chose to eliminate the cryoshroud when rebuilding the system (see Appendix A), which opened up table space around the chamber as well as improved conductance between the TSP manifold and the main chamber. After the rebuild, we were unable to do as hot of a bake as when this system was first built because we had introduced indium joints in several places (see Appendices A and C). As a result of this, our room temperature vacuum pressure was initially much worse than we had hoped. This was solved by adding an additional 40 L/s Agilent VacIon ion pump to the TSP manifold portion of the vacuum chamber.

6.2 Cryostat

Prior to March 2021, this experiment was cooled with a Janis ST400 flow cryostat. However, due to rising helium costs, we switched over to a ColdEdge Stinger ultra-low-vibration closed-cycle cryocooler. The Stinger bayonet can be inserted directly into the existing flow cryostat connection, so disruption to the rest of the apparatus was minimal. However, reaching a sustainable operating point with the new system required substantial trial and error, which is detailed in Appendix B.

Our Stinger system (shown in Fig. 6.4) consists of two compressors, a heat exchanger that handles the water cooling for both compressors (protecting them from the building cooling water), and the coldhead itself. While the coldhead needs to be in the lab, the other equipment can be

Figure 6.3: Photograph of the experimental apparatus showing one high-field and one shim coil, some magnesium optics, the imaging system, ion pumps and TSP, and one of our microwave PDQs.



installed outside the lab to reduce crowding, heat load, and noise in the science spaces. The coldhead is mounted inside a sound-insulating box that we packed with additional noise-damping material to further decrease noise in the lab, and two sets of Aeroquip lines run to the two compressors in the galley. These lines are wrapped in sound-damping insulation as well. The two compressors run two recirculating helium systems: one at 350 psi run by a Sumitomo F-70 compressor, and one at 100 psi run by a Sumitomo HC4E compressor.

The high-pressure system provides the bulk of the cooling power by running a standard Gifford-McMahon (GM) coldhead. GM coldheads and their close cousins, pulse tube (PT) cryocoolers, are both established technologies for closed-cycle, low-temperature, low-vibration cryogenic applications. However, they are mechanically operated and thus vibrations on the order of a few to tens of microns are unavoidable [194, 195]. Additionally, they are quite large, PTs must be mounted in a specific orientation, and in general the entire chamber must be designed around the cryostat (see [137] for details of such a system).

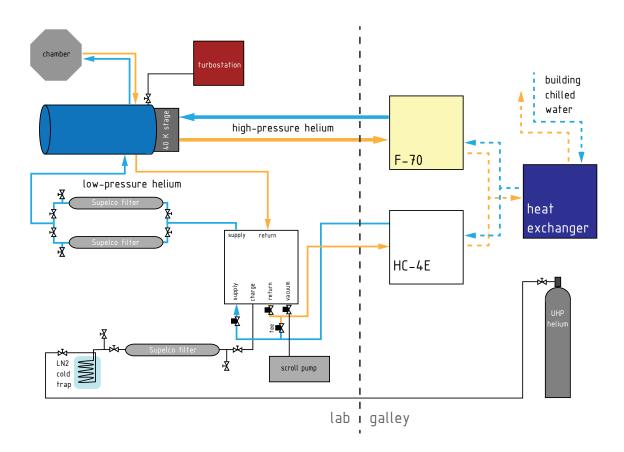
The innovation of the Stinger is to use a GM cryocooler to cool a secondary, lower-pressure helium system which cools the system of interest via a flexible bayonet. In this way, the vibrations of the GM coldhead are very well isolated from the system being cooled. Additionally, the bayonet can be installed using only a single port of an octagon just as our flow cryostat used, and so puts far fewer restraints on the geometry of the system. The Stinger provides additional cooling via Joule-Thomson expansion through a capillary at the end of the bayonet [196].

When the Stinger system is working, it keeps our trap between 8-10 K which gives us single-ion lifetimes on the order of a day and two-ion (same-species) lifetimes of hours. The flow cryostat operated closer to 12-15 K, with two-ion lifetimes of 30 minutes to an hour. We attribute the extended ion lifetime to the improved cryopumping, especially of hydrogen, at this lower temperature.

6.3 High-field coils

We operate with a quantization field of 212.8 G in order to have a magnetic field-insensitive 'clock' transition in magnesium. This field is produced by running approximately 60 A through a

Figure 6.4: Plumbing diagram for both helium systems as well as water cooling for a Stinger cryocooler. The F-70 compressor runs the Gifford-McMahon (GM) cryocooler using 350 psi of helium which cools the lower pressure (~ 100 psi) system circulating by the HC4E compressor through three stages of heat exchangers inside the 'blue can' vacuum housing. This lower pressure line passes through a Joule-Thomson capillary at the tip of the bayonet where the gas expansion provides additional cooling.



pair of foil-wound copper coils from Stangenes Industries, described in [163]. These coils are water cooled and the current is actively stabilized with an Oxford-style servo built by David Allcock and Shaun Burd and further finessed by Laurent Stephenson [197].

The stability of this magnetic field is one of the primary limitations on our experiment at the moment. The ion is so sensitive to this field that we can measure the duty cycle of the chiller that cools the coils using the ion.⁵ In the near future, we plan to move to a permanent magnet setup which has been designed, built, and tested by Christina Bowers.

6.4 Shim coils

In addition to the high field coils which create our quantization field, we have a pair of small shim coils along the two perpendicular directions. Because our state preparation depends on the purity of our circular polarization, which is to say its alignment with our quantization field, we can run small ($\lesssim 1$ A) currents through these coils to 'shim' up the quantization field to align with the laser beam. In practice, we only use these coils when very high fidelity state preparation is desired, as for the randomized benchmarking data reported in [188] (and a second set of experiments, publication of which is in preparation). These small coils are not water cooled or stabilized.

6.5 Imaging system

One of the key improvements we were able to enact when we broke vacuum was to change to a re-entrant top viewport that allowed us to use an achromatic Schwarzschild objective (identical to that used [185] and described in [198]). The old imaging system, described in [163] and shown in Fig. 6.5, was composed of off-the-shelf optics and required ≈ 5 mm of vertical translation to switch between focusing calcium and magnesium wavelengths. This made switching the species we were imaging inconvenient. The new objective was designed to be achromatic at beryllium wavelengths (235 nm and 313 nm) and also works well at 280 nm and 397 nm. However, it has a fairly short

 $^{^{5}}$ We can also detect metal drawers opening and closing and metal chairs being moved around the lab. Quantum sensing!

working distance (≈ 32 mm from the ion to the closest surface of the objective), which is why we needed to replace our top viewport to be able to use it. Replacing the top viewport also had the advantage of removing the four argon ion guns which took up a substantial amount of the space around our chamber (see Appendix A for more details).

After the objective, we have a pair of slits we use for blocking out stray light and a flipper mirror that controls whether the collected light goes to a photomultiplier tube (PMT) (Hamamatsu R7600P-203) or camera (Andor DV887ECS-UVB). We use the camera for laser alignment and debugging, but for day-to-day operation and data taking we use counts gathered from the PMT. We have moderate amounts of coma and spherical aberration in this imaging system, as can be seen in Fig. 6.5, but all data is taken as PMT counts, so we have not found it necessary to correct it.

We use a variety of different filters in the imaging system depending on whether we are using one or both species. For calcium-only operation, we use a Semrock FF01-395/11 line filter which lets through the 397 nm fluorescence but no other wavelengths. For magnesium-only and both-species operation we use a ThorLabs FGUV5 bandpass filter. This allows through the 280 and 285 nm magnesium wavelengths, but blocks roughly 80% of the 397 nm calcium fluorescence. However, given that the SNR of our calcium detection remains the same, we find this compromise acceptable for now.

We have a third filter which we use only for checking the calcium neutral fluorescence at 423 nm - an obsolete Edmunds UV pass, 520-640 nm block. While this does allow through both the calcium and magnesium wavelengths, it passes far too much room light to be useful for ion imaging.

In the future, it may be desirable to add a second photomultiplier tube and use a dichroic filter to split the 280 nm and 397 nm light, enabling truly simultaneous detection rather than the alternate pulsing detection currently used. Another upgrade would be to determine a set of filters that would allow only the 280 nm and 397 nm light to reach the camera for imaging of mixed species crystals. Currently the powerful 866 nm and 854 nm repump beams for calcium are above the blocking range of the filters we use. The PMT is not sensitive to these wavelengths and so the scatter does not cause any issues for data taking, but being able to see both species of ion on the

camera at the same time could be a helpful diagnostic for debugging mixed species loading (see Chapter 7).

6.6 Laser systems

6.6.1 Magnesium

While the eventual goal of this experiment is to have a mixed species experiment where we no longer need the more troublesome set of lasers (in this case the magnesium ones), we will always need the photoionization laser and find it useful to have a resonant laser for cooling and readout on magnesium-only experiments or verifying mixed species ones.

6.6.1.1 Photoionization (285 nm)

Previous work on this experiment used an MPB Raman fiber amplifier and a LAS GmbH WaveTrain doubler to generate the 285 nm light necessary for magnesium photoionization, as described in [163]. However, the power output and stability of that system had been trending downwards for years and in 2022 it was replaced with a VECSEL-based system. From a homemade 1140 nm VECSEL in the lab next door, we fiber over ≈ 500 mW which is doubled to ≈ 100 mW of 570 nm light using a waveguide doubler (NTT Electronics WH-0570-000-A-B-C). This light is then fed into a homemade BBO-based cavity doubler [199] which can produce up to 5 mW of 285 nm light. To minimize trap charging, we use only 300-500 μ W of this available light, but the power and stability offered by this system are a huge improvement over the old one. This light is focused to a $\sim 15 \ \mu$ m waist in the loading zone of the trap.

⁶As promised in [163]!

⁷Thankfully trying to put 285 nm light through a fiber takes care of most of this excess power!

Figure 6.5: On the left, our old imaging system, with Zemax layout as well as images of the trap and a calcium ion as they appeared on our camera. The old imaging system was not achromatic and thus required us to translate the whole system up and down to focus different wavelengths. On the right, our new imaging system gives us markedly clearer images of our trap surface and ions. The achromaticity allows us to see all of our wavelengths on our camera without adjusting the focus, which is helpful for mixed species imaging and detection.

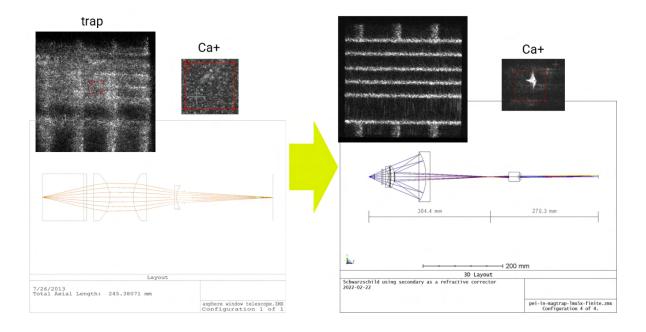
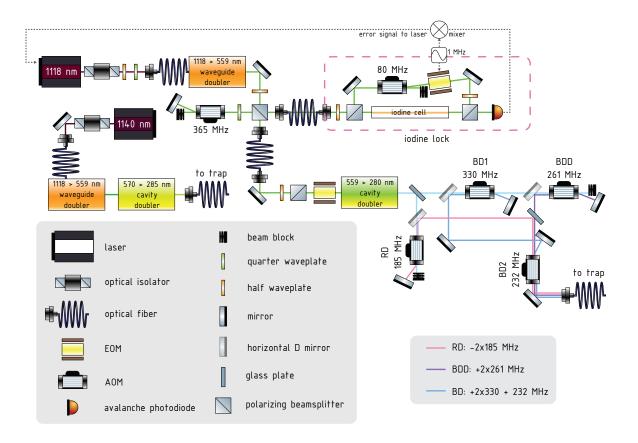


Figure 6.6: Our magnesium laser systems for photoionization, Doppler cooling, and detection. Additional mirrors used to provide sufficient degrees of freedom to align the AOMs and couple into the fibers are not shown in order to reduce diagram complexity. Both laser systems begin with an infrared VECSEL, are doubled into the green using an NTT waveguide doubler, and then doubled again into the UV with a bowtie cavity doubler. Our photoionization light is sent straight to the trap, while our cooling and detection light is further split into three beams: the blue Doppler (BD), blue Doppler detuned (BDD), and red Doppler (RD) for detection, cooling, and repumping. These names are explained in the main text. Figure made using the ComponentLibrary by Alexander Franzen.



6.6.1.2 Resonant (280 nm)

The VECSEL-based laser system resonant with the ${}^2S_{1/2} \leftrightarrow {}^2P_{3/2}$ transition at 280 nm remains the same as was described in [139] and [163].⁸ The 280 nm light provides Doppler cooling, repumping, and state detection for the magnesium ions and the nomenclature for the different beams is an artifact of the group's history. The blue Doppler (BD) beam does low-power, near-resonant Doppler cooling and resonant detection, the blue Doppler detuned (BDD) beam provides high-power, far-detuned Doppler cooling and optical pumping, and the red Doppler (RD) repumps population out of the upper manifold. More details of how these lasers are used for state preparation can be found in Chapter 4, so here we just explain their generation, as depicted in Fig. 6.6.

A homemade 1118 nm VECSEL⁹ emits ≈ 700 mW of infrared light which is fed into a waveguide doubler (NTT Electronics WH-0559-000-A-B-C) to produce ≈ 280 mW of 559 nm light. At this stage, a few mW picked off and used to lock the VECSEL to an iodine spectrometer. The rest of the light is fibered to a different portion of the optical table where it enters a homemade BBO-based doubling cavity [199] that converts roughly 130 mW of green light into 3 mW of 280 nm light. From here that light is split into three different paths which all eventually combine into the same fiber to the trap, shown in Fig. 6.6. The RD path picks off $\approx 10\%$ of the light, feeds it through a double pass AOM at -185 MHz, and then into the trap fiber. The BD/BDD lines are more intertwined: when the BD1 AOM is off, all of the light passes through the AOM unmodulated and enters the +261 MHz BDD double pass after which it combines with the RD, passes through the turned-off BD2 AOM, and on into the fiber towards the trap. When the BD1 AOM is activated, the light is double-passed at +330 MHz and then bounced through the turned-on BD2 single pass +232 MHz AOM and into the trap fiber. The BDD light is at full power only when BD1 and BD2

⁸Possibly the only part of the experiment I inherited that has not been removed or rebuilt, although I have had to realign it more than once.

⁹Our other homemade VECSELS are very robust creatures, but this one in particular is a bit of a diva. I include a few notes on her ideal operating conditions for future occupants of this lab, or owners of other fussy VECSELs. She is very sensitive to acoustic noise, and so she lives in a foamboard box, and her breadboard rests on a Sorbathane bed. There are water absorption lines near 1118 nm that can cause the laser to become unhappy (as evidenced by a weird or wobbly iodine spectrum) if her desiccant packets aren't changed every six months or so, and she is also constantly flushed with a gentle flow of dry nitrogen. Finally, she takes roughly three hours to completely thermally stabilize even though every component of her cavity is individually temperature controlled. But she's worth it.

AOMs are off, and the BD light only reaches the trap when the BD1 and BD2 AOMs are both on. ¹⁰ The BDD light is thus 370 MHz red of the BD, and the RD is 1.262 GHz red of the BD. We perform intensity stabilization on the BD by feeding back to the BD1 double-pass AOM from a photodiode that monitors picked off light just before the trap. This light used to be carried to the trap via a bare, home-connectorized, solarization-protected fiber, but that fiber died and has been replaced with a custom patch cord from AlphaNov. Transmission through this fiber is extremely position-dependent, we believe due to the microbending sensitivity observed in [200].

While we hope to not need this laser system at all eventually, we want to retain the ability to run magnesium-only gates and experiments, and it could be used to verify that our quantum logic operations with calcium are working as expected.

6.6.1.3 Raman (280 nm)

The VECSEL-based Raman laser system for ground-state cooling of magnesium as described in [139] has been almost completely disassembled and used for parts.¹¹ We have kept the VECSEL as a backup in case the resonant VECSEL ever fails.

6.6.2 Calcium

Our helper species, ⁴⁰Ca⁺, requires a lot of lasers. ¹² Fortunately, all of the required wavelengths are available as direct diodes from Toptica. Our 397, 423, 729, 854, and 866 nm lasers are all Toptica DLPros while our 375 nm laser is a Toptica iBeam Smart. ¹³ The 397, 854, and 866 nm lasers are all locked to temperature-stabilized, ULE cavities designed, built, and tested by David Allcock. Our 729 nm laser is stabilized via an offset lock to the ion clocks' 729 nm light, which is in turn stabilized to a temperature-controlled sphere cavity. [201] Our temperature-controlled cavities are not as stable as one might prefer for ideal experimental operation, so we also feed back on the cavity

¹⁰If this is confusing to read, imagine trying to align it.

¹¹Sorry Shaun.

¹²An unkind person may note that we added six new lasers (actually eight before we jettisoned the 393/850 shelving scheme) to our 'laser-free' gate experiment in order to get rid of two.

¹³Although we may eventually move to a Vexlum 729 nm laser to get more power.

piezos using a wavemeter lock to compensate any slow drifts. The 423 nm photoionization laser is directly locked to our wavemeter, while our 375 nm laser is allowed to run freely. The calcium lasers are located on the opposite side of the lab from the trap chamber due to space constraints, and the light is carried between tables via 20 m fibers (ThorLabs PM-S350-HP for the blue light and PM780-HP for the red light).

6.6.2.1 Resonant (397 nm)

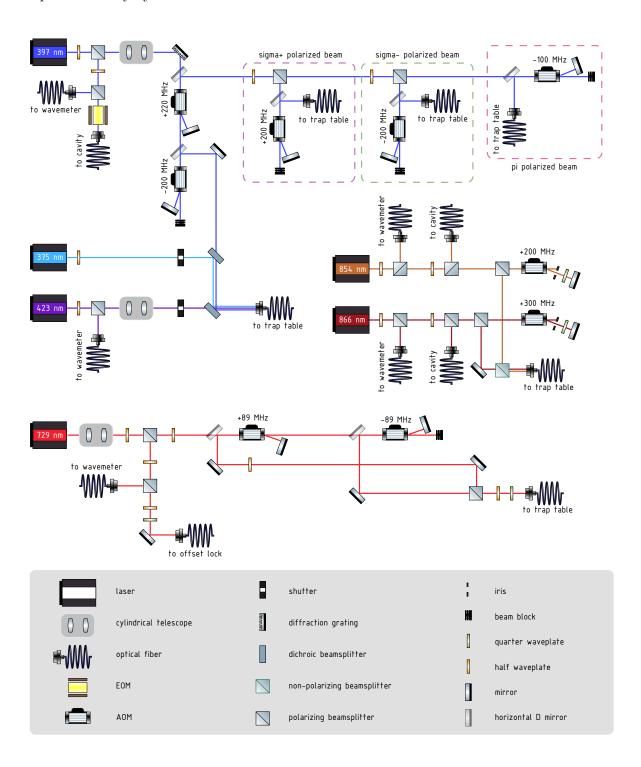
The 397 nm laser system is our most complicated, as it provides both near- and far-detuned Doppler cooling and fluorescence state detection for calcium, as well as the probe beam for electromagnetically-induced transparency (EIT) cooling. To drive our cycling transitions (see Fig. 5.5), we need σ^+ - and σ^- -polarized beams separated by nearly 800 MHz in frequency due to our 212.8 G magnetic field. We also need a weak π -polarized beam 200 MHz above the σ^- beam for EIT cooling, and finally a far-detuned, mixed-polarization beam 400 MHz red of our σ^- beam for far-detuned Doppler cooling. All of these frequencies are sourced from the same laser using a maze of AOMs and traveling to the trap via four separate fibers (see Fig. 6.7).

Once at the trap table, our σ -polarized beams are combined into a short photonic crystal fiber and then the mode is cleaned by focusing it through a pinhole. This beam enters the chamber through the north viewport, along the quantization field. The far-detuned beam arrives at the table in the same fiber as the photoionization lasers, and so enters the chamber from the northwest. The π beam comes in through the east viewport for maximum projection along the axial direction. See Fig. 6.8 for a full diagram of the beam orientations relative to the vacuum chamber and trap.

6.6.2.2 Photoionization (375 nm and 423 nm)

Our two photoionization wavelengths for calcium, 375 nm and 423 nm, are combined with the far-detuned 397 light on the calcium laser table (see Fig. 6.7), and then with the 866 nm and 854 nm lasers on the trap table. The 423 nm light is loosely locked (± 10 MHz) to the wavemeter, while the 375 nm laser is left free running. Both photoionization beams are controlled with SRS shutters

Figure 6.7: Schematic diagram of the generation and frequency modulation of our calcium lasers. All lasers shown are Toptica DLPros except for the 375 nm laser which is a Toptica iBeam Smart. Additional mirrors used to provide sufficient degrees of freedom to align the AOMs and couple into the fibers are not shown in order to reduce diagram complexity. Figure made using the ComponentLibrary by Alexander Franzen.



on the calcium table. The far-detuned light is switched using its AOM. We typically use around 200 uW of 423 nm light and 350 uW of 375 nm light at the trap.

Because we require that the 375, 397, and 423, and 866 nm beams all focus in approximately the same location for efficient ionization and loading, care must be taken with the optics in this beam path. The collimator we initially used (ThorLabs AC618APC-A) to launch the 375/397/423 beam suffered from rather serious achromaticity across that wavelength range, resulting in our beams focusing tens of microns apart, so we switched to a reflective collimator (ThorLabs RC04APC-F01) which performed substantially better. Additionally, we found that the Semrock FF510-Di02 dichroic initially used to combine the 854/866 and 375/397/423 beam paths produced stripes in the profile of the blue beams that we found to be undesirable. We replaced this with a custom dichroic from Layertech that does not do that. Finally, the achromatic doublet used to focus these beams is only achromatic near the center, so extreme care had to be taken to overlap the blue and red beams as close to the center of the lens as possible.

6.6.2.3 Repump (866 nm and 854 nm)

Our calcium manipulation scheme requires two infrared repump lasers as shown in Fig. 5.5. Both lasers are switched via (AA Opto-Electronics) AOMs and they travel to the trap in the same fiber, where they join with the photoionization light to enter the vacuum chamber via the northwest viewport.

6.6.2.4 Quadrupole (729 nm)

The 729 nm laser drives the ${}^4\mathrm{S}_{1/2} \leftrightarrow {}^3\mathrm{D}_{5/2}$ narrow-line quadrupole transition used for shelving our Zeeman qubit as well as resolved sideband and EIT cooling.[202] It is offset locked to the ion clocks' 729 nm laser, which is in turn locked to an ultra-high-finesse sphere cavity. The light is split into two tones, which we call 'upper' and 'lower', meant to shelve into the +5/2 and +3/2 sublevels respectively. These tones are combined into a single optical fiber on the calcium table, as shown in Fig. 6.7, and after the fiber the light is combined with the 397 π beam and enters the trap through

the eastern viewport, as shown in Fig. 6.9.

6.6.3 Beam pointing

As can be seen in Fig. 6.8, this experiment requires that we be able to control the position of ten different laser beams, coming in from up to five different directions, across the full extent of our 330 μ m trapping region. We also need to be able to individually align each wavelength without disturbing the others. This is performed with a combination of actuated lenses and piezo mirror mounts (ThorLabs Polaris K1S2P). In front of each viewport (except for the southeast, since the magnesium photoionization light does not need to move to different zones), the final focusing lens is mounted on a stage whose horizontal translation is controlled by an actuator (ThorLabs Z812B and Z825B, controlled via ThorLabs K-Cubes). We have found that these provide consistent, repeatable movement across a wide enough range to easily cover the whole trapping zone for each combined beam path. To individually control beam pointing, we use piezo mirrors placed before colors are combined in each path. While these do experience some hysteresis and have a much more limited range than the Tdcs, they are generally stable and have more than enough range to tweak the overlap of any single beam without having to disturb others in the same path, as can be seen in Fig. 6.9.

6.7 Trap

The work in this thesis was performed in two different versions of the same trap: before January 2022 the trap used is the same as in [139] and [163], but after April 2022 we had a new trap from a new fabrication run. We do not see significant differences between the two traps except that the new one does not have the damage on the side opposite the load zone that prevented transport to the far side of the trap (see Appendix A for more details), so in general I will not specify which version of the trap was used for any given experiment, though keep in mind that heating rates may not be identical with past work done using this apparatus.

In the new trap, the heating rates for a single calcium ion are 400 quanta/s on the 2.5 MHz

Figure 6.8: Overhead view showing the layout of the trap and its orientation in the octagon, with the entrance port of each laser beam, the direction of the quantization field, and trap axes all indicated. The numbers along the trap axis are the indices we use to measure position: 0 is the leftmost point at which we can trap, and 120 is the rightmost. Each index is separated by 2.75 μ m. The numbers in the corner show the number of our dc electrodes which run 1 to 18, and on the right our two rf and three microwave electrodes are also labeled. Figure modified from Laurent Stephenson.

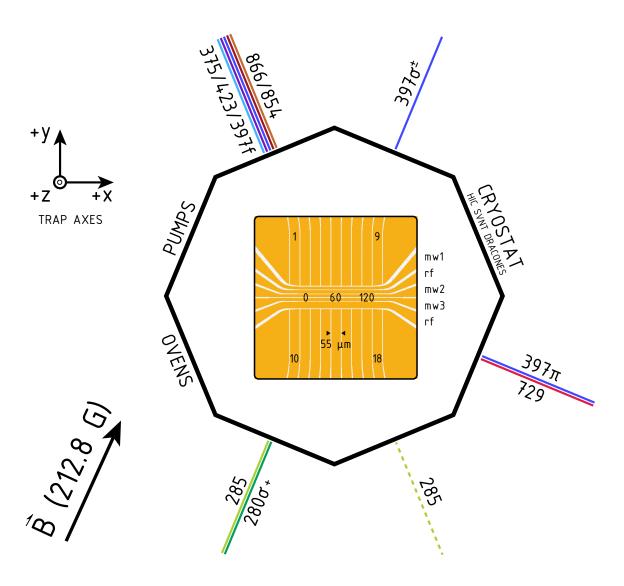


Figure 6.9: Schematic showing the optical layout on the trap table of all our different beam paths. Excess mirrors necessary to control all degrees of freedom have been omitted to help decrease diagram complexity. The orientation is the same as in Fig. 6.8. The dotted lines in the 285 nm beam path show the two different possible paths used to do Doppler-free vs. velocity-sensitive loading (see Chapter 7 for more details). As much as possible, each beam path has its own piezo mirror to allow for independent steering, and all focusing lenses are mounted on translatable stages to allow beam pointing across the whole trapping region, except for the 285 nm-only port which has a stationary lens.

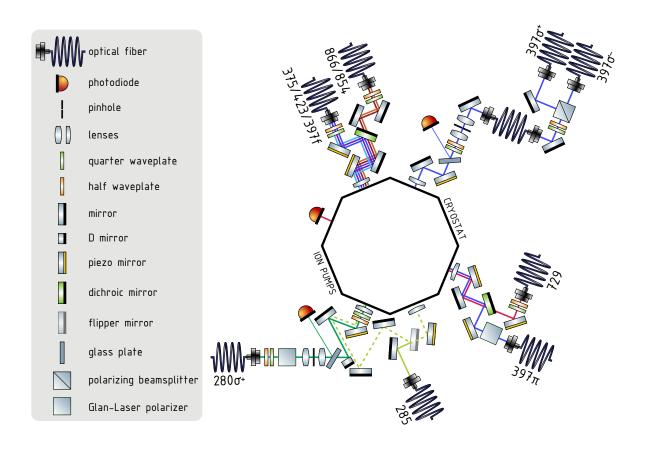
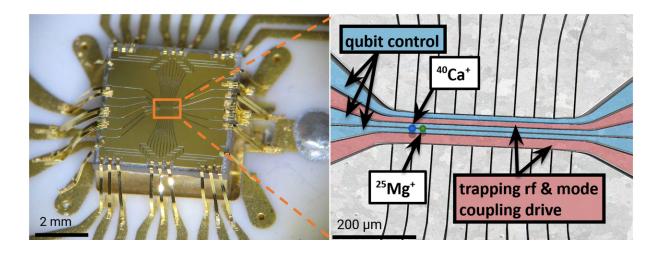
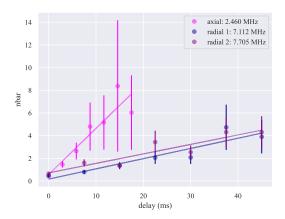


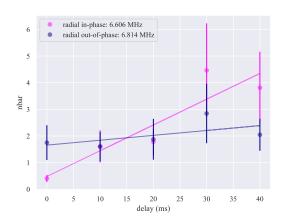
Figure 6.10: A photograph (left) and a scanning electron microscope image (right), which has been colored to emphasize the different electrodes. Ions are not to scale.



axial mode and 80-90 quanta/s on the 7.1 and 7.7 MHz radial modes. For two calcium ions, we measured 100 quanta/s on the radial in-phase mode and 18 quanta/s on the out-of-phase mode. This is worse than the old trap, which reported 60 quanta/s on a 2.4 MHz axial in-phase mode and 18 quanta/s on a 7.3 MHz radial in-phase mode for two magnesium ions. The heating rate on the 7.1 MHz axial out-of-phase mode was so low that a good measurement could not be made, estimated at < 2 quanta/s [163].

Figure 6.11: Heating rate measurements with a single calcium ion (left) and a calcium-calcium crystal (right).





The trap is made by electroplating 8 μ m of gold onto a 560 μ m thick sapphire substrate to create the pattern of electrodes shown in Fig. 6.8. We have 18 dc electrodes to provide axial confinement and shuttling. Two rf and three microwave electrodes run along the trap axis, providing radial confinement and delivering our microwave control currents respectively. We use three microwave electrodes to allow enough degrees of freedom to create the fields and field gradients that drive our gates, as discussed in detail in [163].

6.8 Trap rf generation

Trap rf generation is the same as in [139, 163] (Anapico APSIN3000), as is the passive stabilization.

One key improvement made for easier mixed species operation was to add a digital attenuator in place of the standard 12 dB attenuator previously used. Because our two ion species have different preferences for the strength of the trapping rf, this Artiq-integrated digital attenuator makes it easy to switch between different levels. This will allow us to eventually automate rapid changes in rf level, as other experiments in the group do during their mixed species loading where different rf levels are used during different portions of their merging.

6.8.1 Rf modulation for parametric micromotion compensation

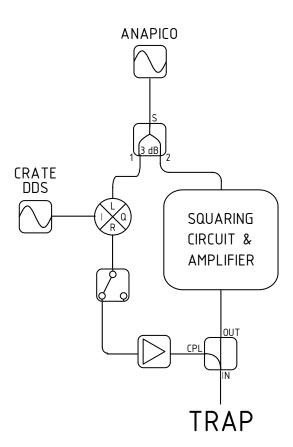
One innovation from Oxford that has been integrated into our experiment thanks to Laurent Stephenson is the parametric micromotion compensation.[187] Details of the actual micromotion calibration are in Chapter 4, so this section will limit itself to the hardware necessary to apply the modulation.

The modulation circuit was added as an alternate path to the squaring circuit used to stabilize the rf during experiments, as shown in Fig. 6.12. The modulation drive is generated by an Urukul DDS and then combined with the trap rf in a switchable way via a mixer (Minicircuits ZFM-3H-S+) and sent to the trap via a directional coupler (Minicircuits ZFBdc20-61HP-S+). Much care was taken to ensure that adding this modulation does not compromise the rf stability when the modulation is not on and we have seen no adverse effects. The modulation sidebands are 50 dB below the trap rf level, which is more than enough to see modulation signals.

6.9 dc voltage generation

The dc voltages which supply our axial confinement and allow us to transport ions back and forth along our trap axis were previously generated using PDQs [203]. While they often worked correctly, occasionally they would enter a failure mode wherein they needed to be manually restarted and physically disconnected from the trap 3-7 times to regain normal function. We chose to replace these with a different homemade voltage generation box, known as the Jonas (Keller) Digital-Analog Converters (JDACs). While they provide slightly worse resolution than the PDQs,

Figure 6.12: Schematic of the rf modulation circuit used to perform the parametric micromotion compensation. Modulation is generated by a DDS which can optionally be combined with the trap rf by means of a mixer and directional coupler. Testing and setup performed by Laurent Stephenson.



they are considerably more reliable.

The JDACs consist of two AD5370 evaluation board which each provide 40 16-bit DACs, plus one homemade LVDS/TTL buffer board per evaluation board that provides the SPI communication with the FPGA for programming. We only use 20 of these output channels to provide the voltages for our 18 dc electrodes plus a bias on the mesh and on the rf electrodes. The 18 dc electrodes provide axial confinement as well as transport and merging of different wells. The mesh provides a 'lid' on our trapping potentials which is very useful to counteract the occasionally enormous stray fields we generate using 280 nm light 30 μ m above our trap surface. The voltage bias on the rf provides one method for tilting the radial confinement of the trap, changing the splitting of our radial modes as well as their projection along our microwave fields. It also has some counterintuitive effects on well depth (see Chapter 7 for more details).

The 18 trap voltages go through RC filters both external to the vacuum chamber and again inside, which filter out technical noise which could cause motional heating. One upgrade that was made during the trap rebuild (see Appendix A) was to increase the number of additional input lines which bypass these external filters from one to three - we use these lines to apply 'tickle' voltages resonant with our motional modes or their difference frequencies, and bypassing one level of filtering means we do not have to drive as hard to see a response on the ion. In previous iterations of this trap, our only tickle line was connected to electrode 4 (for electrode naming conventions, see Fig. 6.8), just left of center on the top row of electrodes. This location provided decent projection of the tickle field along all three motional modes for ions on the left side of the trap as well as in the center. We added lines to electrodes 1 and 9 (at both ends of the top row of electrodes) to enable good projection along all three motional modes for ions at any position along the trap axis.

While developing our mixed species transport and merging operations, we found it necessary to apply larger shims than we have historically needed in order to transport over the large stray fields that occur at the boundaries of our load zones (see Chapter 7 for more details). To allow this, we installed a 20 channel amplified buffer box, also homemade (Analog Devices SSM2019)

amplifiers). 14 This increased our voltage throw on each dc electrode to +/-10 V (we did not amplify our mesh or rf bias voltages).

6.10 Microwaves

The existing microwave structure was designed to address the hyperfine transitions of magnesium shown in Fig. 5.1, between 1.3 and 2.2 GHz [163]. These microwaves were heavily filtered to prevent undesired signals from reaching the trap, so modification of the microwave system to incorporate calcium required a substantial overhaul. The calcium Zeeman qubit is 596 MHz at 212.8 G, so we needed the capability to drive both the carrier and sideband transitions in this range. The redesign and rebuild of the microwave system was done by Laurent Stephenson and Daniel Slichter.

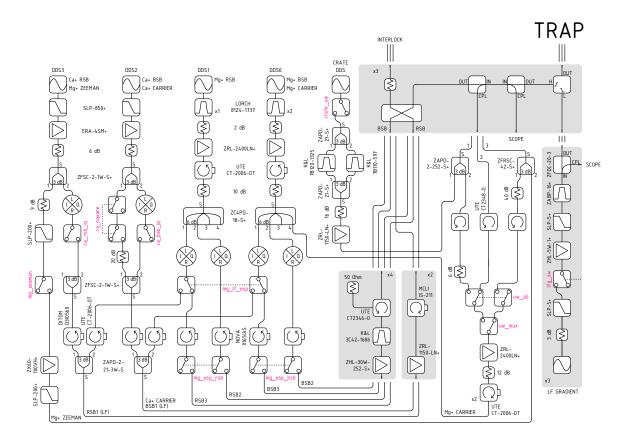
We did not add any additional AD9914 DDS to the system, but rather forced the existing DDS' to serve dual purposes. Rather than the carrier, BSB, RSB, and auxiliary DDS we had for magnesium-only operation, we now need carrier, BSB, and RSB DDS' for each species, plus we added the capability to directly drive the transitions between Zeeman sublevels. The BSB and carrier share a DDS for both species, and the calcium RSB also acts as the magnesium Zeeman DDS. The full schematic of signal generation, amplification, mixing, and switching is shown in Fig. 6.13. We also have one of the DDS' from our Artiq crate plugged directly into the trap to allow an alternate method for driving π and $\pi/2$ pulses during same-species gates to avoid the thorny issues involved in phase-tracking the pulses of one DDS acting as two.

6.11 Control system

Our experimental control is performed using the Advanced Real-Time Infrastructure for Quantum physics (Artiq).[204] We have two Artiq crates: the home-built one described in [163] has a new evaluation board, a Xilinx ZC706, but is otherwise the same; and we have added a new crate based on M-LAB's Sinara family of hardware featuring a Kasli FPGA with three Urukul DDS cards, a Sampler Adc module, and a TTL card. This combination adds 12 DDS channels, 8

¹⁴By Andrew Wilson!

Figure 6.13: Schematic showing our scheme for generating and combining the microwave and radiofrequency tones we use for state preparation and entanglement. Figure courtesy of Laurent Stephenson.



Adc channels, and 8 TTLs to the existing control system. The ZC706 crate primarily controls the magnesium portions of the experiment and all of the microwave equipment. The Kasli crate was added in 2021 and the two crates are connected using the distributed real-time in/out (DRTIO) framework available in version 7 of Artiq, with the Kasli serving as a satellite. The upgrade of our experiment control from Artiq 2 to Artiq 7 (and now 8), as well as the porting of (almost) all the experiment control code from the in-house 1DScan framework to the Oxford-developed NDScan framework [205] was performed by my computer angel, postdoc extraordinaire, Laurent Stephenson. While NDScan has a steep learning curve, it is very powerful and flexible.

The primary selling point is that it allows 'n-dimensional' scans in any parameter, while 1DScan experiments had to be built around scanning a single parameter. This streamlines the experiment writing process, and the fragment-based nature of NDScan enforces a modularity in the codebase that helps prevent issues such as slightly different state preparation and cooling methods being used in different experiments, something we often had in the old codebase depending on who wrote a given experiment and when.

7

Effects of stray potentials on mixed-species loading and transport

In Chapter 2, we introduced the concepts of loading and transport as essential operations for a QCCD-like trapped ion processor, and also discussed the benefits and drawbacks of using surface electrode traps. Loading and transport in surface traps can be challenging due to their shallow well depth, and the stray electric potentials which can develop when charges accumulate on the trap surface. Although calling them stray potentials is most accurate as the field curvatures have an impact on the ions as well as the fields, colloquially the term 'stray fields' is used, and we measure them at the ion position in units of volts per meter. Stray field curvatures are not typically measured. These stray electric fields can be compensated for, as discussed in Chapter 5, but in every trap there is a limit to the amount of voltage that can be applied to any given electrode, and thus a limit to the size of stray field that can be compensated. Uncompensated stray fields push ions out of the rf null, causing excess micromotion which both impairs laser cooling efficiency, and, for multi-ion crystals, can parametrically heat ions out of the trap [97]. Surface traps also experience anomalous heating which can excite ions out of their wells if sufficient laser cooling is not applied.

In the apparatus used in this thesis the standard issues of charging and heating are exacerbated by the ion's proximity to the surface as well as the short-wavelength lasers needed to cool and detect $^{25}\text{Mg}^+$. Compensation for these stray fields can be challenging is because the segmented control electrodes used for transport and stray field compensation are relatively large and distant compared to the ion-surface distance—more distant electrodes mean more voltage is needed to compensate a given field at the ion position.

Stray electric potentials are a consistent issue in surface electrode traps with small ion-to-surface distances due to the shallower trap depths and surfaces in closer proximity compared to blade-and rod-style 3D traps which typically have larger dimensions. It is commonly believed that these stray electric potentials are created by contamination or stray charges accumulating on the surface of the trap electrodes or in the gaps between them, though the exact mechanisms are not known and could be different for different trap materials and designs. Nominally each electrode is held at a static voltage by an external voltage source, but deviations from the expected trap potentials can be observed by noting the differences between simulated and observed ion positions and motional mode frequencies and are commonly attributed to 'patch potentials' on the trap surface. Most systematic studies of charging in surface traps have focused on laser-induced charging [121, 206–208], although in the trap used in this work that does not seem to be the primary mechanism.

This chapter has a slightly different structure than the others because systematic studies of loading and charging were never one of the primary scientific goals for this apparatus. The investigation and characterization of the stray fields in this trap and their effects were only undertaken when they were actively obstructing other scientific objectives. Because of this, this chapter is presented in roughly chronological order. We start with an overview of our trap's unique features and landmarks, then discuss calcium-only operation, for which stray fields were less problematic (though not absent). The reintroduction of magnesium into the system caused a variety of problems, the effects and characterization of which form the bulk of this chapter. Finally, we discuss some strategies for mitigating these stray electric potentials and their effects. It is likely that many of these effects are specific to our trap and ion species, but we hope that what we describe will motivate

further studies of non-laser based sources of charging in surface traps.

7.1 Trap geometry

The trap used in this experiment is a relatively small linear surface electrode trap. The central trapping zone is shown in Fig. 7.1. It was fabricated in-house at NIST by Daniel Slichter by electroplating an 8 μ m thick layer of gold onto a 560 μ m thick sapphire substrate. There are 18 control electrodes arranged in two rows on either side of the trap axis, although the trap axis is not centered between the rows by design. Each control electrode is 50 μ m wide and the gaps between all electrodes are 5 μ m. The axial rf and microwave electrodes have widths of 34.2, 14.6, or 16.8 μ m, as shown in Fig. 7.1.

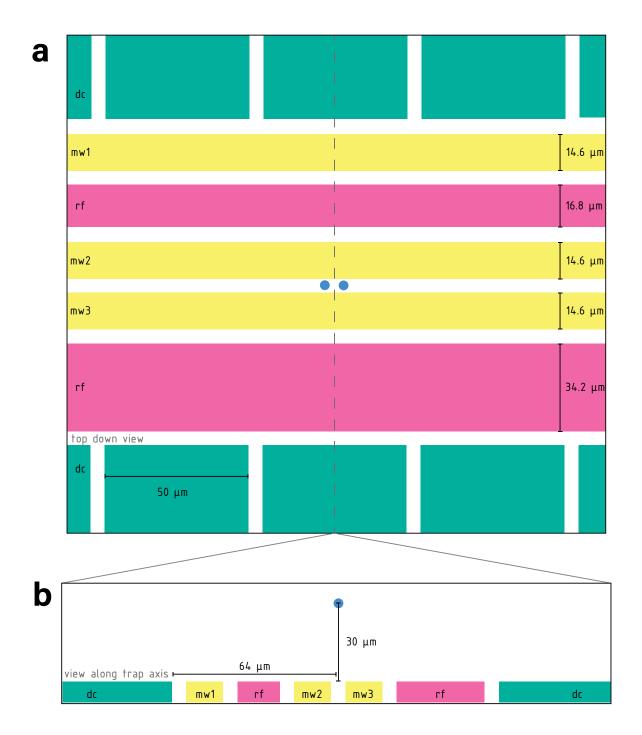
Though the full width of the trapping zone along the trap axis is 495 μ m, in order to create a simple well we need six electrodes: a high-low-high potential configuration on each side of the trap axis. Thus the farthest out we can trap corresponds to the center of the second electrode from the edge on both sides, and our most distant trapping zones are separated by only 330 μ m. In order to keep track of ion, beam, and stray field positions we artificially divide this usable trapping distance into 121 indices, each index being separated by 2.75 μ m. Index 0 is the far left side of the trap, index 60 is the center, and index 120 is the far right side. We use a right-handed coordinate system with the positive x axis pointing along the trap axis towards higher indices, as shown in Fig. 7.1.

7.2 Historical loading and charging in this trap

This trap has an ion-to-surface distance of 30 μ m and historically used magnesium, which requires photoionization light at 285 nm and cooling and resonant and light at 280 nm, as described in Chapter 5. We observe substantial charging of the trap surface that causes both static and dynamic stray electric fields, primarily associated with loading. The source of these stray potentials is not

¹A perfectly symmetric trap geometry would reduce the size of the optimization problem when generating voltage sets, and would make it simple to keep the ion position in the y-direction along the rf null, but we needed an asymmetric configuration for a couple of reasons. The rf electrodes were made different widths to tilt the radial modes off of perfect vertical and horizontal orientations, so the microwave gradients generated by the three microwave electrodes have projection along both modes. The three microwave electrodes are needed to give sufficient degrees of freedom to null the magnetic field while maximizing the magnetic field gradient needed to drive gates.

Figure 7.1: Due to the small ion-to-surface distance and the three axial microwave electrodes, the control electrodes are much farther from the ions than the ion-to-surface distance. This means that larger voltages must be applied to compensate a given field at the ion position. (a) Layout of the central trapping zone with dimensions labeled in μ m. All gaps are 5 μ m. (b) View along the trap axis. Electrode dimensions and the ion-to-surface distance are to scale but the ion itself is not.



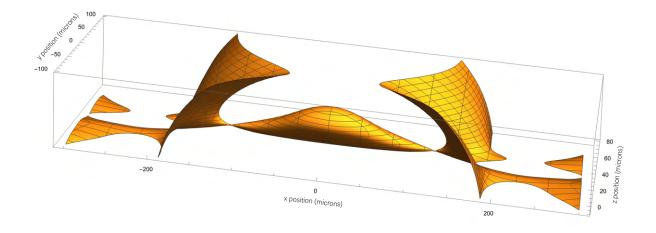
fully understood—proposed mechanisms include photoelectric charging of electrodes, contamination of electrode surfaces with other substances, and charge accumulation on exposed dielectric surfaces in the gaps between electrodes. The ion-to-surface distance in this trap is also among the lowest in any published ion trap, meaning the ions are closer to these surface charges. Before the new mesh was added in the rebuild (see Appendix A), which allows application of a uniform out-of-plane electric field across the whole trap, it was possible to charge the trap so strongly from a few failed loads that out-of-plane stray fields could no longer be compensated. Charging from loading magnesium creates a stray potential well with an axial frequency of ~ 1.5 MHz with all control electrodes grounded. The intentionally applied axial confinement typically used in this trap creates as well with 2.2-2.6 MHz axial frequency, meaning the stray potentials are on the same scale as the applied potentials.

A previous generation of this experiment, which also used magnesium ions in a surface trap with an ion-to-surface distance of 30 μ m designed for microwave and radiofrequency magnetic field gradient interactions, performed a careful characterization of the stray potentials observed in that trap, which is similar to the one used in this experiment but had only a single well [209], and found that they caused variations in the axial trap frequency on the order of hundreds of kHz. They also observed that the bulk of the charging occurred during the ion loading process.

The successor to that trap, the one used in this apparatus, was thus intended to have a loading zone on one end and an experiment zone in the center, to allow experiments to be conducted away from the stray charges induced in the loading zone. However, attempts to transport ions to the center zone for experiments were abandoned in late 2017. Transport required compensating the stray fields along the entire transport path, which was much more time-consuming than compensating the stray fields in just the loading zone. Additionally, even with this tedious compensation along the entire path, the transport eventually stopped working completely for reasons that are not understood. At the time it was believed to be due to a loss of Doppler cooling power from degradation in the fiber that delivered the light to trap (common in the homemade fibers we use for UV light [210]), but replacing that fiber did not restore the ability to transport ions [163]. The experimental results presented in [163] and [139] were all obtained using ions trapped in the load zone on the edge of

the trap. There are some disadvantages to this, namely edge effects on the microwave electrodes causing less uniform fields than in the center, but they were outweighed by the convenience of not calibrating transport operations.

Figure 7.2: The 30 meV equipotential surface for the 'bathtub' well designed for loading in our trap. This creates a trapping volume that is both very large and very deep (for our trap; 3D traps and even other surfaces traps have both larger and deeper wells), allowing us to directly load two calcium ions into the same well. However, even with this optimized well we could not directly load a mixed-species crystal.

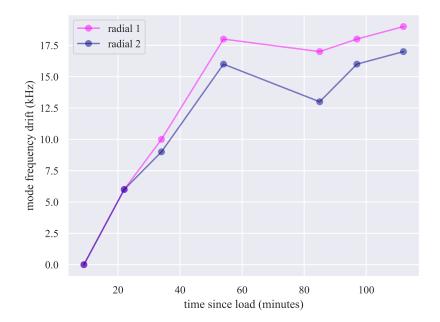


Another factor to consider when comparing to historical loading and charging is that we have observed that stray charges dissipate more slowly at cryogenic temperatures, and with the new cryostat (see Appendix B), we spend more time at lower temperatures than past generations of the experiment, potentially exacerbating stray field accumulation.

7.3 Calcium-only operation

For a period of more than a year, while we were working on consistently loading and controlling calcium at 212.8 G, we did not load magnesium. We observed much less stray charging during this time, which we attributed to the less energetic wavelengths needed for calcium versus magnesium. We did see slow drifts of our radial mode frequencies following loading, as shown in Fig. 7.3.

Figure 7.3: Following a calcium load, we see slow drifts in the radial mode frequencies on the timescale of an hour. The roughly equal drift in both radial modes and the timescale suggest the relaxation of some charging from loading.



7.3.1 Loading

As discussed in Chapter 2, ions with different charge-to-mass ratios require different parameters for stable trapping. We also found that our ion species prefer quite different parameters from each other for loading. Calcium loads most efficiently with larger mesh voltages (4-6 V, compared to 0 or 1 V for magnesium), as well as at a reduced axial scale (2.2 MHz axial frequency for loading versus 2.6 MHz normal operation). In physical terms, this creates a larger, shallower trapping volume along the trap axis with a strong electric field 'lid' from the mesh (1000-1500 V/m) holding the ions down. Once the ions are trapped, we increase the confinement so the ion is more tightly localized and compensate our stray fields, using the procedures described in Chapter 5, which typically requires reducing the mesh from the loading value. With magnesium we could directly load two ions into the same well, with relaxed radial confinement being helpful but not necessary [163]. However, repeated attempts to directly load two calcium ions into the same well were not successful until we started using a relatively large and deep (30 meV) bathtub well over the whole central area of the trap (see Fig. 7.2). Here, we note that 'deep' refers to the energy of particle that can be successfully trapped and 'large' refers to the spatial extent of the well.

Because of these difficulties directly loading multiple ions into a single well, and the different loading parameters for the different ion species, for mixed species loading we planned to load the two species into separate wells and then combine them, a process we call shuttle loading. The 'shuttling' refers to the step of moving the first ion out of the loading zone so it will not be kicked out by the second ion. Once we have loaded two ions into separate wells, we can combine them by merging the wells. This can be done in a slow and controlled manner, as opposed to direct loading which loads hot ions out of the thermal distribution produced by the ovens.

Shuttle loading requires single ion transport, which had not been done in this trap in six years, as well as multi-well trapping and transport potentials, which had never been done in this trap. It also requires merging the separate wells, which can be done asymmetrically (one ion remains stationary while the other is moved towards it) or symmetrically (both ions are moved towards

each other at equal rates), as shown in Fig. 7.7. This merge step is the most delicate part of the operation, as will be described below. Development and testing of this protocol was performed using calcium ions.

7.3.2 Transport

The first technique needed for shuttle loading was single ion transport, the history and basic principles of which were introduced in Chapter 2. While our trap has the capability to transport ions between zones due to its segmented electrodes, it had not been done since 2017. Laurent Stephenson developed our new transport waveforms² and we demonstrated transport with a single calcium ion in January 2023. This initial shuttling did not use any stray field compensation along the transport path, but smoothly ramped the voltages between the compensated the wells at both ends.

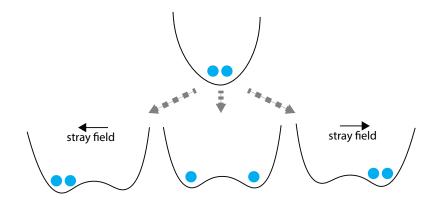
7.3.3 Shuttle loading

Our initial strategy for shuttle loading of two calcium ions was to load one ion in the center of our trap, shuttle it to one end, add a second well in the center, then recombine the two wells (the asymmetric merge pictured in Fig. 7.7). This required calibrating the axial stray field in the merge zone, as it can have a strong effect on merging ions, as illustrated in Fig. 7.4 and Fig. 7.5. After the initial calibration, we were able to reliably shuttle load pairs of calcium ions and did not need to recalibrate stray fields unless there was a substantial external perturbation to the trap. Our first shuttle load of two calcium ions occurred in February 2023, the week following our first successful transport. Similar to our initial single-well transport, we were still only compensating at the endpoints of the transport, on the basis that those were the only locations where the ions spent any substantial time. In April 2023 we added a couple of intermediate compensation points in hopes of making a smoother ramp and thus a more robust shuttle load. This calibration was

²And indeed all of the waveforms for the many different variations of transport and merging that will be described in this chapter. Thank you Loz!

not redone until mid-June, and that new calibration remained valid until we started trying to load magnesium in earnest in August.

Figure 7.4: Even very small stray fields along the trap axis can affect splitting (and the time reversed operation, merging). A difference of a few V/m can push both ions to one side or the other of the split.



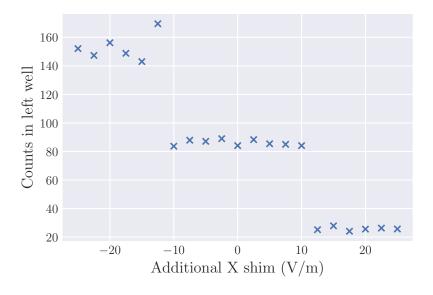
7.4 Mixed-species operation

We resumed attempts to work with magnesium in August of 2023. We planned to use the same transport and merging waveforms developed for calcium shuttle loading, which meant moving our magnesium loading to the center of the trap. This was also appealing as we thought we could calibrate the whole process using calcium ions, and then simply load a magnesium ion in place of the second calcium to obtain a mixed-species ion crystal.

7.4.1 Effects of magnesium loading on transport

We soon found that loading magnesium into the center of our trap created substantial stray axial confinement, on the order of 1.5 MHz axial frequency for a magnesium ion. The depth of the stray potential well, plotted in Fig. 7.6, was such that we could no longer transport the ions out of the center well at all. Transport out of the center is necessary for multi-ion loading as we needed the first ion out of the way in order to load the second. One way of understanding how strong these

Figure 7.5: We scan an additional x shim ('splitting tilt' added on top of the base stray field compensation) while repeatedly splitting and merging a pair of calcium ions. Both plots show the data from the same experiment, displayed in different ways. We set the detection beam in the left well, then split and detect in the split configuration. The left plot shows the number of counts detected in the left well, with the three distinct steps representing both ions on the left, one ion on the left, and both ions on the right. This raw data are included to show how sharp the step is each point represents 200 splits and merges, so it is very deterministic based on the shim value. On the right, we convert that raw data into whether or not the split was successful, i.e. one ion on each side of the well. We see that there is a 20 V/m window for successful splitting.



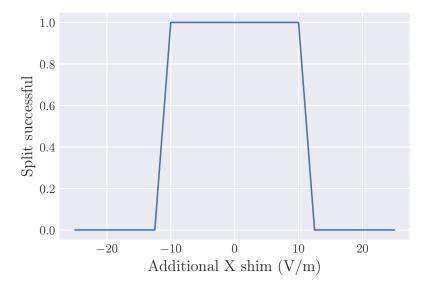
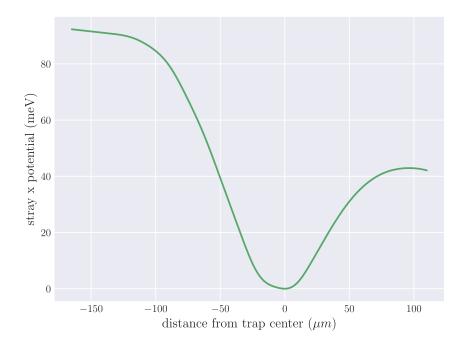


Figure 7.6: The stray potential in x as a result of loading magnesium in the center of the trap, calculated by integrating the stray fields measured along the trap axis. The axial stray field proved most problematic as it created a strong stray well that was initially impossible to transport ions out of because it was stronger than the applied transport wells.

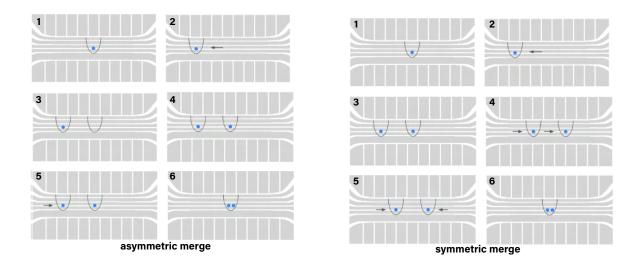


stray fields in the axial direction can be is that, when uncompensated, we observed that they could push ions up to 20 μ m away from the nominal well center.

In the previous generation of this experiment which used only magnesium, they performed a careful measurement of the stray potential. Warring et al. [209] found that their trap simulations were not accurate unless they added in a stray potential with electric field values and curvatures reverse engineered from the mismatch between expected and measured ion motional mode frequencies and rotation angles. Adding in these measurements allowed them to estimate motional frequencies and mode angles to within 10%, though they still saw day-to-day stray field variations on the order of a few tens of volts per meter. They proposed a few potential sources of this stray potential including laser-induced charging and magnesium contamination of trap surfaces. This trap only had a single well, so this characterization was only performed at one point. For successful transport and merging in our trap, we have found that we need to calibrate our stray fields coarsely along the entire transport path and finely within the merge zone (the region over the center electrode, $\pm 27.5 \ \mu m$ from the center well). Here, coarse compensation means within a few hundred V/m, and fine calibration means within \sim 50 - 100 V/m in the radial direction and \sim 20 - 40 V/m along the trap axis. The axial direction is more sensitive for reasons that will be discussed in detail below. The necessary calibration precision is difficult to quantify exactly because the symptom of miscalibrations is ion loss during transport or merging, and reloading shifts the fields. Additionally, the stray fields we observe are 3-5 times larger in magnitude than those observed in [209]. One possible reason for this discrepancy is that the trap in [209] operated at room temperature, whereas the current trap is used at ~ 10 K, slowing the dissipation of stray charges.

Eventually we were able to overcome the effects of these stray fields by installing amplifiers in our dc voltage generation chain to double the voltage available on each electrode to +/-10 V, allowing us to apply stronger shims. With transport capabilities restored, we hoped to be able to load and merge a mixed-species ion crystal.

Figure 7.7: We have tried merging our wells at several different locations along the trap axis. The two most successful strategies were the asymmetric merge at the trap center, pictured on the left, which worked very well for calcium-only operation. However, once magnesium had been reintroduced we found that the symmetric merge strategy was more successful as it provides the most independent control of both ions and the symmetry constraints can help reduce the size of the optimization performed when developing new wells. We also attempted asymmetric merges at 55 μ m out from the center (one electrode over) and 110 μ m out (two electrodes over) with much less success.



7.4.2 Effects of magnesium loading on merging

While doubling our voltages allowed us to compensate for the stray potential sufficiently to transport the ion out of it, merging two wells (and the mirror operation, splitting) is extremely sensitive to stray fields. In order to merge two wells into each other, the axial confinement necessarily needs to be relaxed, and any stray axial field can be thought of as a 'tilt' in the quadratic well. This tilt can cause both ions to slosh into one side or the other during a splitting operation, or to 'roll down' the potential with enough energy to knock the other ion out during a merge. The mixed-species surface trap in the group at ETH Zurich reported that stray fields as small as 5 V/m could ruin their splitting operation [35]. We experimentally determined that there is a range of roughly 20 V/m where we can successfully split two calcium ions, as seen in Fig. 7.5, with mixed-species crystals likely being even more sensitive due to the differential effects of stay fields on the different ion species.³

We hoped that the stray potential well from loading magnesium in the center would reach a steady state, but the high sensitivity of the merge and the volatility of the stray fields from magnesium loading proved to be incompatible. In November, we decided to move the magnesium loading zone out to the farthest possible edge of the trap. Due to the small trap size, however, this was only 165 μ m from the trap center. This did cause a substantial shift in our stray fields, as can be seen in Fig. 7.8 and Fig. 7.9. The effect of magnesium loading creating a large stray potential well was consistent: the stray well observed here was strong enough to impede transport, and to trap in with all control electrodes grounded. Surface charging can be dissipated by warming the trap up, as evidenced by the orange trace from February 2024, taken after we had performed a light on-table bake (3.5 hours at 375 K) and before any magnesium load attempts. However, after a few days of loading both ion species the stray potentials returned to approximately the same levels as before. As can also be seen in Fig. 7.8, there remains a smaller potential well in the center. It is not clear if this effect is from calcium loading at that position or if it is residual from having loaded

³This work was done later, using the symmetric merge strategy and not the asymmetric merge in use at this time, but I believe the result stands.

magnesium in the center. Even with magnesium at the far edge of the trapping region, the effects of a single magnesium load on the axial compensation needed to split are measurable.

We tried performing our merge operation one electrode right of the center (55 μ right of center), chosen to try and use the stray potential landscape to help rather than hinder the merge. At the time this decision was made, this position was a small peak in the axial potential landscape, which we hoped would work similarly to the narrow 'splitting' electrode some traps use to push ions apart as they split and slow them as they merged. However, the stray potential landscape in this trap turned out to be too volatile to rely on as a portion of the merge strategy. Eventually we decided to move the merge operation to the center of the trap, where we have the most control, and make it symmetric, as shown in Fig. 7.7. We also hoped that the extra symmetry available might help make waveform optimization easier.

We initially had some success with this strategy, but like all of our merge strategies it was only intermittent. Sometimes a careful calibration of the full trap would stay good for days, allowing many successful mixed species merges with no tedious recalibration required. Other times a calibration would work one time and then not again. We have been unable to correlate the charge landscapes which allow mixed-species merges with any particular variable. The ion losses nearly always occur over the center electrode. We believe that once the ions are both on top of the same electrode, we lose too much of the control necessary to compensate the stray fields at both ion positions if they are not relatively smooth. Even though the ions are separated by just 55 μ m at this point, we have measured spatial variation in shim fields across that distance in the hundreds of V/m. Radial stray fields also have different effects on ions of different masses, meaning mixed-species merges are even more sensitive than the single-species merges. We worry that the uncompensated excess micromotion creates heating during transport that causes the ions to leave the trap. We hope that the calcium beam renovations currently underway, which will allow for more cooling power, will be enough to allow merging even with some uncompensated stray fields.

Figure 7.8: The stray potential along the trap axis (x axis) as a result of loading magnesium on the left side of the trap, generated by integrating over the stray fields measured along the trap axis, recorded over a period of months. Pink stars mark the charge landscapes where we could successfully merge a mixed-species crystal. The raw shim field values in both x and z are shown in Fig. 7.9.

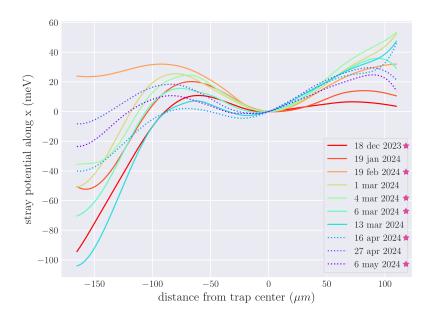
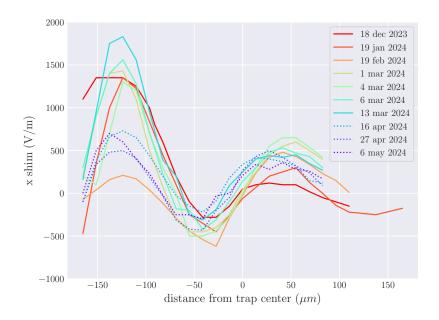
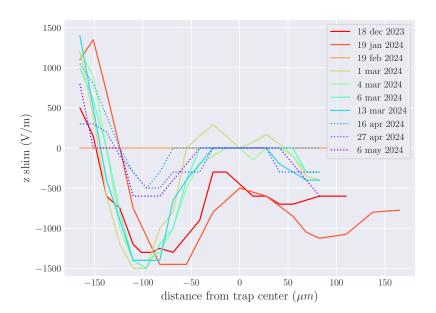


Figure 7.9: The stray fields along x and z as a result of loading magnesium on the left side of the trap. Stray fields along y are typically an order of magnitude smaller and thus not regularly measured. The general shape of the fields is fairly consistent but the overall size is quite large. The orange line (February) was taken right after a light on-table bake of the system performed in an attempt to reduce background gas pressure. This notably relaxed the stray potentials but they quickly returned to their previous levels after we resumed loading of both species. The dotted lines are from after we switched to velocity-selective magnesium loading. Velocity-selective magnesium loading reduced the magnitude of the stray potentials in the magnesium loading zone, but did not affect the stray potentials outside the magnesium loading zone. This is consistent with the hypothesis that the magnesium oven alone creates the axial stray fields that we observed affecting the calcium-calcium splitting operations in Fig. 7.10.





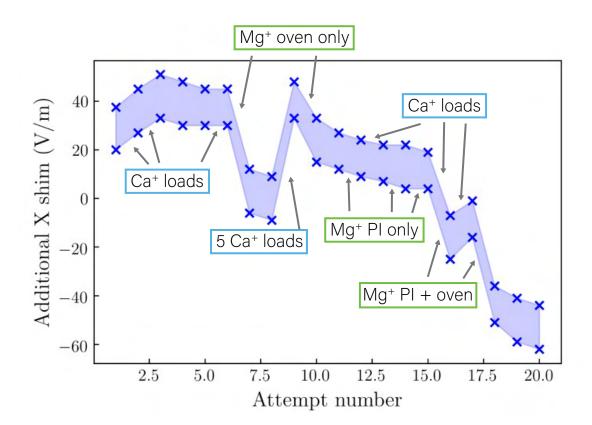
7.4.2.1 Quantifying the effects of magnesium loading on splitting and merging

To assess the effects of magnesium loading on the split/merge operation in the center of the well, we loaded a pair of calcium ions and performed several thousand split-merges while scanning the axial shim. During a period when we had performed no magnesium loads in the past week, we were able to split and merge at the same voltages thousands of times, with the successful split shim range being constant the whole time. Two minutes of magnesium loading was enough to disrupt this.

Figure 7.10 shows the effects of various operations on the edges of the 'plateau of successful splitting' between two calcium ions. At the beginning of the day, a pair of calcium ions was trapped in the center of the trap, and an experiment as is shown in Fig. 7.5 was performed to ensure that the splitting and merging was behaving and to measure the edges of the splitting shim plateau. This experiment has a fairly high probability of ion loss between submission of the experiment and execution of the first split-merge, which is why there are so many calcium load attempts interspersed throughout. We believe this to be due to the lack of idle cooling during the experiment's long compile time, something which is intrinsic to Artiq use, because there were no ion losses experienced during the actual split-merge experiments.

We had planned to systematically test the effect of magnesium loads on the calcium-calcium splitting in order to get an idea of the size of the effect and characterize the relaxation time we had anecdotally observed. However, there were several unexpected observations made during these measurements. It had been noted previously that magnesium load attempts tended to pull the edge of the splitting plateau towards lower axial shim fields. We expected that when the splitting plateau went towards higher axial shim fields, it was due to relaxation of those stray fields, but during these measurements we noticed that it seemed to be directly correlated to calcium loading. Of the first three calcium loads, the first two took nearly one minute to load an ion after the oven had warmed up, whereas the third one loaded in about 10 seconds. The first two loads had much more of an effect than the third one, possibly because there was time for many more untrapped ions to be

Figure 7.10: Measuring the effects of magnesium loading at the far left edge of the trap (three electrodes over) on calcium-calcium splitting in the center of the trap (165 μ m away). Each point on the x axis here corresponds to a scan as shown in Fig. 7.5 where the edges of the plateau have been extracted and plotted, demonstrating the extreme variability in the plateau of successful splitting.



created. Another piece of evidence in favor of the calcium loading causing charging effects is the incident where some technical difficulties required five attempts to reload a pair of calcium ions, leading to a dramatic shift in the splitting plateau towards higher axial shim fields.

The second unexpected effect was that of the magnesium oven alone. We wanted to test each part of the magnesium load separately (oven, photoionization, and cooling light) and measure relaxation times of the resulting stray fields. We chose to do the oven alone first, as we expected to see no effect and thus to not have to wait for it to relax. We were shocked to see that running the magnesium oven with no photoionization light created a substantial shift in the splitting plateau—an operation that was being performed 165 μ m away. Repeating this experiment replicated the result. Testing just the photoionization light (with the oven off) backed up this finding: the stray fields that were affecting the merge zone so drastically were not a result of light-induced charging, but something to do with the oven. It could be that the oven was hot enough to be producing ions, or electrons, or other charged particles. This effect was very surprising and demands further study. Interestingly, the calcium oven does not seem to create the same effect.

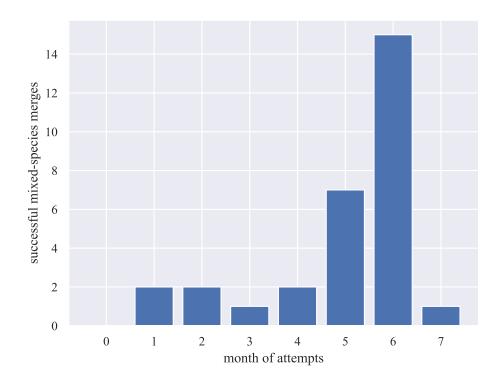
We do not have such rigorous quantification of the effects of magnesium loading in the actual load zone on the edge of teh trap, but we have observed changes in the z shim field of up to 500 V/m in a single day when beginning to load magnesium again after a period of calcium-only operation.

7.4.3 Merge-free loading attempts

Due to the difficulties described above, we have tried many different loading strategies that did not involve transport and merging of different ion species. Perhaps the most obvious of these is directly loading both species into the same well. We have tried this across multiple decades of rf attenuation as well as different axial confinements with no success. We have tried designing deeper small wells, and also directly loading into our large bathtub well. We have tried 'party loading' where we run both ovens and photoionization beams at the same time. The result of all of these has been to alternately load one ion species and then the other, with each new ion kicking the previous

⁴Name credit to Justin Niedermeyer.

Figure 7.11: We began trying to load mixed-species crystals in earnest in October 2023 (month 0 on this plot). Our first successful loads were in November but difficulties persisted and we rarely had more than a few days per month when we could successfully merge mixed-species crystals. The high number of successful loads in April comes from the week after we started doing velocity-selective magnesium loading.

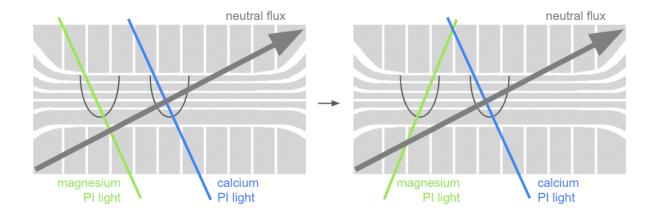


one out. We have also tried bringing mixed species crystals within the boundaries of the bathtub well in separate wells, then ramping into the bathtub well. This was also unsuccessful.

7.5 Strategies for managing and mitigating stray fields

We have come up with a variety of strategies to try to both manage and mitigate the stray fields in our trap. We will list a few that we have tried, and then a few more drastic measures that have not yet been attempted as of this writing.

Figure 7.12: Initially both species were loaded in a Doppler-free configuration, where the neutral flux and photoionization beams were nominally perpendicular (neglecting imperfect collimation of the atomic beam), as shown on the left. The ovens are mounted next to each other, so the neutral flux direction of both species is the same. We switched in April 2024 to a velocity-selective configuration, where the neutral magnesium atoms would now experience a Doppler shift that we could take advantage of to avoid ionizing the hotter parts of the thermal distribution. This was done by moving our magnesium photoionization beam so it entered the trap through a different viewport.



7.5.1 Velocity-selective loading of magnesium

One hypothesis for the source of our charging from loading is that the stray charges are the result of untrapped ions embedding onto the trap surface. If this is the case, then increasing our ratio of ions trapped to ions created should help reduce our charging per load. In order to test this hypothesis, we moved our magnesium photoionization beam from its Doppler-free configuration.

This was done by relocating the beam to enter along with the resonant 280 nm light, creating a 22.5° angle between the oven flux and the photoionization light, as shown in Fig. 7.12.

Because the light was previously resonant with the Doppler-free neutral flux, i.e. the atoms with approximately zero speed along the beam direction (neglecting imperfections in neutral beam collimation), the resonance is now shifted to be centered around the slowest neutral atoms. At our normal loading photoionization power of $\sim 300~\mu\text{W}$ of 285 nm light (roughly $100*I_{SAT}$ for that transition) we should have roughly 800 MHz of power broadening on top of the 79 MHz natural linewidth of that transition [186], with the possibility of increasing that power broadening by turning the photoionization light up.⁵ We estimate based on these parameters that this Doppler-selective loading most strongly ionizes those neutral atoms traveling at or below about 125 m/s.⁶

The effect of this change can be seen in the dotted lines in Fig. 7.8. This reduced the stray axial fields near the magnesium load zone by roughly 60%, despite not changing any other loading parameters. This seems to support the hypothesis that a considerable fraction of our charging in the load zone is not laser-induced but from untrapped charges embedding in the trap surface.⁷

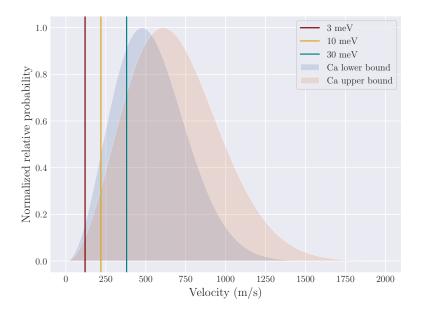
Following this change, we experienced a brief period where we successfully loaded and merged over a dozen mixed species crystals. However, after roughly a week and a half, the merging stopped working. A weekend warmup to room temperature did not relax our fields as much as we had hoped and allowed one final successful merge but even later attempts that same day were unsuccessful, and there have been no successful merges since then. We suspect that we reach a point where the radial stray field differences across the width of the electrode becomes too large for us to successfully compensate. These fields can be relaxed by warming up to room temperature for a few days, but they quickly return once we resume ion loading. They can be stabilized by avoiding loading

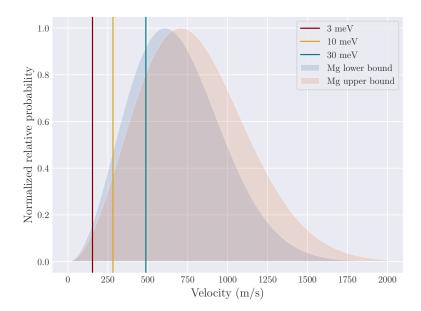
 $^{^{5}}$ Although we do think we have SOME amount of charging from the photoionization light, so probably diminishing returns here.

 $^{^6}$ This is ignoring the transverse Doppler shift which in the beryllium photoionization studies performed in Appendix E was found to be on the order of a GHz, but still a good approximation to first order.

⁷It must be noted that we would only expect these changes to affect the charging in the loading zone, as that is where the photoionization lasers create charging, either via laser-induced charging on the electrodes or untrapped charged particles from photoionization. We would not expect velocity-selective loading to affect the charging far outside the load zone which we observed from the magnesium oven, which is consistent with our measurements that this charging occurs regardless of whether the photoionization lasers are on.

Figure 7.13: The thermal velocity distributions produced by the ovens for both species, with energies corresponding to three well depths that are typical for our trap. For both species, even our deepest well leaves the majority of ions produced untrappable. Another factor to consider is that empirically, surface traps can only trap ions with energy on the order of a tenth of the actual well depth [93].





magnesium but that makes mixed species operation difficult.

7.5.2 Radical stray field mitigation strategies

A few more radical strategies for managing the stray charge have been suggested. Keeping the ion loading on a regular duty cycle (e.g. one load per hour per species) has helped to keep the fields stable. Another proposed strategy has been to try to saturate the charging on the trap surface. In theory, there is a limit to the amount of stray charge that can be accumulated on the surface. By rastering the magnesium photoionization light across the trap surface, perhaps while also running the oven, we could attempt to saturate the charging on the trap surface. A large concern with this strategy is that we would still have the rapid drifting of the stray fields, just at a much larger magnitude. There is also a small risk of depleting the magnesium oven unnecessarily.

7.6 Summary of stray field effects

The primary takeaway from this chapter is that when stray fields are of the same order of magnitude as applied fields, they must be taken into account for trap simulations to be useful. The stray fields in this trap are likely caused by a combination of factors, and exacerbated by the cryogenic operation and low ion-to-surface distance required by our science objectives. The low ion-to-surface distance places lasers (and untrapped ions) closer to the trap surface where they can impinge and cause problems, and also keeps trapped ions closer to these stray surface charges. Cryogenic operation slows dissipation of stray charges. We witness stray potential wells created where we load magnesium, which can be reduced using velocity-selective photoionization, but the magnesium oven alone also seems to create charging across a large portion of the trap, which is unaffected by this technique. While we hope to eventually stop using the 280 nm lasers required for cooling of magnesium ions, the photoionization light will still be required, as will the magnesium oven which seems to cause stray fields in ways that are not well understood.

⁸Although stability in a regime where we cannot merge is of debatable use.

8

Laser-free quantum logic

The development of quantum logic spectroscopy, where a secondary species can be used to prepare and read out a species that lacks a suitable transition for direct fluorescence readout via their shared motion [38, 84], enabled trapped ion clocks to rank among the best in the world [39]. However, the technique has proven useful far beyond the initial intended scope, as it has been invaluable to expanding quantum control techniques to more exotic particles such as molecules [40, 41] and highly charged ions [42]. Quantum logic readout has also enabled the highest detection fidelities ever recorded [85]. Quantum logic readout can ideally be a quantum nondemolition measurement, where the information being measured is not destroyed by the act of measurement (beyond projection), enabling arbitrarily high detection fidelities because the measurement can be repeated until a desired confidence level is reached [84, 85]. However, when implemented with Raman beams or on optical qubits, photon scattering or decay of metastable qubit states during this process is still a source of error. Performing quantum logic readout with laser-free sidebands is free of these restrictions and could achieve even higher detection fidelities in fewer rounds thanks to the absence of photon scattering and decay errors which can cause leakage out of the qubit manifold.

Traditional laser-based quantum logic techniques were covered in Chapter 4, so each quantum

logic section will contain just a brief review of the laser-based scheme to facilitate comparison to the laser-free scheme as implemented on our specific ion species. In general the techniques are very similar, just exchanging laser-based sidebands for laser-free, though there are a few technical details to be considered. We will discuss the potential advantages of such schemes, then conclude by presenting experimental demonstration of several preliminary steps towards laser-free quantum logic. This technique is essential to the goal of the all-electronic qubit, as it enables state preparation and detection with no lasers required for the qubit species.

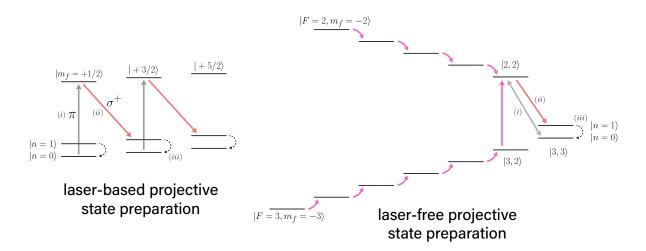
8.1 Projective state preparation

In laser-based quantum logic, the equivalent of optical pumping is performed using alternating carrier and sideband pulses with specific polarization (combined with cooling on the helper species) to ensure population moves irreversibly in one direction along the manifold. In magnesium, we can move population around within the hyperfine manifold using microwaves, progressively sweeping population out of each state and being sure none is left behind.

We begin by assuming that the population in the data ion is spread among all the states in the ground-state hyperfine manifold, and we want to sweep it all into $|3,3\rangle$, the bright state. In order to move population from $|2,2\rangle$ to $|3,3\rangle$, first the shared motion is cooled to the ground state and the helper ion is prepared in its electronic ground state. Next, a red sideband is driven on the $|2,2\rangle \leftrightarrow |3,3\rangle$ transition of the data ion. This will take population in $|2,2\rangle |n=0\rangle$ to $|3,3\rangle |n=1\rangle$ state, but leave population in $|3,3\rangle |n=0\rangle$ undisturbed. Cooling on the helper ion will then collapse the populations into the ground state of $|3,3\rangle$. We can proceed backwards through the manifold, moving each state into $|2,2\rangle$ one by one using the Zeeman microwaves (and a carrier pulse between $|3,2\rangle$ to $|2,2\rangle$ for population in the F=3 half of the manifold), then clear the population into $|3,3\rangle |n=0\rangle$ using these steps. The reason we need to be so careful about clearing out one state at a time is that the Zeeman microwaves that move population from the left to the right will also move population right to left if the higher states are not emptied, causing imperfect state preparation.

One potential imperfection is that the red sideband which takes $|2,2\rangle |n=0\rangle$ to $|3,3\rangle |n=1\rangle$

Figure 8.1: Schematic comparing the laser-based projective state preparation of $^{27}\text{Al}^+$ from [38] to our proposed laser-free procedure for $^{25}\text{Mg}^+$. While the exact steps shown here are species-specific, the general principles are widely applicable to different ion species. In both figures, grey and red arrows denote the carrier and red sidebands (laser-based on the left, microwave/rf on the right), while the dotted lines represent ground-state cooling on the helper ion. The pink arrows in the right-hand diagram are the Zeeman microwaves (see Chapter 5) which are used sequentially to move population to the right. For details of the exact energy level splittings of magnesium, see Chapter 5. The Zeeman 'microwaves' have frequencies of 75-135 MHz, while the carrier and sideband transitions between $|2,2\rangle$ and $|3,2\rangle$ or $|3,3\rangle$ are ~ 1.4 GHz.



will also take any leakage population in $|3,3\rangle | n = 1\rangle$ (for example, from imperfect ground state cooling or heating during the procedure) up to $|2,2\rangle | n = 0\rangle$. To mitigate this, we can perform an additional carrier π pulse between $|3,3\rangle \leftrightarrow |2,2\rangle$ before the red sideband so that the larger of the two populations will be in $|2,2\rangle$ to start. This is how we end up at the final process illustrated in Fig. 8.1. The numbered pulses in the figure show the process of getting population from $|2,2\rangle$ to $|3,3\rangle$, which is repeated every time, and the un-numbered pulses are run in sequence to successively put each state's population into $|2,2\rangle$ where it can be combined in the population in the $|3,3\rangle$ target state.

8.2 Projective state readout

Readout of the data ion via the helper ion is similar to the projective state preparation, although there are a few key differences. Quantum logic readout can be done in a way that is 'quantum nondemolition' (QND), where after the initial projection of the data ion into one of its qubit states, it is in an eigenstate of the measurement operator and therefore not disturbed by repeated rounds of measurement. This is important because it means that the overall measurement fidelity can be arbitrarily high even if the fidelity of a single measurement is not.

The first step of this protocol, just as for state preparation, is to cool the shared motion of both ions to its ground state. In the example of $^{27}\mathrm{Al}^+$, the upper qubit level is not involved in the sideband transitions that are used to transfer information into the motional state, but for our magnesium case that is not true, so we first want to use our carrier microwave pulses to shelve the population in the dark state to $|2,-1\rangle$. A carrier π pulse then swaps population from $|3,3\rangle$ into $|2,2\rangle$, and a red sideband drives population in $|2,2\rangle$ $|n=0\rangle$ to $|3,3\rangle$ $|n=1\rangle$, mapping the bright state population into the motional state of the ions. Next a red sideband on the calcium will take population in $|-1/2\rangle$ $|n=1\rangle$ to $|+1/2\rangle$ $|n=0\rangle$ while leaving population in $|-1/2\rangle$ $|n=0\rangle$ undisturbed. The population in the bright state of magnesium has thus been mapped onto the state of the calcium ion. The population that started in the $|3,3\rangle$ $|n=0\rangle$ state of magnesium was

¹This is also the first step of fluorescence-based magnesium detection, but for different reasons.

returned there and so the measurement process can be repeated until the desired detection fidelity is reached.

8.2.1 Zeeman microwaves

In Chapter 5 we introduced the Zeeman microwaves,² which range in frequency from 75 MHz to 135 MHz and are used to move population around within either the F=2 or F=3 manifold for magnesium.

We had hoped to be able to clear out the F=3 and F=2 manifolds simultaneously using ROTTEN composite pulses [211] but we found that π pulse times for the near-degenerate transitions were not close enough, so we have to do them sequentially. Still, the π times for these transitions are on the order of a few microseconds, so the procedure can be performed fairly quickly. Spectroscopy of the in-manifold transitions demonstrating our ability to prepare and drive population anywhere within the ground-state hyperfine manifold of magnesium is shown in Fig. 8.2 and Fig. 8.3.

8.3 Advantages of laser-free quantum logic

The highest detection fidelities yet measured in trapped ions were achieved using quantum logic-based projective readout [85]. A dominant error in trapped ion quantum logic readout (and in laser-based trapped ion computation in general) is photon scattering: stray resonant photons emitted by an ion involved in an operation can be absorbed by spectator ions, destroying their quantum information. Scattering events can also cause leakage out of the qubit manifold in hyperfine qubits. This leakage error can be mitigated by dividing the hyperfine manifold into orthogonal subspaces and fixing laser parameters to predominantly limit scattering of each state within its given manifold. This technique enabled Erickson et al. to reach the record detection fidelities reported in [85]. However, the main sources of imperfection can be avoided if operations on the data ion do not scatter photons, as is true for quantum logic readout based on microwave and rf sidebands.

 $^{^{2}}$ More accurately called the Zeeman rf, as the traditionally accepted lower cutoff for microwaves is 300 MHz. Please forgive this lapse in nomenclature.

Figure 8.2: Spectroscopy of the six Zeeman transitions in the F=3 half of the ground-state hyperfine manifold of $^{25}{\rm Mg}^+$. Each transition is driven with a 10 $\mu{\rm s}$ pulse, with the amplitude of the drive adjusted to perform a π pulse on resonance.

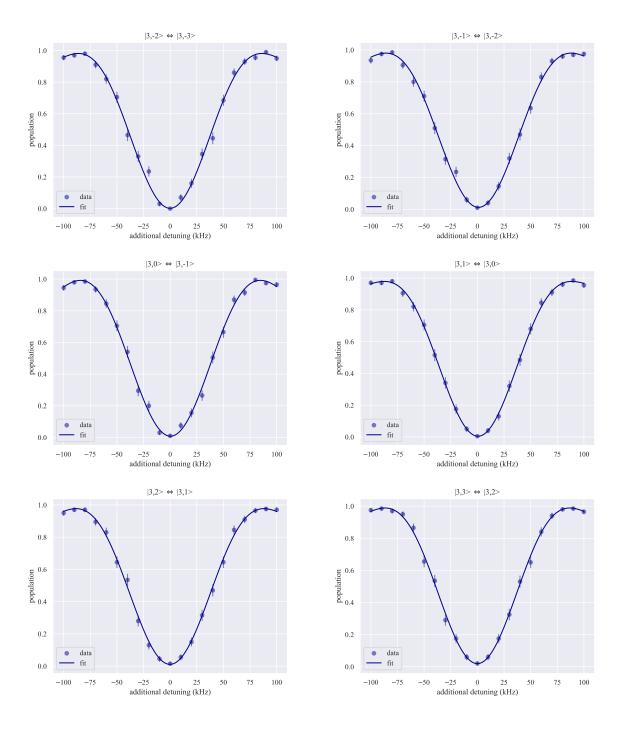
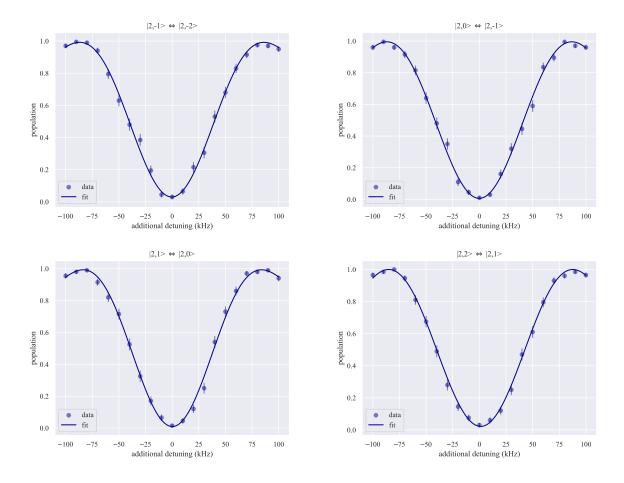


Figure 8.3: Spectroscopy of the four Zeeman transitions in the F=2 half of ground-state hyperfine manifold of $^{25}{\rm Mg}^+$. Each transition is driven with a 10 microsecond pulse, with the amplitude of the drive adjusted to perform a π pulse on resonance.



This lack of decay and thus scattering offers another advantage as well. We stated earlier that the detection efficiency can be improved to the desired threshold by increasing the number of rounds of detection performed. The number of rounds that can be repeated, and thus the detection fidelity, can be limited by the lifetime of the upper state (for example in $^{27}\text{Al}^+$, the decay of the $^{3}P_{0}$ state) [84]. Hyperfine qubits have the advantage that the shelf effectively does not decay back into the bright state and thus the dark state lifetime will not be the limit for the detection fidelity.

8.4 Mixed-species laser-free entanglement for quantum logic readout

As promising as this scheme for laser-free quantum logic might seem, there are still important limitations to be faced. The readout scheme as described above relies on perfect ground state cooling of the bus mode, with errors in the ground state cooling translating into population errors. Just as the MS gate is more robust to warmer ions because it does not rely on perfect single sideband pulses like the CZ gate, quantum logic readout can be performed using a mixed-species entangling gate and thus gain insensitivity to errors in ground-state cooling [181]. Laser-free mixed-species entangling gates could enable high-fidelity readout that is robust to ion temperature and photon scattering errors. This could become an important strategy to increase the robustness of quantum logic to experimental imperfections.

The remainder of this chapter is dedicated to the preliminary steps we have demonstrated towards this goal, and the next (and penultimate) chapter of this thesis presents the theoretical proposal for such a mixed-species laser-free entangling gate using a mixed-species extension of the near-motional frequency oscillating magnetic field gradient gate from [8, 149], together with parametric mode-mode coupling to help overcome the large differences in radial mode participation in mixed-species crystals.

8.5 Preliminary steps

In order to eventually perform mixed-species quantum logic (without an entangling gate) we need to demonstrate: ground-state cooling of a mixed-species crystal, laser-free sideband pulses on

both ion species, and potentially mode-mode coupling (which could be used to speed up mixed-species operations which are hampered by unbalanced mode participations). For the mixed-species entangling gate, we need to demonstrate single-species entangling gates in addition to parametric mode-mode coupling. Explanation of these preliminary steps and demonstrations follow.

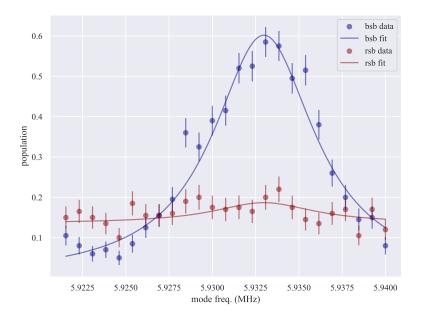
8.5.1 Sympathetic cooling

An important first step towards mixed-species quantum logic is the ability to ground-state cool the modes of interest. Ultimately we wish to stop using our magnesium Doppler cooling lasers and rely solely on sympathetic ground-state cooling using the calcium lasers. Performing broadband cooling with EIT first could help reduce the total cooling duration as well. At the time of writing, we have successfully demonstrated ground-state cooling of a calcium-dominant radial mode in a mixed-species crystal using first Doppler cooling and then resolved sideband cooling with the 729 nm quadrupole laser.

Due to the imbalanced mode participations, only calcium-dominant modes can be efficiently cooled with the calcium lasers, but there are ways to speed up cooling of magnesium-dominant modes by transferring heat into calcium-dominant modes where it can be cooled away.

8.5.2 Parametric mode-mode coupling

One strategy for mitigating the uneven participation in mixed-species radial modes is parametric mode-mode coupling. Though the technique of parametrically exciting specific motional modes in large clouds of ions is not new [212], the application of this technique to systems of just a few ions in surface electrode traps is not yet widespread [28, 30, 213, 214]. Both Gorman et al. and Hou et al. use this technique for cooling: by parametrically coupling two modes, heat can be transferred out of one that may be difficult to cool due to poor participation by the coolant ion into a well-cooled mode, cooling both modes efficiently [28, 30]. It has also been used to demonstrate coherent mode-mode coupling [213, 214], but has not yet been used in the context of a two-qubit gate. This technique can be used to compensate in any arena where uneven mode participation



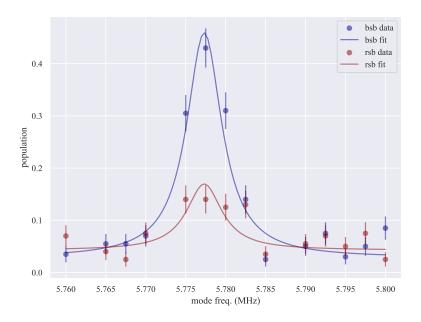


Figure 8.4: Sideband ratios measured with the 729 nm quadrupole laser on the calcium ion of a calcium-magnesium crystal, demonstrating near-ground state cooling of a calcium-dominant radial mode. The upper plot is from the first time we achieved this, in December 2023, with an estimated $\bar{n} \sim 0.45$ and the lower plot is from the second time, in April 2024, with an estimated $\bar{n} \sim 0.6$. These likely overestimate the final temperature due to the technical background from thermal effects in the AOMs use to control the 729 nm cooling and shelving pulses.

is an issue; we will focus on laser-free quantum logic readout and laser-free entangling gates with mixed-species crystals. That first application will be addressed in this chapter, and the second is the focus of the next chapter. First, we will examine the general theory, then we will look at the experimental implementation and preliminary results on calcium ion crystals.

8.5.2.1 Theory

In order to couple two motional modes, we need an electric potential with curvature along both modes to be coupled at the position of the ion [28], i.e. a nonzero $\alpha_{ij,n}$ where

$$\alpha_{ij,n} = \frac{\partial^2 U}{\partial x_i \partial x_j}|_{r=r_0,n}.$$
(8.1)

Here, U is the potential, x_i and x_j are the unit vectors of modes i and j, r_0 is the ion position, and n is the ion index. This curvature then must be modulated at the difference frequency of the modes, i.e. $\omega_c = |\omega_i - \omega_j|$.

Following [28], this creates a coupling Hamiltonian for a parametric drive that is on resonance, after making a rotating wave approximation that assumes the drive frequency is much less than the mode frequencies we get

$$\hat{H}_c = \hbar g_{ij} (\hat{a}\hat{b}^\dagger + \hat{a}^\dagger \hat{b}) \tag{8.2}$$

where the coupling constant g_{ij} between modes i and j, for n ions, from [30], is

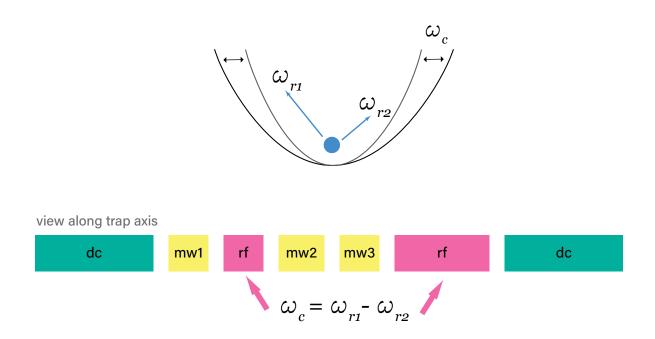
$$g_{ij} = \sum_{n} \frac{Q_n \alpha_{ij,n}}{4M_n \sqrt{\omega_i \omega_j}} \xi_{i,n} \xi_{j,n}.$$
 (8.3)

Here $\xi_{i,n}$ is the normalized participation of ion n in mode i, M_n and Q_n are the charge and mass of ion n, and ω_i is the motional mode frequency. Higher frequency modes will have slower coupling, and small mode participations reduce the coupling rate.

8.5.2.2 Experimental implementation

The choice of modes to be coupled determines the complexity of the experimental implementation. In a typical surface electrode trap with linear rf electrodes along the trap axis and segmented dc electrodes, applying a potential modulation to the rf electrodes will create a curvature modulation that can couple some of the radial modes to each other, whereas choosing an off-axis dc electrode will create a curvature modulation that can couple an axial to a radial mode. These are the two configurations used in [28] and this work, but of course more complicated curvatures may be needed, especially for longer ion chains. In [30], a cubic axial curvature is used to couple modes of a three-ion crystal with odd and even mirror symmetries. This is done by applying a phase-synchronized oscillating voltage with pre-calculated amplitudes onto 12 dc electrodes to generate a curvature that has suitable spatial dependence. To couple the radial modes of one or two ions, as we demonstrate below, quadratic potentials are sufficient and straightforward to generate.

Figure 8.5: We can couple two motional modes by modulating a potential curvature with projection on both modes at the difference frequency of the modes.



Another difficulty encountered when applying these modulation voltages is that ion traps are,

in general, specifically designed to filter oscillating electrical signals (other than the trap rf) from reaching the trap electrodes because such signals can drive motional heating. The supply lines for dc electrodes typically have multiple stages of low-pass filters to reduce any electrical noise near the motional frequencies from reaching the trap electrodes, and the trap rf passes through a resonator that strongly attenuates frequencies that are not resonant. This filtering has limited the coupling achievable in previous demonstrations to a few 10's of kHz [28, 30].

Past work in this apparatus using parametric drives for squeezing [215–217] made use of a resonator circuit to supply the parametric drive with enough power to reach the ions even through the rf resonator. While this remains an option if we desire stronger coupling in the future, the current coupling drive is applied directly onto the rf electrodes. We also have the option of applying the parametric drive to one of the three dc electrodes which have an external 'tickle' connection that bypasses one level of external filtering. This could be used if we wanted to couple an axial to a radial mode, for instance.

In order to determine which electrodes to use to generate optimal coupling we turned to the trap simulations used to calculate our shim voltages. These simulations take a 2D model of our trap and calculate the potentials and curvatures at the trap center generated by a unit potential on each electrode. The simulations tell us that the desired yz curvature needed to couple two radial modes for an ion at the center of the trap can be created by any of the rf or dc electrodes. However, there are some technical aspects that need to be considered: the applied voltage is going to be oscillating in a frequency range close to our motional modes, and we want to make sure we are not creating a strong oscillating electric field that could off-resonantly drive ion motion. Although the dc electrodes could nominally generate a stronger curvature per unit potential, the filtering in the frequency range we need to address is very strong. Gorman et al. applied voltage to opposing electrodes to cancel out the electric field [28], but the rf electrodes have already been designed to have an electric field null at the ion position.

The results of the trap simulations predict a yz curvature at the center of the trap of -0.159 V/m² per volt applied on the rf electrodes. While the exact coupling constant depends on the mode

frequencies and participations, for typical values in our trap we can estimate a coupling rate of ~ 3 MHz/V for a calcium-calcium crystal and ~ 200 kHz/V for a calcium-magnesium crystal. The large discrepancy is due to the small participations for mixed-species modes, and the precise values of the curvature seen by each mode will depend on the mode rotation, but this provides a basis to work from.

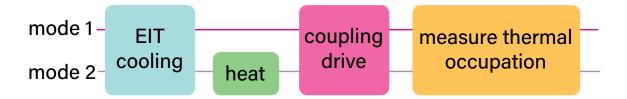
8.5.2.3 Results

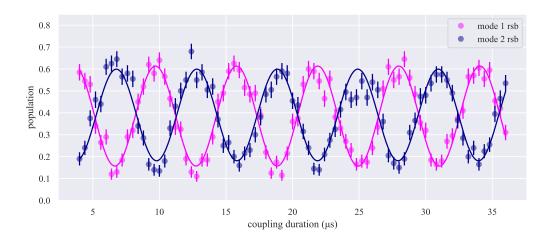
We performed demonstrations of parametric mode-mode coupling between the two radial modes of a single calcium ion and between the in-phase and out-of-phase radial modes of a calcium-calcium crystal. We generate the parametric coupling drive using one channel of a PDQ, the same type of fast DAC as is used to generate our 5 MHz gradient tones. The drive is connected to the rf electrodes on the same port as our rf bias voltage. At full power, the drive is 2 V peak-to-peak at the input to the trap. The data presented below was taken with amplitudes of 0.4 to 0.6 V peak-to-peak because the coupling was too fast for our pulse shaping at full amplitude. We shape the rising and falling edges of the coupling drive pulses to reduce off-resonant motional excitation, and as we cannot create ramps shorter than 1 μ s, we reduce the drive strength such that the swap period is at least $\sim 2 \mu$ s.

We measured the coupling rates between the two radial modes of a single calcium, and between two radial modes of a calcium-calcium crystal using the same experimental protocol, shown in Fig. 8.6. First, both modes are cooled near the ground state using EIT cooling. Next, we apply a few blue sideband pulses on one mode to provide motion for the modes to exchange. The data for a single ion are shown in Fig. 8.6 and for a two-ion crystal in Fig. 8.7. For the two-ion exchange data, we pushed the coupling duration out past 3 ms in order to demonstrate that the interaction remains coherent over ms, which will be necessary for coupling modes during a gate operation.

The peak-to-peak voltage numbers above are for the value applied to the input, but the actual voltage seen by the ion is much lower due to attenuation by the rf resonator. We observed coupling rates in a calcium-calcium crystal of 165 kHz with an applied voltage of 400 mV peak-to-peak.

Figure 8.6: Top: pulse sequence for the parametric mode-mode coupling experiments described in this section. Both modes are cooled near to the ground state with broadband EIT cooling, then a variable number of blue sideband pulses (usually 1-3) is applied to one mode to provide motion for the modes to exchange. The parametric coupling drive is then applied for a variable duration and the thermal occupation of both modes is measured. Bottom: coherent exchange of motion between the two radial modes of a single calcium ion. This experiment used a drive strength of 0.4 V_{pp} at the input, and achieved a swap duration of $\sim 3~\mu s$, or a coupling rate of 165 kHz.





From this, we can estimate that the actual voltage reaching the trap is 55 mV, for an attenuation of roughly 17 dB through the rf resonator. The expected maximum coupling rate for a calcium-calcium crystal is thus ~ 800 kHz, and for a calcium-magnesium crystal we would expect a maximum coupling rate of ~ 55 kHz. We can amplify the drive if we desire stronger coupling rates for mixed-species gates.

8.5.3 Mode-mode coupling for quantum logic

At first the application of this mode-mode coupling technique for quantum logic operations may not be obvious. We need to drive laser-free sidebands on the *same* mode of the shared motion to effect these quantum logic operations, in spite of the uneven participations in the radial modes we are constrained to use. Once we also consider the relative slowness of laser-free sidebands, driving sidebands of a single motional mode on both ion species could be prohibitively slow, especially given the limited coherence time of the field-sensitive qubits in calcium and magnesium. Our proposed workaround is to use mode-mode coupling to swap the quantum of motional excitation from one species' dominant mode into the other species' dominant mode between the sideband pulses on those modes, as shown in Fig. 8.9. This way each sideband can be driven on a mode that the ion participates strongly in. We believe that the quality and speed of our swap pulses will be sufficient for this application.

8.5.4 Faster magnesium-magnesium near-motional frequency oscillating magnetic field gradient gates

In spring 2020, after completing data taking for [8], we sought to speed up the magnesium-magnesium gate. Because the strength of the gradient puts a speed limit on the strength of the gate [166], this could not be achieved by turning up the microwave power: we needed to increase the strength of the gradient. Given that the gate presented in [8, 163] used ~ 1.2 A of current in each of the three gradient electrodes, we worried that the heat load on the trap could be significant. However, after careful calibration and re-nulling of the magnetic field that accompanies the oscillating

Figure 8.7: Coherent exchange of approximately a single quantum of motion between an in-phase and out-of-phase mode of a calcium-calcium crystal is shown at three different times. If we want to use the coupling drive during gate interactions, we want it to remain coherent for longer than the length of a possible gate interaction, so this demonstration of coherent exchange out to 3 ms is an important proof of principle. Some heating of both modes can be seen in the increasing amplitude as the scan time increases, which is to be expected. These demonstrations used a coupling drive strength of 0.6 V_{pp} and produced a coupling rate of 270 kHz, or a swap duration just under 2 μ s.

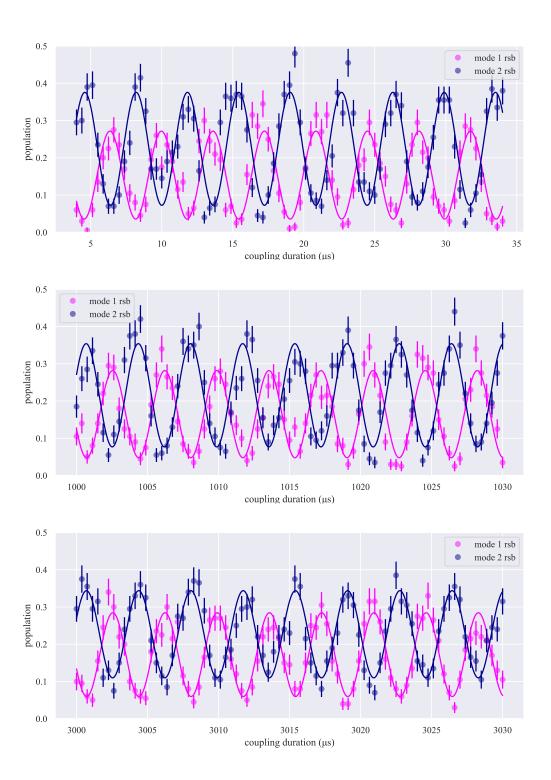


Figure 8.8: Plot of the red sidebands for mode 1 and mode 2 as the coupling frequency is scanned over the mode difference frequency. As with all of these scans, we start by intentionally adding motion mode 2 so that there is energy in the system to be exchanged. A Gaussian fit to the center of both peaks returns a FWHM of roughly 80 kHz.

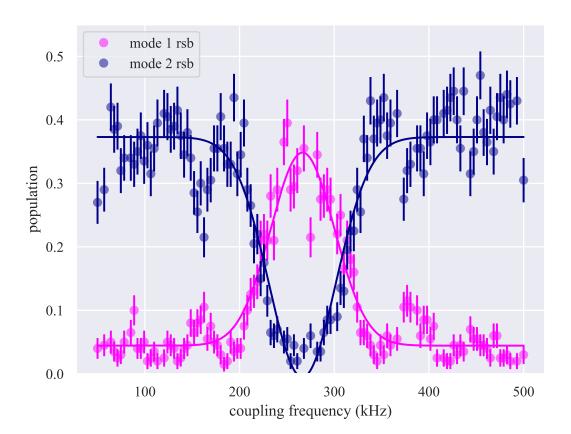
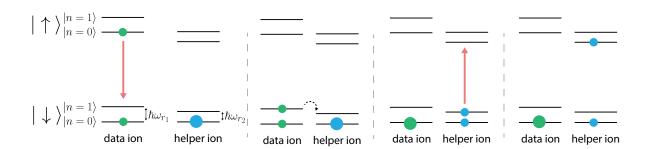


Figure 8.9: Illustration showing how the state of the data ion can be mapped onto the state of the helper ion using mode-mode coupling between two separate motional modes. For laser-free sidebands with uneven mode participation, this allows the operation to perform each sideband on a mode where the ion strongly participates, speeding up the overall operation. The red arrows represent the red sidebands, as in Fig. 4.2, and the black dotted line here represents the motional swap pulse.



gradient following the process described in [163], we were able to achieve gate durations of 350 μ s for an eight-loop gate, although fidelities were not quite as high as the slower gates. Our peak bootstrapped, SPAM-corrected fidelity³ for these fast gates was 99.7 $^{+0.15}_{-0.14}\%$.

These fast gates were performed the day before our lab was shut down for four months due to the COVID-19 pandemic. Although at the time this was a speed record for high-fidelity laser-free entangling gates, we never wrote up the result as we assumed that we would soon be able to perform randomized benchmarking on the gate interaction using calcium for sympathetic cooling and report the speed record in that manner along with a more rigorous measurement of the fidelity than we were able to provide in [8]. However, we have not yet been able to achieve a three-ion crystal (two magnesium ions to perform the gate and one calcium ion for sympathetic cooling) to perform this work and so these fast gates were never published. Recent work has demonstrated high-fidelity laser-free gates with gate durations of 310 μ s [45] and 120 μ s [9].

8.5.5 Calcium-calcium near-motional frequency oscillating magnetic field gradient gate

Because the Zeeman qubit in 40 Ca⁺ is sensitive to magnetic fields, we can perform the near-motional frequency oscillating magnetic field gradient gate between two such qubits. We spent less than one day on this gate before we needed to turn our attention to other matters and so there is extensive optimization that could still be done. Fidelities were \sim 75% with a gate time of 800 μ s, but this was with no Doppler cooling, away from the IDD point that provides insensitivity to qubit frequency noise, on the center-of-mass mode which experiences higher heating rates than the rocking mode, and used only two loops rather than the eight loops used in the high-fidelity demonstrations from [8].

³Using the same analysis methods as [8].

Figure 8.10: Parity flop of the state produced from our faster gates (350 μ s gate duration) using a pair of magnesium ions.

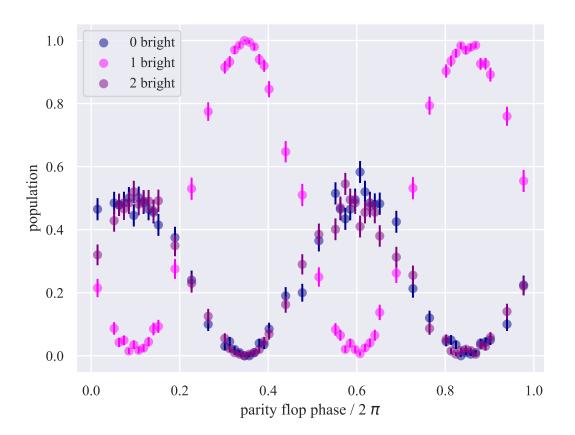
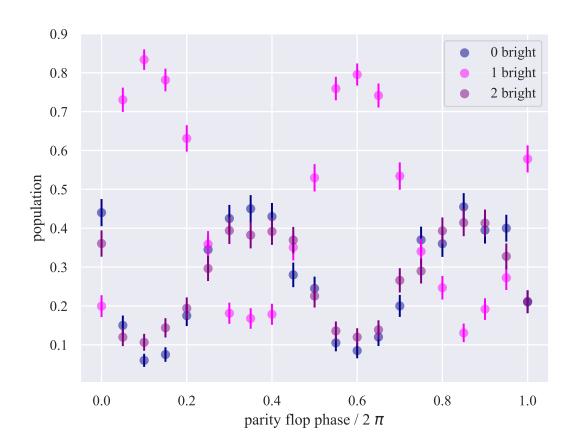


Figure 8.11: Parity flop of the state produced from the near-motional frequency oscillating magnetic field gradient gate scheme on a pair of calcium ions, which had not been previously demonstrated. Many calibrations were not ideal when this data was taken.



Mixed-species laser-free entangling gates

While laser-based mixed-species entangling gates have been demonstrated between different isotopes [182] and different ion species [183, 184], as of this writing no one has demonstrated a mixed-species laser-free entangling gate. In order to understand the challenges involved, we need to return to the details of the motional modes of mixed species ion crystals and also remember that the microwave electrode geometry of this trap restricts the gates to radial modes. In this chapter we will start with the basic near-motional frequency oscillating magnetic field gradient gate as described in Chapter 3 and [8, 149, 163], operating on a single mode of a single-species two-ion crystal. From there we will extend the gate to multiple species beginning with the most straightforward case and gradually relaxing assumptions.

Due to the difficulties loading and keeping mixed-species ion crystals described in Chapter 7, this chapter remains theoretical, but simulations are interspersed throughout. The parameter space for these gates is very large, and these simulations are intended more as a proof of principle that good entangling gate fidelities should be achievable with reasonable gate durations and realistic experimental parameters for our system than a thorough investigation of optimal gate parameters.

9.1 Single-species near-motional frequency oscillating magnetic field gradient gate

We described the basic gate Hamiltonian and the interactions it can generate in Chapter 3.1, where we used the notation from [149] and [8]. In this chapter, we will continue to use that notation and expand upon it when necessary. For easy reference, we will reproduce here the ion frame Hamiltonian

$$H(t) = 2\hbar\Omega_{\mu}\cos(\delta t)\hat{\sigma}_{x} + 2\hbar\Omega_{g}\cos(\omega_{g}t)\hat{\sigma}_{z}[\hat{a}e^{-i\omega_{r}t} + \hat{a}^{\dagger}e^{i\omega_{r}t}]. \tag{9.1}$$

Recall that Ω_{μ} is the microwave Rabi rate, Ω_{g} is the gradient Rabi rate, δ is the sideband detuning, ω_{g} is the gradient oscillation frequency, and ω_{r} is the motional mode frequency. To transform into the bichromatic interaction picture, we take the operator (from [149] Eq. 10)

$$\hat{U}(t) = e^{-i\frac{2\Omega_{\mu}\sin\delta t}{\delta}\hat{\sigma}_i} \tag{9.2}$$

and apply the operators \hat{U} and \hat{U}^{\dagger} such that

$$\hat{H}_I \approx \hat{U}^{\dagger} \hat{H} \hat{U}. \tag{9.3}$$

After making the approximations and selecting appropriate parameters as covered in Chapter 3.1 and in more detail in [149], we get

$$\hat{H}_I \approx \hbar \Omega_g J_2(\frac{4\Omega_\mu}{\delta}) \hat{S}_z[\hat{a}e^{i\Delta t} + \hat{a}^\dagger e^{-i\Delta t}], \tag{9.4}$$

where δ is the microwave sideband detuning from the qubit frequency, ω_r is the chosen motional mode frequency, \hat{a} and \hat{a}^{\dagger} are the annihilation and creation operators the for motional mode, \hat{S}_i is the multi-ion spin operator $\hat{S}_i = \sum_i \hat{\sigma}_i$, and Δ is the gate detuning such that $\Delta = 2n\delta - (\omega_r - \omega_g)$.

9.2 Extending the gate to multiple ion species

In theory, this gate should be straightforward to extend to multiple ion species. Because the strong oscillating gradient provides the spin-motion coupling regardless of species, a second species with a different qubit frequency only requires one additional pair of microwave sidebands.

In contrast to a mixed-species laser-based gate where lasers of different wavelength need to be phase-stabilized to each other, the technical generation of the control fields is much simpler. However, there are a number of additional factors that complicate the actual gate interaction.

Adding the second pair of microwave sidebands detuned around the new species' qubit frequency gives the ion frame Hamiltonian

$$\hat{H}(t) = 2\hbar\Omega_{\mu_1}\cos(\delta t)\hat{\sigma}_{x1} + 2\hbar\Omega_{\mu_2}\cos(\delta t)\hat{\sigma}_{x2} + 2\hbar\cos(\omega_g t)(\Omega_{g1}\xi_1\hat{\sigma}_{z1} + \Omega_{g2}\xi_2\hat{\sigma}_{z2})[\hat{a}e^{-i\omega_r t} + \hat{a}^{\dagger}e^{i\omega_r t}]$$

$$(9.5)$$

where the microwave Rabi rates (Ω_{μ}) may be different but could in principle be matched, the δ 's are the same because we have just one shared motional mode ω_r , and we are now being explicit about which ion species the spin operators are acting on. The coefficients ξ_i describe the normalized mode participation of each ion in the shared motional mode. The gradient Rabi rates Ω_{gi} are not necessarily the same; they depend on $d\omega_0/dB_z$ which is close but not identical for our two species.

In the bichromatic interaction picture, it is easy to see that these mode participations modify the effective gate speed

$$\hat{H}_{I} \approx \hbar (\Omega_{g1} J_{2}(\frac{4\Omega_{\mu_{1}}}{\delta}) \xi_{1} \hat{\sigma}_{z1} + \Omega_{g2} J_{2}(\frac{4\Omega_{\mu_{1}}}{\delta}) \xi_{2} \hat{\sigma}_{z2}) [\hat{a} e^{i\Delta t} + \hat{a}^{\dagger} e^{-i\Delta t}]. \tag{9.6}$$

If we had access to a radial mode with roughly equal mode participation for both ion species, $\xi_1 \approx \xi_2$, then we could adjust the strength of the microwave sidebands on each ion species to compensate for any slight imbalances in the mode participation or gradient Rabi rate, similar to

¹These sidebands could also be in the rf frequency range, though I will use the term microwave in this chapter for simplicity.

what [183] did for their mixed-species laser-based gate.

However, as detailed in Chapter 4, the radial mode participations are not roughly equal in this apparatus. For the mass imbalance between $^{25}\text{Mg}^+$ and $^{40}\text{Ca}^+$ and standard trapping parameters the radial modes have participation ratios in the range $\sim 10:1$ to $\sim 20:1$ —roughly an order of magnitude worse than for the axial modes often used in mixed-species entangling gates. While it is possible to do a gate on one of these modes, it would require a gate duration roughly 10x longer than for a same-species gate.

We plan to speed up the mixed-species entanglement by picking two modes, one that is calcium-dominant and one which is magnesium-dominant, and then using a parametric coupling drive [28, 30, 212] to couple those modes during the gate. We will start by assuming that we have completely independent motional modes for each species (i.e. the calcium does not participate at all in the magnesium-dominant mode and vice versa), and then later account for non-independent modes.

9.2.1 Independent motional modes with added coupling

The simplest coupled two-mode gate is one in which neither ion participates in the other ion's mode. This may or may not be a good assumption depending on your system. For us, the participation of the non-dominant ion in the other's mode hovers around the 5-10% range over a large range of trap parameters. For now, we will assume complete independence of the motional modes.

We are going to begin in the ion frame (i.e. the interaction picture with respect to the bare ion Hamiltonian) with the Hamiltonian from [149], extend it to multiple species with different motional modes this time, then add a coupling term.

We start in the ion frame with

$$\hat{H}_1(t) = 2\hbar\Omega_{\mu 1}\cos(\delta_1 t)\hat{\sigma}_{x1} + 2\hbar\Omega_{g1}\cos(\omega_g t)\hat{\sigma}_{z1}[\hat{a}e^{-i\omega_{r1}t} + \hat{a}^{\dagger}e^{i\omega_{r1}t}]$$

$$(9.7)$$

and

$$\hat{H}_2(t) = 2\hbar\Omega_{\mu 2}\cos(\delta_2 t)\hat{\sigma}_{x2} + 2\hbar\Omega_{g2}\cos(\omega_g t)\hat{\sigma}_{z2}[\hat{b}e^{-i\omega_{r2}t} + \hat{b}^{\dagger}e^{i\omega_{r2}t}]$$
(9.8)

where 1 and 2 denote our two ion species, which now have different mode frequencies ω_{ri} and thus microwave detunings (δ_i) , in addition to their different microwave Rabi rates $(\Omega_{\mu i})$, and gradient Rabi rates (Ω_{gi}) . Additionally, the annihilation and creation operators are different for the two modes $(\hat{a} \text{ for mode 1 and } \hat{b} \text{ for mode 2})$. We are essentially doing parallel single-ion gates. We can transform into the bichromatic interaction picture easily because all the operators for separate species commute with each other, and if we fix the sideband detunings (δ_i) to be equal, which can be done by placing the gradient frequency exactly between the two mode frequencies, we get

$$\hat{H}_{I} \approx \hbar \Omega_{g1} J_{2} \left(\frac{4\Omega_{\mu 1}}{\delta}\right) \hat{\sigma}_{z1} \left[\hat{a}e^{i\Delta t} + \hat{a}^{\dagger}e^{-i\Delta t}\right] + \hbar \Omega_{g2} J_{2} \left(\frac{4\Omega_{\mu 2}}{\delta}\right) \hat{\sigma}_{z2} \left[\hat{b}e^{i\Delta t} + \hat{b}^{\dagger}e^{-i\Delta t}\right]. \tag{9.9}$$

This interaction cannot generate entanglement between the ions because we have no terms that couple them. We can now add a coupling term, derived in Chapter 8

$$\hat{H}_c = \hbar g_{ab} (\hat{a}\hat{b}^\dagger + \hat{a}^\dagger \hat{b}), \tag{9.10}$$

and transform this complete ion frame Hamiltonian into the bichromatic interaction picture. We did this transformation for the gradient term above, and since the coupling Hamiltonian and bichromatic Hamiltonian commute with each other, the beamsplitter term is unchanged by the transformation. The final, full Hamiltonian is

$$\hat{H}_{I} \approx \hbar \Omega_{g1} J_{2}(\frac{4\Omega_{\mu 1}}{\delta}) \hat{\sigma}_{z1} [\hat{a}e^{i\Delta t} + \hat{a}^{\dagger}e^{-i\Delta t}] + \hbar \Omega_{g2} J_{2}(\frac{4\Omega_{\mu 2}}{\delta}) \hat{\sigma}_{z2} [\hat{b}e^{i\Delta t} + \hat{b}^{\dagger}e^{-i\Delta t}] + \hbar g(\hat{a}\hat{b}^{\dagger} + \hat{a}^{\dagger}\hat{b}).$$
(9.11)

It is going to be easier to understand the dynamics if we go into the dressed basis of the coupled modes, where the Hamiltonian will be diagonalized [28].

9.2.1.1 Dressed mode basis

In terms of the standard basis vectors \hat{a} and \hat{b} , the dressed mode vectors are [28]

$$\hat{c} = \frac{\hat{a} + \hat{b}}{\sqrt{2}} \tag{9.12}$$

and

$$\hat{d} = \frac{\hat{a} - \hat{b}}{\sqrt{2}} \tag{9.13}$$

so conversely

$$\hat{a} = \frac{\hat{c} + \hat{d}}{\sqrt{2}} \tag{9.14}$$

and

$$\hat{b} = \frac{\hat{c} - \hat{d}}{\sqrt{2}}.\tag{9.15}$$

Putting first just the coupling term into this basis

$$\hat{H}_C = \frac{\hbar g}{2} \left(\hat{c}^\dagger \hat{c} - \hat{d}^\dagger \hat{d} \right). \tag{9.16}$$

And now the rest of the Hamiltonian

$$\begin{split} \hat{H}_{I} \approx & \hbar \Omega_{g1} J_{2}(\frac{4\Omega_{\mu 1}}{\delta}) \hat{\sigma}_{z1} \left[\frac{\hat{c} + \hat{d}}{\sqrt{2}} e^{i\Delta t} + \frac{\hat{c}^{\dagger} + \hat{d}^{\dagger}}{\sqrt{2}} e^{-i\Delta t} \right] \\ & + \hbar \Omega_{g2} J_{2}(\frac{4\Omega_{\mu 2}}{\delta}) \hat{\sigma}_{z2} \left[\frac{\hat{c} - \hat{d}}{\sqrt{2}} e^{i\Delta t} + \frac{\hat{c}^{\dagger} - \hat{d}^{\dagger}}{\sqrt{2}} e^{-i\Delta t} \right] + \hbar g(\hat{c}^{\dagger} \hat{c} - \hat{d}^{\dagger} \hat{d}) \\ &= \frac{\hbar}{\sqrt{2}} \left(\Omega_{g1} J_{2}(\frac{4\Omega_{\mu 1}}{\delta}) \hat{\sigma}_{z1} + \Omega_{g2} J_{2}(\frac{4\Omega_{\mu 2}}{\delta}) \hat{\sigma}_{z2} \right) [\hat{c}e^{i\Delta t} + \hat{c}^{\dagger} e^{-i\Delta t}] \\ &+ \frac{\hbar}{\sqrt{2}} \left(\Omega_{g1} J_{2}(\frac{4\Omega_{\mu 1}}{\delta}) \hat{\sigma}_{z1} - \Omega_{g2} J_{2}(\frac{4\Omega_{\mu 2}}{\delta}) \hat{\sigma}_{z2} \right) [\hat{d}e^{i\Delta t} + \hat{d}^{\dagger} e^{-i\Delta t}] + \hbar g(\hat{c}^{\dagger} \hat{c} - \hat{d}^{\dagger} \hat{d}). \end{split}$$

If we define each effective gate speed separately we can simplify the expression, so let

$$\tilde{\Omega}_1 \equiv \frac{\Omega_{g1}}{\sqrt{2}} J_2(\frac{4\Omega_{\mu 1}}{\delta}) \tag{9.18}$$

and

$$\tilde{\Omega}_2 \equiv \frac{\Omega_{g2}}{\sqrt{2}} J_2(\frac{4\Omega_{\mu 2}}{\delta}) \tag{9.19}$$

giving us

$$\hat{H}_{I} \approx \hbar \left(\tilde{\Omega}_{1} \hat{\sigma}_{z1} + \tilde{\Omega}_{2} \hat{\sigma}_{z2} \right) \left[\hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t} \right] + \hbar \left(\tilde{\Omega}_{1} \hat{\sigma}_{z1} - \tilde{\Omega}_{2} \hat{\sigma}_{z2} \right) \left[\hat{d}e^{i\Delta t} + \hat{d}^{\dagger}e^{-i\Delta t} \right] + \hbar g(\hat{c}^{\dagger}\hat{c} - \hat{d}^{\dagger}\hat{d}),$$

$$(9.20)$$

the independent mode Hamiltonian in the dressed mode basis. To see how the coupling interacts with these terms, it will be helpful to go into the interaction picture with respect to the coupling term now.

9.2.1.2 Coupling interaction picture

To go into the interaction picture with respect to the coupling Hamiltonian, we define the operator

$$\hat{U} = \exp\left(-i\hat{H}_c t/\hbar\right) = \exp\left(-ig(\hat{c}^{\dagger}\hat{c} - \hat{d}^{\dagger}\hat{d})t\right)$$
(9.21)

and now we transform the other part of the Hamiltonian:

$$\hat{H}_{CI} = \hat{U}^{\dagger} \hat{H}_{I} \hat{U}
= e^{ig(\hat{c}^{\dagger} \hat{c} - \hat{d}^{\dagger} \hat{d})t} \Big[\hbar \left(\tilde{\Omega}_{1} \hat{\sigma}_{z1} + \tilde{\Omega}_{2} \hat{\sigma}_{z2} \right) \left[\hat{c} e^{i\Delta t} + \hat{c}^{\dagger} e^{-i\Delta t} \right]
+ \hbar \left(\tilde{\Omega}_{1} \hat{\sigma}_{z1} - \tilde{\Omega}_{2} \hat{\sigma}_{z2} \right) \left[\hat{d} e^{i\Delta t} + \hat{d}^{\dagger} e^{-i\Delta t} \right] \Big] e^{-ig(\hat{c}^{\dagger} \hat{c} - \hat{d}^{\dagger} \hat{d})t}
= \hbar \left(\tilde{\Omega}_{1} \hat{\sigma}_{z1} + \tilde{\Omega}_{2} \hat{\sigma}_{z2} \right) e^{ig\hat{c}^{\dagger} \hat{c}t} \left[\hat{c} e^{i\Delta t} + \hat{c}^{\dagger} e^{-i\Delta t} \right] e^{-ig\hat{c}^{\dagger} \hat{c}t}
+ \hbar \left(\tilde{\Omega}_{1} \hat{\sigma}_{z1} - \tilde{\Omega}_{2} \hat{\sigma}_{z2} \right) e^{-ig\hat{d}^{\dagger} \hat{d}t} \left[\hat{d} e^{i\Delta t} + \hat{d}^{\dagger} e^{-i\Delta t} \right] e^{ig\hat{d}^{\dagger} \hat{d}t}.$$
(9.22)

The \hat{c} and \hat{d} portions commute with each other and are going to be the same up to a sign so we will only work out one explicitly. Taking only the relevant portion and applying the Baker-Campbell-Hausdorf theorem which states that $e^A B e^{-A} = B + [A, B]/2 + ...$ we obtain

$$e^{ig\hat{c}^{\dagger}\hat{c}t}\left[\hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t}\right]e^{-ig\hat{c}^{\dagger}\hat{c}t} = \hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t} + \frac{1}{2}\left[ig\hat{c}^{\dagger}\hat{c}t, \hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t}\right] + \dots$$
(9.23)

Calculating just the first commutator gives

$$\begin{bmatrix}
ig\hat{c}^{\dagger}\hat{c}t, \hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t}
\end{bmatrix} = igt\left[(\hat{c}^{\dagger}\hat{c}\hat{c} - \hat{c}\hat{c}^{\dagger}\hat{c})e^{i\Delta t} + (\hat{c}^{\dagger}\hat{c}\hat{c}^{\dagger} - \hat{c}^{\dagger}\hat{c}^{\dagger}\hat{c})e^{-i\Delta t} \right]
= -igt\left[\hat{c}e^{i\Delta t} - \hat{c}^{\dagger}e^{-i\Delta t} \right].$$
(9.24)

We can use this to calculate the next commutator, between the above result and the original matrix A:

$$\left[ig\hat{c}^{\dagger}\hat{c}t, -igt[\hat{c}e^{i\Delta t} - \hat{c}^{\dagger}e^{-i\Delta t}]\right] = -(igt)^{2} \left[(\hat{c}^{\dagger}\hat{c}\hat{c} + \hat{c}\hat{c}^{\dagger}\hat{c})e^{i\Delta t} - (\hat{c}^{\dagger}\hat{c}\hat{c}^{\dagger} + \hat{c}^{\dagger}\hat{c}^{\dagger}\hat{c})e^{-i\Delta t} \right]
= (-igt)^{2} [\hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t}]$$
(9.25)

where we can see it has come back to the original term. We can now see that this series of commutators is going to be cyclical and in fact is going to form a Taylor series, and the d portion of the Hamiltonian will be the same but with an opposite sign because A has the opposite sign:

$$\hat{H}_{CI} = \hbar \left(\tilde{\Omega}_1 \hat{\sigma}_{z1} + \tilde{\Omega}_2 \hat{\sigma}_{z2} \right) \sum_{n=0} \frac{(-igt)^n}{n!} \left[\hat{c}e^{i\Delta t} + (-1)^n \hat{c}^{\dagger} e^{-i\Delta t} \right]$$

$$+ \hbar \left(\tilde{\Omega}_1 \hat{\sigma}_{z1} - \tilde{\Omega}_2 \hat{\sigma}_{z2} \right) \sum_{n=0} \frac{(igt)^n}{n!} \left[\hat{d}e^{i\Delta t} + (-1)^n \hat{d}^{\dagger} e^{-i\Delta t} \right].$$

$$(9.26)$$

Recognizing the Taylor expansion of $e^x = \sum x^n/n!$ allows us to write this as

$$\hat{H}_{CI} = \hbar \left(\tilde{\Omega}_1 \hat{\sigma}_{z1} + \tilde{\Omega}_2 \hat{\sigma}_{z2} \right) \left[\hat{c}e^{i(\Delta - g)t} + \hat{c}^{\dagger}e^{-i(\Delta - g)t} \right]$$

$$+ \hbar \left(\tilde{\Omega}_1 \hat{\sigma}_{z1} - \tilde{\Omega}_2 \hat{\sigma}_{z2} \right) \left[\hat{d}e^{i(\Delta + g)t} + \hat{d}^{\dagger}e^{-i(\Delta + g)t} \right].$$

$$(9.27)$$

It is now apparent that the coupling rate modifies the gate detuning, in a way that is opposite for each mode.

9.2.2 Most complicated: non-independent motional modes with added coupling

Everything we've done so far has assumed that we have completely independent motional modes - one in which only magnesium participates, and one in which only calcium participates. Finally, we relax this assumption to account for both ions participating in both modes and repeat the transformations above to see how this affects the gate Hamiltonian.

We can parametrize the uneven participation as an angle θ such that

$$\vec{r}_a = \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} \tag{9.28}$$

and

$$\vec{r_b} = \begin{bmatrix} \sin(\theta) \\ -\cos(\theta) \end{bmatrix} \tag{9.29}$$

describe our two modes and the basis states are the local modes for calcium and magnesium respectively. When θ is $\pi/2$, the participation of both ions in both modes is equal. We will consider cases where θ is small, reflective of our actual mode participations.

Adding this into the (ion frame, non-dressed basis) Hamiltonian looks like

$$\hat{H}_{I} \approx \hbar \Omega_{g} J_{2} \left(\frac{4\Omega_{\mu 1}}{\delta}\right) (\cos(\theta)\hat{\sigma}_{z1} + \sin(\theta)\hat{\sigma}_{z2}) [\hat{a}e^{i\Delta t} + \hat{a}^{\dagger}e^{-i\Delta t}]$$

$$+ \hbar \Omega_{g} J_{2} \left(\frac{4\Omega_{\mu 2}}{\delta}\right) (\sin(\theta)\hat{\sigma}_{z1} - \cos(\theta)\hat{\sigma}_{z2}) [\hat{b}e^{i\Delta t} + \hat{b}^{\dagger}e^{-i\Delta t}].$$

$$(9.30)$$

We can now perform a gate without adding the coupling drive, as there is participation of each ion in both modes, as shown in Fig. 9.2. However, as mentioned in the introduction, the gate duration will be inversely proportional the mode participation $\sin(\theta)$. If we add in the coupling drive here, we can speed up the mixed-species entangling gates; the results of such simulations are shown in Fig. ??.

9.2.2.1 Dressed mode basis

We once again want to look at this Hamiltonian in the dressed mode basis, so write

$$\hat{H}_{I} \approx \hbar \Omega_{g} J_{2} \left(\frac{4\Omega_{\mu 1}}{\delta}\right) \left(\cos(\theta)\hat{\sigma}_{z1} + \sin(\theta)\hat{\sigma}_{z2}\right) \left[\frac{\hat{c} + \hat{d}}{\sqrt{2}} e^{i\Delta t} + \frac{\hat{c}^{\dagger} + \hat{d}^{\dagger}}{\sqrt{2}} e^{-i\Delta t}\right]$$

$$+ \hbar \Omega_{g} J_{2} \left(\frac{4\Omega_{\mu 2}}{\delta}\right) \left(\sin(\theta)\hat{\sigma}_{z1} - \cos(\theta)\hat{\sigma}_{z2}\right) \left[\frac{\hat{c} - \hat{d}}{\sqrt{2}} e^{i\Delta t} + \frac{\hat{c}^{\dagger} - \hat{d}^{\dagger}}{\sqrt{2}} e^{-i\Delta t}\right].$$

$$(9.31)$$

Now using the same effective Rabi rates as before we can write

$$\hat{H}_{I} \approx \hbar \left(\tilde{\Omega}_{1}(\cos(\theta)\hat{\sigma}_{z1} + \sin(\theta)\hat{\sigma}_{z2}) + \tilde{\Omega}_{2}(\sin(\theta)\hat{\sigma}_{z1} - \cos(\theta)\hat{\sigma}_{z2}) \right) \left[\hat{c}e^{i\Delta t} + \hat{c}^{\dagger}e^{-i\Delta t} \right]
+ \hbar \left(\tilde{\Omega}_{1}(\cos(\theta)\hat{\sigma}_{z1} + \sin(\theta)\hat{\sigma}_{z2}) - \tilde{\Omega}_{2}(\sin(\theta)\hat{\sigma}_{z1} - \cos(\theta)\hat{\sigma}_{z2}) \right) \left[\hat{d}e^{i\Delta t} + \hat{d}^{\dagger}e^{-i\Delta t} \right] + \hbar g(\hat{c}^{\dagger}\hat{c} - \hat{d}^{\dagger}\hat{d}).$$
(9.32)

So just as in the other basis (as it should be), this creates small (if θ is small) perturbations to the effective gate speed.

9.2.2.2 Coupling interaction picture

The motional portion of the Hamiltonian is independent of the actual mode participations, so the same transformation into the coupling interaction picture is valid, and the dressed, coupled mode-basis Hamiltonian is

$$\hat{H}_{CI} = \hbar \left(\tilde{\Omega}_1(\cos(\theta)\hat{\sigma}_{z1} + \sin(\theta)\hat{\sigma}_{z2}) + \tilde{\Omega}_2(\sin(\theta)\hat{\sigma}_{z1} - \cos(\theta)\hat{\sigma}_{z2}) \right) \left[\hat{c}e^{i(\Delta - g)t} + \hat{c}^{\dagger}e^{-i(\Delta - g)t} \right]$$

$$+ \hbar \left(\tilde{\Omega}_1(\cos(\theta)\hat{\sigma}_{z1} + \sin(\theta)\hat{\sigma}_{z2}) - \tilde{\Omega}_2(\sin(\theta)\hat{\sigma}_{z1} - \cos(\theta)\hat{\sigma}_{z2}) \right) \left[\hat{d}e^{i(\Delta + g)t} + \hat{d}^{\dagger}e^{-i(\Delta + g)t} \right]$$

$$(9.33)$$

for a coupled, mixed-species gate with non-independent motional modes.

9.3 Simulations

The simulations in this section are performed in the dressed basis and coupling interaction picture. All simulation parameters are normalized to the effective Rabi frequency on ion 1 $\tilde{\Omega}_1$, which is not assumed to be equal to $\tilde{\Omega}_2$. For the estimated gate parameters given in this section, we pick $\tilde{\Omega}_1 = 10$ kHz, based on values from [8].

The plots in this section show gate detuning scans, where we are scanning the parameter Δ . In this experiment, this is how gates are experimentally calibrated due to thermal effects in the amplifiers used to generate the gradients. Scanning gate detunings instead of durations means that pulse durations are uniform during the scan and do not cause distortions. We can select the gate detuning that will create entanglement for a given gate duration by selecting the point at which the $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ populations are each 0.5 and the $|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle$ populations are 0. This is the point where our given gate parameters have generated the Bell state $|\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle$, the goal of this entangling interaction.

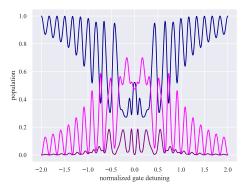
9.3.1 Independent motional modes, with and without coupling

If we assume the modes are independent and do not apply a coupling drive, we cannot generate mixed-species entanglement. However, applying a 2 kHz coupling drive could allow us to drive a mixed-species gate in less than 300 μ s with reasonable experimental parameters, as simulated in Fig. 9.1. However, in this apparatus, the modes are not actually independent and the uneven participations are large enough to make the assumption a bad one.

9.3.2 Non-independent motional modes, no coupling

For these simulations, we use an estimated participation of each ion in its non-dominant mode of 10% (see Chapter 4 for more details on the source of this value). This allows us to see entanglement without adding the coupling drive, but it requires quite a long gate duration of roughly 2.9 ms, as simulated in Fig. 9.2.

Figure 9.1: Gate detuning scans for independent mixed-species motional modes, without coupling on the left and with coupling on the right. Without coupling we cannot generate mixed-species entanglement, regardless of gate duration. We show a scan with a duration equivalent to 600 μ s. Adding a weak coupling drive creates an entangling gate: adding a 2 kHz coupling drive, we can see a point where a high-fidelity Bell state could be generated (marked with a black dotted line) at estimated gate duration of 290 μ s.



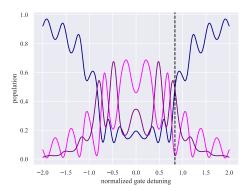
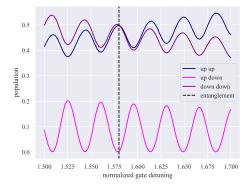
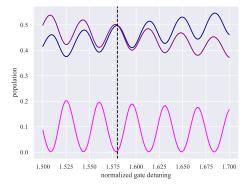


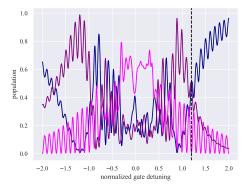
Figure 9.2: Gate detuning scans for non-independent mixed-species motional modes, without coupling. Without coupling we *can* create a high-fidelity Bell state if both ions participate in both modes, but the gates are slow due to the weak participation of each ion in the other's dominant mode. These scans have an estimated gate duration of 2.9 ms. The scan on the right zooms in on the entanglement point from the left scan.

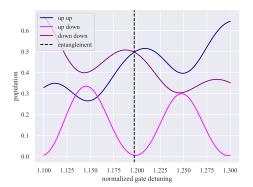




Adding a coupling drive of 2 kHz allows us to create a high-fidelity Bell state with a shorter interaction duration, as shown in Fig. 9.3.

Figure 9.3: Gate detuning scans for non-independent mixed-species motional modes, with coupling. Adding a 2 kHz coupling drive can speed up the interaction duration needed to create a high-fidelity Bell state. These scans have an estimated gate duration of 1 ms, nearly 3x faster than without the coupling drive.





9.4 Alternate method: PHRAP-style twisting

It has become clear over the course of this thesis that a mixed-species laser-free entangling gate with no mode-mode coupling is going to be an order of magnitude slower than a same-species laser-free gate. However, a parametric drive is not the only method of mode-mode coupling that could be used to increase gate speed.

We discussed with R. Tyler Sutherland the possibility of using the so-called phonon rapid adiabatic (PHRAP)-style twisting scheme developed for cooling [218] to speed up mixed-species laser-free gates. The technique entails either increasing the axial confinement or relaxing the radial confinement until one axial and one radial mode are near-degenerate. Then a static electric field is applied to push the ions out of the rf null. Because of the mass dependence of the pseudopotential, the different species experience different displacements resulting in a twisted ion crystal. This projects the Coulomb force along both modes, generating mode-mode coupling. We could then

perform a gate on the coupled modes. This is an intriguing possibility and the necessary trap parameters seem to form a stable trap configuration, however due to the loading and lifetime difficulties we were unable to attempt such a gate. The technique only works for mixed species crystals as it relies on the differential response to applied radial electric fields, and so we could not test it on a single-species crystal, but it would require fewer applied fields and thus less calibration than the scheme described above.

10

Conclusions & future directions

In this work, we have presented a substantial rebuild of the existing system for laser-free gates on $^{25}\mathrm{Mg^+}$ to improve existing functionality and expand the system to add a second ion species. We overhauled the vacuum chamber, cryostat, laser systems, imaging system, control hardware and software, and demonstrated operation with $^{40}\mathrm{Ca^+}$ ions at a novel magnetic field, including laser-free entanglement and cooling. We also demonstrated mixed-species loading, transport, and cooling in this system for the first time, the closest to a trap surface these operations have been performed. ¹

We also presented theoretical proposals and simulations demonstrating feasibility for a mixedspecies ion system with one completely laser-free species. However, the substantial difficulties and instabilities in the mixed species loading and operation did not allow for experimental demonstration of these schemes on the timescale of this thesis.

10.1 Suggestions for improvements to this apparatus

There is still lots of work that could be done to improve the existing apparatus to a potentially workable point. Investigations are currently underway to find ways to stabilize and otherwise manage

¹As far as I'm aware!

the large and rapidly drifting stray fields which plagued mixed-species loading and transport. An interesting first step would be to change the calcium photoionization beams to also selectively ionize only the slower neutral atoms and see the effects this has on the charge landscape. If this is able to reduce the charging around the calcium load zone by the same amount as it did for the magnesium load zone, roughly 60%, then it is possible that the charge landscape on the right-hand side of the trap might be stable enough to resume merging on that side.

It is also possible to integrate periodic warmup cycles and perhaps even more on-table bakes to relax the stray fields that seem to accumulate. These could be performed over the weekend with minimal disruptions to uptime during the week. Of course, this apparatus has demonstrated itself to be state-of-the-art at single species operations and there is a lot of work that could still be done on parallel single species operations using the multi-well and transport capabilities that we have recently developed. In particular, there is a lot of interesting work to be done building on the squeezing and amplification interactions performed by Shaun Burd [139] in this apparatus. However, the most valuable lessons learned here are for future generations of these experiments.

10.2 Suggestions for future generations of laser-free mixed-species experiments

The bulk of our difficulties came from our trap layout, chosen ion species, and loading configurations, only one of which was feasible to change on the timescale of this PhD.

Issues with the trap design were primarily with insufficient electrodes to independently control the two wells as they approached each other, as well as the electrodes being quite large and distant from the ion. A trap which had more finely spaced electrodes that were closer to the ions, and perhaps a dedicated 'split' electrode such as in the original double well trap [219] would make compensation of ions during the merge process more feasible. Another small change that could end up making a big difference would be to increase the ion-electrode height. While our laser-free gate speed goes as as $1/d^2$ where d is the ion height above the trap, effects such as heating and sensitivity to stray charge also increase with decreased ion height. Embedding the microwave electrodes below

the surface to allow higher gate drive currents while shielding the ion from the accompanying oscillating electric field that can drive off-resonant spin flips has long been an objective for future generations of this trap. Doing this would also allow the DC electrodes to be placed closer to the trap axis, potentially decreasing voltage needed per electrode to compensate a given stray field. It seems plausible that we have so much difficulty with stray charges accumulated during the loading process compared to others because we are forced to create ions so close to the surface.

While changing our magnesium loading to be velocity-selective substantially decreased the stray fields around the magnesium load zone, the difference in magnitude of stray fields around the magnesium versus calcium load zones indicates that there is still a difference between the two species. Previous work has shown that shorter wavelength lasers cause more photoelectric stray fields than longer wavelengths, and it is possible that we are seeing multiple effects contributing to the load zone stray fields. A change of species to one with longer wavelength photoionization could potentially further improve trap stability.

More drastic modifications to the experiment design could include techniques such as MOT loading, which should improve loading rate while reducing load zone charging, or perhaps a truly separate load zone that is actually far enough from the experiment zone to avoid charging there. A load zone with a substantially higher ion height than the experiment zone could offer a compromise between charging by loading near the trap surface and slowing the gates with a higher ion height. Different load configurations involving neutral flux traveling in such a way that untrapped charged particles are less likely to impact the trap surface is another option: many traps use atomic beams which emerge from behind or beneath the trap such that the flux travels away from any trap surfaces.

In conclusion, this apparatus has demonstrated and continues to demonstrate high-fidelity laser-free operations on single species, and has demonstrated preliminary mixed-species capability which can be expanded upon. We have also learned a great deal about ways to improve future generations of laser-free experiments.

References

- [1] R. P. Feynman, "Simulating physics with computers", Int. Jour. of Theor. Phys. **21**, 467 (1982) (Cited on pp. 1, 9).
- [2] D. P. DiVincenzo, "Two-bit gates are universal for quantum computation", Phys. Rev. A 51, 1015 (1995) (Cited on pp. 1, 10).
- [3] J. Preskill, "Quantum computing in the nisq era and beyond", Quantum 2, 79 (2018) (Cited on p. 1).
- [4] W. Paul and H. Steinwedel, Zeitschrift für Naturforschung A 8, 448 (1953) (Cited on p. 2).
- [5] 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 (2016) (Cited on pp. 2, 10, 27).
- [6] 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 ⁹Be⁺ ion qubits", Phys. Rev. Lett. 117, 060505 (2016) (Cited on pp. 2, 10, 26, 238).
- [7] C. R. Clark, H. N. Tinkey, B. C. Sawyer, A. M. Meier, K. A. Burkhardt, C. M. Seck, C. M. Shappert, N. D. Guise, C. E. Volin, S. D. Fallek, H. T. Hayden, W. G. Rellergert, and K. R. Brown, "High-fidelity bell-state preparation with ⁴⁰Ca⁺ optical qubits", Phys. Rev. Lett. **127**, 130505 (2021) (Cited on p. 2).
- [8] 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 trapped ion qubits", Nature 597, 209 (2021) (Cited on pp. 2, 4, 5, 9, 10, 27, 30, 35, 37, 39, 137, 145, 149, 152, 153, 162).
- [9] C. M. Löschnauer, J. M. Toba, A. C. Hughes, S. A. King, M. A. Weber, R. Srinivas, R. Matt, R. Nourshargh, D. T. C. Allcock, C. J. Ballance, C. Matthiesen, M. Malinowski, and T. P. Harty, "Scalable, high-fidelity all-electronic control of trapped-ion qubits", arXiv, 2407.07694 (2024) (Cited on pp. 2, 4, 9, 10, 26, 30, 39, 149).
- [10] K. A. Gilmore, M. Affolter, R. J. Lewis-Swan, D. Barberena, E. Jordan, A. M. Rey, and J. J. Bollinger, "Quantum-enhanced sensing of displacements and electric fields with twodimensional trapped-ion crystals", Science 373, 673 (2021) (Cited on p. 2).
- [11] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, and C. Monroe, "Fault-tolerant control of an error-corrected qubit", Nature 598, 281 (2021) (Cited on p. 2).
- [12] L. Postler, S. Heuβen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt, P. Schindler, M. Müller, and T. Monz, "Demonstration of fault-tolerant universal quantum gate operations", Nature 605, 675 (2022) (Cited on p. 2).

- [13] M. DeCross, R. Haghshenas, M. Liu, E. Rinaldi, J. Gray, Y. Alexeev, C. H. Baldwin, J. P. Bartolotta, M. Bohn, E. Chertkov, J. Cline, J. Colina, D. DelVento, J. M. Dreiling, C. Foltz, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, D. Gresh, A. Hall, A. Hankin, A. Hansen, N. Hewitt, I. Hoffman, C. Holliman, R. B. Hutson, T. Jacobs, J. Johansen, P. J. Lee, E. Lehman, D. Lucchetti, D. Lykov, I. S. Madjarov, B. Mathewson, K. Mayer, M. Mills, P. Niroula, J. M. Pino, C. Roman, M. Schecter, P. E. Siegfried, B. G. Tiemann, C. Volin, J. Walker, R. Shaydulin, M. Pistoia, S. A. Moses, D. Hayes, B. Neyenhuis, R. P. Stutz, and M. Foss-Feig, "The computational power of random quantum circuits in arbitrary geometries", arXiv, 2406.02501 (2024) (Cited on pp. 2, 10).
- [14] G. Gabrielse, X. Fei, L. Haarsma, S. L. Rolston, R. Tjoelker, T. A. Trainor, H. Kalinowsky, J. Haas, and W. Kells, "First capture of antiprotons in an ion trap: progress toward a precision mass measurement and antihydrogen", Physica Scripta 1988, 36 (1988) (Cited on p. 2).
- [15] D. Leibfried, "Quantum state preparation and control of single molecular ions", New Journal of Physics 14, 023029 (2012) (Cited on p. 2).
- [16] C. Matthiesen, Q. Yu, J. Guo, A. M. Alonso, and H. Häffner, "Trapping electrons in a room-temperature microwave paul trap", Phys. Rev. X 11, 011019 (2021) (Cited on p. 2).
- [17] S. A. King, L. J. Spieß, P. Micke, A. Wilzewski, T. Leopold, E. Benkler, R. Lange, N. Huntemann, A. Surzhykov, V. A. Yerokhin, J. R. Crespo López-Urrutia, and P. O. Schmidt, "An optical atomic clock based on a highly charged ion", Nature **611**, 43 (2022) (Cited on p. 2).
- [18] M. Fan, R. A. Ready, H. Li, S. Kofford, R. Kwapisz, C. A. Holliman, M. S. Ladabaum, A. N. Gaiser, J. R. Griswold, and A. M. Jayich, "Laser cooling and trapping of ²²⁴Ra⁺", Phys. Rev. Res. 5, 043201 (2023) (Cited on p. 2).
- [19] M. S. Safronova, V. A. Dzuba, V. V. Flambaum, U. I. Safronova, S. G. Porsev, and M. G. Kozlov, "Highly charged ions for atomic clocks, quantum information, and search for α variation", Phys. Rev. Lett. **113**, 030801 (2014) (Cited on p. 2).
- [20] D. Budker, P. W. Graham, H. Ramani, F. Schmidt-Kaler, C. Smorra, and S. Ulmer, "Millicharged dark matter detection with ion traps", PRX Quantum 3, 010330 (2022) (Cited on p. 2).
- [21] K. Gaul, N. R. Hutzler, P. Yu, A. M. Jayich, M. Ilia š, and A. Borschevsky, "CP-violation sensitivity of closed-shell radium-containing polyatomic molecular ions", Phys. Rev. A 109, 042819 (2024) (Cited on pp. 2, 42).
- [22] Y. Zhou, J. O. Island, and M. Grau, "Quantum logic control and precision measurements of molecular ions in a ring trap: an approach for testing fundamental symmetries", Phys. Rev. A 109, 033107 (2024) (Cited on pp. 2, 42).
- [23] 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. Nat. Inst. Stand. Technol. 103, 259 (1998) (Cited on pp. 2, 4, 10, 13, 16, 20, 22, 29, 31, 41, 46, 237, 238).
- [24] D. Kielpinski, C. Monroe, and D. J. Wineland, "Architecture for a large-scale ion-trap quantum computer", Nature 417, 709 (2002) (Cited on pp. 2–4, 10, 20, 41, 238).

- [25] R. D. Delaney, L. R. Sletten, M. J. Cich, B. Estey, M. Fabrikant, D. Hayes, I. M. Hoffman, J. Hostetter, C. Langer, S. A. Moses, A. R. Perry, T. A. Peterson, A. Schaffer, C. Volin, G. Vittorini, and W. C. Burton, "Scalable multispecies ion transport in a grid based surfaceelectrode trap", arXiv, 2403.00756 (2024) (Cited on pp. 3, 4, 18, 42).
- [26] 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 (2000) (Cited on pp. 2, 42, 44, 46).
- [27] J. B. Wübbena, S. Amairi, O. Mandel, and P. O. Schmidt, "Sympathetic cooling of mixed-species two-ion crystals for precision spectroscopy", Phys. Rev. A 85, 043412 (2012) (Cited on pp. 2, 16, 42, 44).
- [28] D. J. Gorman, P. Schindler, S. Selvarajan, N. Daniilidis, and H. Häffner, "Two-mode coupling in a single-ion oscillator via parametric resonance", Physical Review A **89** (2014) (Cited on pp. 4, 138, 140–142, 155–157).
- [29] S. A. King, L. J. Spieß, P. Micke, A. Wilzewski, T. Leopold, J. R. Crespo López-Urrutia, and P. O. Schmidt, "Algorithmic ground-state cooling of weakly coupled oscillators using quantum logic", Phys. Rev. X 11, 041049 (2021) (Cited on pp. 4, 47).
- [30] P.-Y. Hou, J. J. Wu, S. D. Erickson, G. Zarantonello, A. D. Brandt, D. C. Cole, A. C. Wilson, D. H. Slichter, and D. Leibfried, "Indirect cooling of weakly coupled trapped-ion mechanical oscillators", Phys. Rev. X 14, 021003 (2024) (Cited on pp. 4, 47, 138, 140–142, 155).
- [31] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland, "Complete methods set for scalable ion trap quantum information processing", Science 325, 1227 (2009) (Cited on p. 4).
- [32] Y. Wan, D. Kienzler, S. D. Erickson, K. H. Mayer, T. R. Tan, J. J. Wu, H. M. Vasconcelos, S. Glancy, E. Knill, D. J. Wineland, A. C. Wilson, and D. Leibfried, "Quantum gate teleportation between separated qubits in a trapped-ion processor", Science 364, 875 (2019) (Cited on p. 4).
- [33] 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, 102 (2019) (Cited on pp. 4, 42).
- [34] K. Sosnova, A. Carter, and C. Monroe, "Character of motional modes for entanglement and sympathetic cooling of mixed-species trapped-ion chains", Phys. Rev. A 103, 012610 (2021) (Cited on pp. 4, 42).
- [35] M. M. van Mourik, "Control of multi-species trapped ions in a scalable quantum computer", PhD thesis (University of Innsbruck, 2023) (Cited on pp. 4, 42, 74, 118).
- [36] M. Marinelli, "Quantum information processing with mixed-species ion crystals", PhD thesis (ETH Zurich, 2020) (Cited on pp. 4, 42).
- [37] J. Wright, C. Auchter, C.-K. Chou, R. D. Graham, T. W. Noel, T. Sakrejda, Z. Zhou, and B. B. Blinov, "Scalable quantum computing architecture with mixed species ion chains", arXiv, 1410.0037 (2014) (Cited on pp. 4, 42).
- [38] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, "Spectroscopy using quantum logic", Science **309**, 749 (2005) (Cited on pp. 4, 41, 47–50, 130, 132).

- [39] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibrandt, "²⁷Al⁺ quantum-logic clock with a systematic uncertainty below 10⁻¹⁸", Phys. Rev. Lett. **123**, 033201 (2019) (Cited on pp. 4, 48, 130).
- [40] F. Wolf, Y. Wan, J. C. Heip, F. Gebert, C. Shi, and P. O. Schmidt, "Non-destructive state detection for quantum logic spectroscopy of molecular ions", Nature 530, 457 (2016) (Cited on pp. 4, 48, 130).
- [41] C.-w. Chou, C. Kurz, D. B. Hume, P. N. Plessow, D. R. Leibrandt, and D. Leibfried, "Preparation and coherent manipulation of pure quantum states of a single molecular ion", Nature 545, 203 (2017) (Cited on pp. 4, 48, 130).
- [42] P. Micke, T. Leopold, S. A. King, E. Benkler, L. J. Spieß, L. Schmöger, M. Schwarz, J. R. Crespo López-Urrutia, and P. O. Schmidt, "Coherent laser spectroscopy of highly charged ions using quantum logic", Nature **578**, 60 (2020) (Cited on pp. 4, 48, 130).
- [43] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland, "Microwave quantum logic gates for trapped ions", Nature 476, 181 (2011) (Cited on pp. 4, 18, 30, 35, 36).
- [44] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe, "Experimental entanglement of four particles", Nature 404, 256 (2000) (Cited on pp. 4, 9, 26).
- [45] M. A. Weber, M. F. Gely, R. K. Hanley, T. P. Harty, A. D. Leu, C. M. Löschnauer, D. P. Nadlinger, and D. M. Lucas, "Robust and fast microwave-driven quantum logic for trapped-ion qubits", arXiv, 2402.12955 (2024) (Cited on pp. 4, 30, 149).
- [46] U. Warring, C. Ospelkaus, Y. Colombe, R. Jördens, D. Leibfried, and D. J. Wineland, "Individual-ion addressing with microwave field gradients", Phys. Rev. Lett. 110, 173002 (2013) (Cited on pp. 5, 17, 37).
- [47] C. Piltz, T. Sriarunothai, A. Varón, and C. Wunderlich, "A trapped-ion-based quantum byte with 10-5 next-neighbour cross-talk", Nature Communications 5 (2014) (Cited on pp. 5, 37).
- [48] S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger, "Trapped-ion quantum logic with global radiation fields", Phys. Rev. Lett. 117, 220501 (2016) (Cited on pp. 5, 30, 38).
- [49] R. Srinivas, C. M. Löschnauer, M. Malinowski, A. C. Hughes, R. Nourshargh, V. Negnevitsky, D. T. C. Allcock, S. A. King, C. Matthiesen, T. P. Harty, and C. J. Ballance, "Coherent control of trapped-ion qubits with localized electric fields", Phys. Rev. Lett. 131, 020601 (2023) (Cited on pp. 5, 37).
- [50] E. R. Hudson and W. C. Campbell, "Laserless quantum gates for electric dipoles in thermal motion", Phys. Rev. A **104**, 042605 (2021) (Cited on p. 5).
- [51] 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 (2020) (Cited on pp. 5, 7, 28).
- [52] 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 (2020) (Cited on pp. 5, 18, 28, 29).

- [53] C. Mordini, A. R. Vasquez, Y. Motohashi, M. Müller, M. Malinowski, C. Zhang, K. K. Mehta, D. Kienzler, and J. P. Home, "Multi-zone trapped-ion qubit control in an integrated photonics qccd device", arXiv, 2401.18056 (2024) (Cited on pp. 5, 7, 18, 29).
- [54] 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 (2021) (Cited on pp. 5, 18).
- [55] L. K. Grover, "A fast quantum mechanical algorithm for database search", arXiv (1996) (Cited on p. 8).
- [56] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer", SIAM Journal on Computing 26, 1484 (1997) (Cited on p. 8).
- [57] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, "Quantum supremacy using a programmable superconducting processor", Nature 574, 505 (2019) (Cited on p. 9).
- [58] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, "Logical quantum processor based on reconfigurable atom arrays", Nature 626, 58 (2024) (Cited on pp. 9, 20).
- [59] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, B. M. Niedzielski, A. Melville, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, J. A. Grover, K. Serniak, and W. D. Oliver, "High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler", Phys. Rev. X 13, 031035 (2023) (Cited on p. 9).
- [60] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, "High-fidelity parallel entangling gates on a neutral-atom quantum computer", Nature 622, 268 (2023) (Cited on p. 9).
- [61] A. Cao, W. J. Eckner, T. L. Yelin, A. W. Young, S. Jandura, L. Yan, K. Kim, G. Pupillo, J. Ye, N. D. Oppong, and A. M. Kaufman, "Multi-qubit gates and 'schrödinger cat' states in an optical clock", arXiv, 2402.16289 (2024) (Cited on p. 9).
- [62] M. H. Abobeih, Y. Wang, J. Randall, S. J. H. Loenen, C. E. Bradley, M. Markham, D. J. Twitchen, B. M. Terhal, and T. H. Taminiau, "Fault-tolerant operation of a logical qubit in a diamond quantum processor", Nature **606**, 884 (2022) (Cited on p. 9).

- [63] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, "Quantum logic with spin qubits crossing the surface code threshold", Nature 601, 343 (2022) (Cited on p. 9).
- [64] A. H. Karamlou, I. T. Rosen, S. E. Muschinske, C. N. Barrett, A. Di Paolo, L. Ding, P. M. Harrington, M. Hays, R. Das, D. K. Kim, B. M. Niedzielski, M. Schuldt, K. Serniak, M. E. Schwartz, J. L. Yoder, S. Gustavsson, Y. Yanay, J. A. Grover, and W. D. Oliver, "Probing entanglement in a 2d hard-core bose-hubbard lattice", Nature 629, 561 (2024) (Cited on p. 9).
- [65]R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalava, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. D. T. Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. G. Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. Mccourt, M. McEwen, A. Megrant, B. M. Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skruzny, V. Smelyanskiy, W. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. V. Heidweiller, T. White, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, and N. Zhu, "Suppressing quantum errors by scaling a surface code logical qubit", arXiv, 2207.06431 (2022) (Cited on p. 9).
- [66] A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman, "An atomic boson sampler", Nature 629, 311 (2024) (Cited on p. 9).
- [67] M. P. da Silva, C. Ryan-Anderson, J. M. Bello-Rivas, A. Chernoguzov, J. M. Dreiling, C. Foltz, F. Frachon, J. P. Gaebler, T. M. Gatterman, L. Grans-Samuelsson, D. Hayes, N. Hewitt, J. Johansen, D. Lucchetti, M. Mills, S. A. Moses, B. Neyenhuis, A. Paz, J. Pino, P. Siegfried, J. Strabley, A. Sundaram, D. Tom, S. J. Wernli, M. Zanner, R. P. Stutz, and K. M. Svore, "Demonstration of logical qubits and repeated error correction with better-than-physical error rates", arXiv, 2404.02280 (2024) (Cited on p. 9).
- [68] P. Benioff, "The computer as a physical system: a microscopic quantum mechanical hamiltonian model of computers as represented by turing machines", Journal of Statistical Physics 22, 563 (1980) (Cited on p. 9).
- [69] Y. I. Manin, Vychislimoe i nevychislimoe (computable and noncomputable), 1980 (Cited on p. 9).

- [70] D. J. Wineland, R. E. Drullinger, and F. L. Walls, "Radiation-pressure cooling of bound resonant absorbers", Phys. Rev. Lett. 40, 1639 (1978) (Cited on p. 9).
- [71] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt, "Optical-sideband cooling of visible atom cloud confined in parabolic well", Phys. Rev. Lett. 41, 233 (1978) (Cited on p. 9).
- [72] D. J. Wineland and W. M. Itano, "Laser cooling of atoms", Phys. Rev. A 20, 1521 (1979) (Cited on p. 9).
- [73] W. Neuhauser, M. Hohenstatt, P. E. Toschek, and H. Dehmelt, "Localized visible Ba+mono-ion oscillator", Phys. Rev. A 22, 1137 (1980) (Cited on p. 9).
- [74] W. M. Itano and D. J. Wineland, "Precision measurement of the ground-state hyperfine constant of ²⁵Mg⁺", Phys. Rev. A **24**, 1364 (1981) (Cited on p. 9).
- [75] J. I. Cirac and P. Zoller, "Quantum computations with cold trapped ions", Phys. Rev. Lett. **74**, 4091 (1995) (Cited on pp. 9, 24).
- [76] A. Sørensen and K. Mølmer, "Quantum computation with ions in thermal motion", Phys. Rev. Lett. 82, 1971 (1999) (Cited on pp. 9, 26, 35).
- [77] K. Mølmer and A. Sørensen, "Multiparticle entanglement of hot trapped ions", Phys. Rev. Lett. 82, 1835 (1999) (Cited on pp. 9, 26, 35).
- [78] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, "Realization of the cirac-zoller controlled-not quantum gate", Nature 422, 408 (2003) (Cited on pp. 9, 24).
- [79] D. P. DiVincenzo, "The physical implementation of quantum computation", Fortschritte der Physik 48, 771 (2000) (Cited on pp. 9, 24).
- [80] J. E. Christensen, D. Hucul, W. C. Campbell, and E. R. Hudson, "High-fidelity manipulation of a qubit enabled by a manufactured nucleus", npj Quantum Information 6, 35 (2020) (Cited on p. 10).
- [81] A. Ransford, C. Roman, T. Dellaert, P. McMillin, and W. C. Campbell, "Weak dissipation for high-fidelity qubit-state preparation and measurement", Phys. Rev. A 104, L060402 (2021) (Cited on p. 10).
- [82] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, "Single ion qubit with estimated coherence time exceeding one hour", Nature Communications 12, 233 (2021) (Cited on p. 10).
- [83] A. D. Leu, M. F. Gely, M. A. Weber, M. C. Smith, D. P. Nadlinger, and D. M. Lucas, "Fast, high-fidelity addressed single-qubit gates using efficient composite pulse sequences", Phys. Rev. Lett. 131, 120601 (2023) (Cited on p. 10).
- [84] D. B. Hume, T. Rosenband, and D. J. Wineland, "High-fidelity adaptive qubit detection through repetitive quantum nondemolition measurements", Phys. Rev. Lett. 99, 120502 (2007) (Cited on pp. 10, 41, 48, 130, 137).
- [85] S. D. Erickson, J. J. Wu, P.-Y. Hou, D. C. Cole, S. Geller, A. Kwiatkowski, S. Glancy, E. Knill, D. H. Slichter, A. C. Wilson, and D. Leibfried, "High-fidelity indirect readout of trapped-ion hyperfine qubits", Phys. Rev. Lett. 128, 160503 (2022) (Cited on pp. 10, 130, 134).

- [86] L. M. Duan, B. B. Blinov, D. L. Moehring, and C. Monroe, "Scalable trapped ion quantum computation with a probabilistic ion-photon mapping", arXiv, 0401020 (2004) (Cited on p. 11).
- [87] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Entanglement of single-atom quantum bits at a distance", Nature 449, 68 (2007) (Cited on p. 11).
- [88] 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 (2020) (Cited on p. 11).
- [89] J. O'Reilly, G. Toh, I. Goetting, S. Saha, M. Shalaev, A. Carter, A. Risinger, A. Kalakuntla, T. Li, A. Verma, and C. Monroe, "Fast photon-mediated entanglement of continuously-cooled trapped ions for quantum networking", arXiv, 2404.16167 (2024) (Cited on p. 11).
- [90] M. G. Raizen, J. M. Gilligan, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Ionic crystals in a linear paul trap", Phys. Rev. A 45, 6493 (1992) (Cited on p. 11).
- [91] W. Paul, "Electromagnetic traps for charged and neutral particles", Rev. Mod. Phys. **62**, 531 (1990) (Cited on pp. 13, 14).
- [92] S. Earnshaw, "On the Nature of the Molecular Forces which regulate the Constitution of the Luminiferous Ether", Trans. Camb. Phil. Soc. 7, 97 (1842) (Cited on p. 13).
- [93] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions", Rev. Mod. Phys. 75, 281 (2003) (Cited on pp. 16, 22, 128).
- [94] H. Dehmelt, "Radiofrequency spectroscopy of stored ions i: storage**part ii: spectroscopy is now scheduled to appear in volume v of this series.", in , Vol. 3, edited by D. Bates and I. Estermann, Advances in Atomic and Molecular Physics (Academic Press, 1968), pp. 53–72 (Cited on p. 16).
- [95] P. K. Gosh, Ion Traps (Oxford University Press, 1995) (Cited on p. 16).
- [96] D. James, "Quantum dynamics of cold trapped ions with application to quantum computation", Applied Physics B: Lasers and Optics 66, 181 (1998) (Cited on p. 16).
- [97] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Minimization of ion micromotion in a Paul trap", Journal of Applied Physics 83, 5025 (1998) (Cited on pp. 17, 104).
- [98] A. Quinn, M. Brown, T. J. Gardner, and D. T. C. Allcock, "Geometries and fabrication methods for 3d printing ion traps", arXiv, 2205.15892 (2022) (Cited on p. 17).
- [99] S. Xu, X. Xia, Q. Yu, S. Khan, E. Megidish, B. You, B. Hemmerling, A. Jayich, J. Biener, and H. Häffner, "3d-printed micro ion trap technology for scalable quantum information processing", arXiv, 2310.00595 (2023) (Cited on p. 17).
- [100] R. Zhuravel, A. Sheffield, M. Straus, A. Menon, and D. Valdez, "A 3d monolithic trap for quantum simulation and computation", in Damop (Bulletin of the American Physical Society, 2024) (Cited on p. 17).

- [101] A. Lovera, E. Casamenti, C. Alfieri, E. Brucke, M. Wagener, S. Welte, P. Leindecker, I. Sergachev, M. Fontboté-Schmidt, J. Home, J. Schmidt, and C. Hempel, "3D femtosecond laser fabrication of glass monolithic ion traps for quantum computing", in Nanoscale and quantum materials: from synthesis and laser processing to applications 2024, Vol. PC12874, edited by A. V. Kabashin, M. Farsari, and M. Mahjouri-Samani (International Society for Optics and Photonics, 2024), PC1287404 (Cited on p. 17).
- [102] 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", arXiv, 0501147 (2005) (Cited on pp. 17, 18, 216).
- [103] 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) (Cited on pp. 17, 216).
- [104] I. H. Zimmermann, G. Zarantonello, J. J. Wu, P.-Y. Hou, A. D. Brandt, D. Leibfried, and D. H. Slichter, "Towards a cryogenic apparatus for beryllium, magnesium, calcium ion trap experiments", in Damop (Bulletin of the American Physical Society, 2024) (Cited on p. 18).
- [105] M. Malinowski, D. Allcock, and C. Ballance, "How to wire a 1000-qubit trapped-ion quantum computer", PRX Quantum 4, 040313 (2023) (Cited on p. 18).
- [106] S. Jain, T. Sägesser, P. Hrmo, C. Torkzaban, M. Stadler, R. Oswald, C. Axline, A. Bautista-Salvador, C. Ospelkaus, D. Kienzler, and J. Home, "Penning micro-trap for quantum computing", Nature 627, 510 (2024) (Cited on p. 18).
- [107] M. A. Weber, C. Löschnauer, J. Wolf, M. F. Gely, R. K. Hanley, J. F. Goodwin, C. J. Ballance, T. P. Harty, and D. M. Lucas, "Cryogenic ion trap system for high-fidelity near-field microwave-driven quantum logic", Quantum Science and Technology 9, 015007 (2023) (Cited on pp. 18, 35).
- [108] J. Stuart, R. Panock, C. Bruzewicz, J. Sedlacek, R. McConnell, I. Chuang, J. Sage, and J. Chiaverini, "Chip-integrated voltage sources for control of trapped ions", Phys. Rev. Appl. 11, 024010 (2019) (Cited on p. 18).
- [109] K. K. Mehta, "Integrated optical quantum manipulation and measurement of trapped ions", PhD thesis (MIT, 2017) (Cited on pp. 18, 29).
- [110] J. T. Merrill, C. Volin, D. Landgren, J. M. Amini, K. Wright, S. C. Doret, C.-S. Pai, H. Hayden, T. Killian, D. Faircloth, K. R. Brown, A. W. Harter, and R. E. Slusher, "Demonstration of integrated microscale optics in surface-electrode ion traps", New Journal of Physics 13, 103005 (2011) (Cited on p. 18).
- [111] A. P. VanDevender, Y. Colombe, J. Amini, D. Leibfried, and D. J. Wineland, "Efficient fiber optic detection of trapped ion fluorescence", Phys. Rev. Lett. 105, 023001 (2010) (Cited on p. 18).
- [112] P. F. Herskind, S. X. Wang, M. Shi, Y. Ge, M. Cetina, and I. L. Chuang, "Microfabricated surface ion trap on a high-finesse optical mirror", Opt. Lett. **36**, 3045 (2011) (Cited on p. 18).
- [113] C. R. Clark, C.-w. Chou, A. R. Ellis, J. Hunker, S. A. Kemme, P. Maunz, B. Tabakov, C. Tigges, and D. L. Stick, "Characterization of fluorescence collection optics integrated with a microfabricated surface electrode ion trap", Phys. Rev. Appl. 1, 024004 (2014) (Cited on p. 18).

- [114] M. Ghadimi, V. Blūms, B. G. Norton, P. M. Fisher, S. C. Connell, J. M. Amini, C. Volin, H. Hayden, C.-S. Pai, D. Kielpinski, M. Lobino, and E. W. Streed, "Scalable ion-photon quantum interface based on integrated diffractive mirrors", npj Quantum Information 3, 4 (2017) (Cited on p. 18).
- [115] A. Van Rynbach, P. Maunz, and J. Kim, "An integrated mirror and surface ion trap with a tunable trap location", Applied Physics Letters 109, 221108 (2016) (Cited on p. 18).
- [116] A. M. Eltony, S. X. Wang, G. M. Akselrod, P. F. Herskind, and I. L. Chuang, "Transparent ion trap with integrated photodetector", Applied Physics Letters **102**, 054106 (2013) (Cited on p. 18).
- [117] W. J. Setzer, M. Ivory, O. Slobodyan, J. W. Van Der Wall, L. P. Parazzoli, D. Stick, M. Gehl, M. G. Blain, R. R. Kay, and H. J. McGuinness, "Fluorescence detection of a trapped ion with a monolithically integrated single-photon-counting avalanche diode", Applied Physics Letters 119 (2021) (Cited on p. 18).
- [118] D. Reens, M. Collins, J. Ciampi, D. Kharas, B. F. Aull, K. Donlon, C. D. Bruzewicz, B. Felton, J. Stuart, R. J. Niffenegger, P. Rich, D. Braje, K. K. Ryu, J. Chiaverini, and R. McConnell, "High-fidelity ion state detection using trap-integrated avalanche photodiodes", Phys. Rev. Lett. 129, 100502 (2022) (Cited on p. 18).
- [119] C. W. Hogle, D. Dominguez, M. Dong, A. Leenheer, H. J. McGuinness, B. P. Ruzic, M. Eichenfield, and D. Stick, "High-fidelity trapped-ion qubit operations with scalable photonic modulators", npj Quantum Information 9, 74 (2023) (Cited on p. 18).
- [120] M. Akhtar, F. Bonus, F. R. Lebrun-Gallagher, N. I. Johnson, M. Siegele-Brown, S. Hong, S. J. Hile, S. A. Kulmiya, S. Weidt, and W. K. Hensinger, "A high-fidelity quantum matter-link between ion-trap microchip modules", Nature Communications 14, 531 (2023) (Cited on p. 18).
- [121] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt, "Trapped-ion probing of light-induced charging effects on dielectrics", New Journal of Physics 12, 093035 (2010) (Cited on pp. 18, 105).
- [122] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibrandt, K. R. Brown, and I. L. Chuang, "Suppression of heating rates in cryogenic surface-electrode ion traps", Phys. Rev. Lett. **100**, 013001 (2008) (Cited on pp. 18, 74).
- [123] J. Chiaverini and J. M. Sage, "Insensitivity of the rate of ion motional heating to trapelectrode material over a large temperature range", Phys. Rev. A 89, 012318 (2014) (Cited on p. 18).
- [124] P. L. W. Maunz, "High optical access trap 2.0.", Sandia Technical Note 1, 10.2172/1237003 (2016) (Cited on p. 19).
- [125] C. Decaroli, "Multi-wafer ion traps for scalable quantum information processing", PhD thesis (ETH Zurich, 2021) (Cited on p. 19).
- [126] S. Auchter, C. Axline, C. Decaroli, M. Valentini, L. Purwin, R. Oswald, R. Matt, E. Aschauer, Y. Colombe, P. Holz, T. Monz, R. Blatt, P. Schindler, C. Rössler, and J. Home, "Industrially microfabricated ion trap with 1 ev trap depth", Quantum Science and Technology 7, 035015 (2022) (Cited on p. 19).

- [127] D. R. Leibrandt, R. J. Clark, J. Labaziewicz, P. Antohi, W. Bakr, K. R. Brown, and I. L. Chuang, "Laser ablation loading of a surface-electrode ion trap", Phys. Rev. A **76**, 055403 (2007) (Cited on p. 19).
- [128] X. Shi, S. L. Todaro, G. L. Mintzer, C. D. Bruzewicz, J. Chiaverini, and I. L. Chuang, "Ablation loading of barium ions into a surface-electrode trap", Applied Physics Letters 122 (2023) (Cited on p. 19).
- [129] 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 (2012) (Cited on pp. 20, 245).
- [130] J. D. Sterk, M. G. Blain, M. Delaney, R. Haltli, E. Heller, A. L. Holterhoff, T. Jennings, N. Jimenez, A. Kozhanov, Z. Meinelt, E. Ou, J. V. D. Wall, C. Noel, and D. Stick, "Multijunction surface ion trap for quantum computing", arXiv, 2403.00208 (2024) (Cited on p. 20).
- S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, [131]B. Bjork, P. Blanchard, M. Bohn, J. G. Bohnet, N. C. Brown, N. Q. Burdick, W. C. Burton, S. L. Campbell, J. P. Campora, C. Carron, J. Chambers, J. W. Chan, Y. H. Chen, A. Chernoguzov, E. Chertkov, J. Colina, J. P. Curtis, R. Daniel, M. DeCross, D. Deen, C. Delaney, J. M. Dreiling, C. T. Ertsgaard, J. Esposito, B. Estey, M. Fabrikant, C. Figgatt, C. Foltz, M. Foss-Feig, D. Francois, J. P. Gaebler, T. M. Gatterman, C. N. Gilbreth, J. Giles, E. Glynn, A. Hall, A. M. Hankin, A. Hansen, D. Hayes, B. Higashi, I. M. Hoffman, B. Horning, J. J. Hout, R. Jacobs, J. Johansen, L. Jones, J. Karcz, T. Klein, P. Lauria, P. Lee, D. Liefer, S. T. Lu, D. Lucchetti, C. Lytle, A. Malm, M. Matheny, B. Mathewson, K. Mayer, D. B. Miller, M. Mills, B. Neyenhuis, L. Nugent, S. Olson, J. Parks, G. N. Price, Z. Price, M. Pugh, A. Ransford, A. P. Reed, C. Roman, M. Rowe, C. Ryan-Anderson, S. Sanders, J. Sedlacek, P. Shevchuk, P. Siegfried, T. Skripka, B. Spaun, R. T. Sprenkle, R. P. Stutz, M. Swallows, R. I. Tobey, A. Tran, T. Tran, E. Vogt, C. Volin, J. Walker, A. M. Zolot, and J. M. Pino, "A race-track trapped-ion quantum processor", Phys. Rev. X 13, 041052 (2023) (Cited on p. 20).
- [132] M. A. Rowe, A. Ben-Kish, B. DeMarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, and D. J. Wineland, "Transport of quantum states and separation of ions in a dual rf ion trap", arXiv, 0205094 (2002) (Cited on p. 20).
- [133] 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 (2009) (Cited on p. 20).
- [134] W. C. Burton, B. Estey, I. M. Hoffman, A. R. Perry, C. Volin, and G. Price, "Transport of multispecies ion crystals through a junction in a radio-frequency paul trap", Phys. Rev. Lett. 130, 173202 (2023) (Cited on p. 20).
- [135] D. Palani, F. Hasse, P. Kiefer, F. Boeckling, J.-P. Schroeder, U. Warring, and T. Schaetz, "High-fidelity transport of trapped-ion qubits in a multilayer array", Phys. Rev. A 107, L050601 (2023) (Cited on p. 20).
- [136] 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 (2012) (Cited on pp. 20, 237).

- [137] S. L. Todaro, "Improved state detection and transport of trapped ion qubits for scalable quantum computing", PhD thesis (University of Colorado at Boulder, 2020) (Cited on pp. 20, 81, 216–219, 227, 228).
- [138] H. HAFFNER, C. ROOS, and R. BLATT, "Quantum computing with trapped ions", Physics Reports 469, 155 (2008) (Cited on pp. 22, 24).
- [139] S. C. Burd, "Squeezing and amplification of trapped-ion motion", PhD thesis (University of Colorado at Boulder, 2020) (Cited on pp. 23, 36, 52, 74, 75, 88, 89, 93, 97, 108, 167).
- [140] C. Monroe, D. M. Meekhof, B. E. King, and D. J. Wineland, "A "schrödinger cat" superposition state of an atom", Science **272**, 1131 (1996) (Cited on p. **23**).
- [141] K. C. McCormick, J. Keller, S. C. Burd, D. J. Wineland, A. C. Wilson, and D. Leibfried, "Quantum-enhanced sensing of a single-ion mechanical oscillator", Nature **572**, 86 (2019) (Cited on p. 23).
- [142] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Laser cooling to the zero-point energy of motion", Phys. Rev. Lett. **62**, 403 (1989) (Cited on p. **23**).
- [143] 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 (1995) (Cited on p. 23).
- [144] C. Fang, Y. Wang, K. Sun, and J. Kim, "Realization of scalable cirac-zoller multi-qubit gates", arXiv, 2301.07564 (2023) (Cited on p. 24).
- [145] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, "Fast quantum logic gates with trapped-ion qubits", Nature 555, 75 (2018) (Cited on pp. 26, 27).
- [146] A. Sørensen and K. Mølmer, "Entanglement and quantum computation with ions in thermal motion", Phys. Rev. A **62**, 022311 (2000) (Cited on p. 26).
- [147] G. Milburn, S. Schneider, and D. James, "Ion trap quantum computing with warm ions", Fortschritte der Physik 48, 801 (2000) (Cited on p. 27).
- [148] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland, "Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate", Nature 422, 412 (2003) (Cited on p. 27).
- [149] R. T. Sutherland, R. Srinivas, S. C. Burd, D. Leibfried, A. C. Wilson, D. J. Wineland, D. T. C. Allcock, D. H. Slichter, and S. B. Libby, "Versatile laser-free trapped-ion entangling gates", New Journal of Physics 21, 033033 (2019) (Cited on pp. 27, 30, 39, 40, 137, 152, 153, 155).
- [150] I. D. Moore, W. C. Campbell, E. R. Hudson, M. J. Boguslawski, D. J. Wineland, and D. T. C. Allcock, "Photon scattering errors during stimulated raman transitions in trapped-ion qubits", Phys. Rev. A 107, 032413 (2023) (Cited on p. 28).
- [151] J. Kwon, W. J. Setzer, M. Gehl, N. Karl, J. Van Der Wall, R. Law, M. G. Blain, D. Stick, and H. J. McGuinness, "Multi-site integrated optical addressing of trapped ions", Nature Communications 15 (2024) (Cited on p. 29).
- [152] F. Mintert and C. Wunderlich, "Ion-trap quantum logic using long-wavelength radiation", Phys. Rev. Lett. 87, 257904 (2001) (Cited on pp. 29, 33, 37).

- [153] D. Leibfried, E. Knill, C. Ospelkaus, and D. J. Wineland, "Transport quantum logic gates for trapped ions", Phys. Rev. A **76**, 032324 (2007) (Cited on p. 29).
- [154] C. Ospelkaus, C. E. Langer, J. M. Amini, K. R. Brown, D. Leibfried, and D. J. Wineland, "Trapped-ion quantum logic gates based on oscillating magnetic fields", Phys. Rev. Lett. **101**, 090502 (2008) (Cited on pp. 30, 33–35, 37).
- [155] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas, "High-fidelity trapped-ion quantum logic using near-field microwaves", Phys. Rev. Lett. 117, 140501 (2016) (Cited on pp. 30, 38).
- [156] H. Hahn, G. Zarantonello, M. Schulte, A. Bautista-Salvador, K. Hammerer, and C. Ospelkaus, "Integrated 9be+ multi-qubit gate device for the ion-trap quantum computer", npj Quantum Information 5, 70 (2019) (Cited on pp. 30, 35, 38).
- [157] G. Zarantonello, H. Hahn, J. Morgner, M. Schulte, A. Bautista-Salvador, R. F. Werner, K. Hammerer, and C. Ospelkaus, "Robust and resource-efficient microwave near-field entangling ⁹Be⁺ gate", Phys. Rev. Lett. **123**, 260503 (2019) (Cited on pp. 30, 39).
- [158] M. Nünnerich, D. Cohen, P. Barthel, P. H. Huber, D. Niroomand, A. Retzker, and C. Wunderlich, "Fast, robust and laser-free universal entangling gates for trapped-ion quantum computing", arXiv, 2403.04730 (2024) (Cited on pp. 30, 38).
- [159] D. J. Griffiths, Introduction to Electrodynamics, 4th ed. (Pearson, 2012) (Cited on p. 31).
- [160] J. Reichel, W. Hänsel, and T. W. Hänsch, "Atomic micromanipulation with magnetic surface traps", Phys. Rev. Lett. 83, 3398 (1999) (Cited on p. 33).
- [161] M. A. Sepiol, "A high-fidelity microwave driven two-qubit logic gate in ⁴³Ca⁺", PhD thesis (University of Oxford, 2016) (Cited on pp. 33, 34, 61, 74).
- [162] H. Hahn, G. Zarantonello, A. Bautista-Salvador, M. Wahnschaffe, M. Kohnen, J. Schoebel, P. O. Schmidt, and C. Ospelkaus, "Multilayer ion trap with three-dimensional microwave circuitry for scalable quantum logic applications", Applied Physics B 125 (2019) (Cited on p. 35).
- [163] R. Srinivas, "Laser-free trapped-ion quantum logic with a radiofrequency magnetic field gradient", PhD thesis (University of Colorado at Boulder, 2020) (Cited on pp. 36, 37, 40, 52, 54, 57, 65, 74, 75, 83, 85, 88, 93, 97, 101, 108, 111, 145, 149, 152, 199, 200).
- [164] S. Weidt, J. Randall, S. C. Webster, E. D. Standing, A. Rodriguez, A. E. Webb, B. Lekitsch, and W. K. Hensinger, "Ground-state cooling of a trapped ion using long-wavelength radiation", Phys. Rev. Lett. 115, 013002 (2015) (Cited on p. 36).
- [165] T. Sriarunothai, G. S. Giri, S. Wölk, and C. Wunderlich, "Radio frequency sideband cooling and sympathetic cooling of trapped ions in a static magnetic field gradient", Journal of Modern Optics 65, 560 (2017) (Cited on p. 36).
- [166] R. Srinivas, S. C. Burd, R. T. Sutherland, A. C. Wilson, D. J. Wineland, D. Leibfried, D. T. C. Allcock, and D. H. Slichter, "Trapped-ion spin-motion coupling with microwaves and a near-motional oscillating magnetic field gradient", Phys. Rev. Lett. 122, 163201 (2019) (Cited on pp. 36, 145).
- [167] M. Johanning, A. Braun, N. Timoney, V. Elman, W. Neuhauser, and C. Wunderlich, "Individual addressing of trapped ions and coupling of motional and spin states using rf radiation", Phys. Rev. Lett. 102, 073004 (2009) (Cited on p. 37).

- [168] D. P. L. Aude Craik, N. M. Linke, T. P. Harty, C. J. Ballance, D. M. Lucas, A. M. Steane, and D. T. C. Allcock, "Microwave control electrodes for scalable, parallel, single-qubit operations in a surface-electrode ion trap", Applied Physics B 114, 3 (2014) (Cited on p. 37).
- [169] A. Khromova, C. Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and C. Wunderlich, "Designer spin pseudomolecule implemented with trapped ions in a magnetic gradient", Phys. Rev. Lett. 108, 220502 (2012) (Cited on p. 38).
- [170] D. T. C. Allcock, W. C. Campbell, J. Chiaverini, I. L. Chuang, E. R. Hudson, I. D. Moore, A. Ransford, C. Roman, J. M. Sage, and D. J. Wineland, "omg blueprint for trapped ion quantum computing with metastable states", Applied Physics Letters 119, 214002 (2021) (Cited on pp. 41, 47).
- [171] N. Chen, L. Li, W. Huie, M. Zhao, I. Vetter, C. H. Greene, and J. P. Covey, "Analyzing the rydberg-based optical-metastable-ground architecture for ¹⁷¹Yb nuclear spins", Phys. Rev. A **105**, 052438 (2022) (Cited on p. 41).
- [172] J. W. Lis, A. Senoo, W. F. McGrew, F. Rönchen, A. Jenkins, and A. M. Kaufman, "Midcircuit operations using the omg architecture in neutral atom arrays", Phys. Rev. X 13, 041035 (2023) (Cited on p. 41).
- [173] O. Băzăvan, S. Saner, M. Minder, A. C. Hughes, R. T. Sutherland, D. M. Lucas, R. Srinivas, and C. J. Ballance, "Synthesizing a $\hat{\sigma}_z$ spin-dependent force for optical, metastable, and ground-state trapped-ion qubits", Phys. Rev. A 107, 022617 (2023) (Cited on p. 41).
- [174] L. Feng, Y.-Y. Huang, Y.-K. Wu, W.-X. Guo, J.-Y. Ma, H.-X. Yang, L. Zhang, Y. Wang, C.-X. Huang, C. Zhang, L. Yao, B.-X. Qi, Y.-F. Pu, Z.-C. Zhou, and L.-M. Duan, "Realization of a crosstalk-avoided quantum network node using dual-type qubits of the same ion species", Nature Communications 15, 204 (2024) (Cited on p. 41).
- [175] D. J. Larson, J. C. Bergquist, J. J. Bollinger, W. M. Itano, and D. J. Wineland, "Sympathetic cooling of trapped ions: a laser-cooled two-species nonneutral ion plasma", Phys. Rev. Lett. 57, 70 (1986) (Cited on p. 46).
- [176] H. Rohde, S. T. Gulde, C. F. Roos, P. A. Barton, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt, "Sympathetic ground-state cooling and coherent manipulation with two-ion crystals", Journal of Optics B: Quantum and Semiclassical Optics 3, S34 (2001) (Cited on p. 46).
- [177] M. D. Barrett, B. DeMarco, T. Schaetz, V. Meyer, D. Leibfried, J. Britton, J. Chiaverini, W. M. Itano, B. Jelenkovi ć, J. D. Jost, C. Langer, T. Rosenband, and D. J. Wineland, "Sympathetic cooling of ⁹Be⁺ and ²⁴Mg⁺ for quantum logic", Phys. Rev. A 68, 042302 (2003) (Cited on p. 46).
- [178] R. Rugango, J. E. Goeders, T. H. Dixon, J. M. Gray, N. B. Khanyile, G. Shu, R. J. Clark, and K. R. Brown, "Sympathetic cooling of molecular ion motion to the ground state", New Journal of Physics 17, 035009 (2015) (Cited on p. 46).
- [179] L. Schmöger, O. O. Versolato, M. Schwarz, M. Kohnen, A. Windberger, B. Piest, S. Feucht-enbeiner, J. Pedregosa-Gutierrez, T. Leopold, P. Micke, A. K. Hansen, T. M. Baumann, M. Drewsen, J. Ullrich, P. O. Schmidt, and J. R. C. López-Urrutia, "Coulomb crystallization of highly charged ions", Science 347, 1233 (2015) (Cited on p. 46).

- [180] S. D. Fallek, V. S. Sandhu, R. A. McGill, J. M. Gray, H. N. Tinkey, C. R. Clark, and K. R. Brown, "Rapid exchange cooling with trapped ions", Nature Communications 15, 1089 (2024) (Cited on p. 47).
- [181] D. Kienzler, Y. Wan, S. D. Erickson, J. J. Wu, A. C. Wilson, D. J. Wineland, and D. Leibfried, "Quantum logic spectroscopy with ions in thermal motion", Phys. Rev. X 10, 021012 (2020) (Cited on pp. 50, 137).
- [182] C. J. Ballance, V. M. Schäfer, J. P. Home, D. J. Szwer, S. C. Webster, D. T. C. Allcock, N. M. Linke, T. P. Harty, D. P. L. Aude Craik, D. N. Stacey, A. M. Steane, and D. M. Lucas, "Hybrid quantum logic and a test of bell's inequality using two different atomic isotopes", Nature 528, 384 (2015) (Cited on pp. 50, 152).
- [183] T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland, "Multi-element logic gates for trapped-ion qubits", Nature 528, 380 (2015) (Cited on pp. 50, 152, 155).
- [184] A. C. Hughes, V. M. Schäfer, K. Thirumalai, D. P. Nadlinger, S. R. Woodrow, D. M. Lucas, and C. J. Ballance, "Benchmarking a high-fidelity mixed-species entangling gate", Phys. Rev. Lett. **125**, 080504 (2020) (Cited on pp. 50, 152).
- [185] T. R. Tan, "High-fidelity entangling gates with trapped-ions", PhD thesis (University of Colorado at Boulder, 2009) (Cited on pp. 51, 83).
- [186] D. N. Madsen, S. Balslev, M. Drewsen, N. Kjærgaard, Z. Videsen, and J. W. Thomsen, "Measurements on photo-ionization of 3s3p¹P₁ magnesium atoms", J. Phys. B: At. Mol. Opt. Phys. 33, 4981 (2000) (Cited on pp. 54, 127).
- [187] D. P. Nadlinger, P. Drmota, D. Main, B. C. Nichol, G. Araneda, R. Srinivas, L. J. Stephenson, C. J. Ballance, and D. M. Lucas, "Micromotion minimisation by synchronous detection of parametrically excited motion", arXiv, 2107.00056 (2021) (Cited on pp. 54, 98).
- [188] A. Kwiatkowski, "Optimized experiment design and analysis for fully randomized benchmarking", PhD thesis (University of Colorado at Boulder, 2024) (Cited on pp. 59, 83).
- [189] D. T. Allock, "Surface-electrode ion traps for scalable quantum computing", PhD thesis (University of Oxford, 2011) (Cited on p. 61).
- [190] D. Kienzler, "Quantum harmonic oscillator state synthesis by reservoir engineering", PhD thesis (ETH Zurich, 2015) (Cited on pp. 61, 68).
- [191] S. Gulde, D. Rotter, P. Barton, F. Schmidt-Kaler, R. Blatt, and W. Hogervorst, "Simple and efficient photo-ionization loading of ions for precision ion-trapping experiments", Applied Physics B 73, 861 (2001) (Cited on p. 63).
- [192] 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 (2018) (Cited on pp. 76, 193).
- [193] T. G. Ballance, J. F. Goodwin, B. Nichol, L. J. Stephenson, C. J. Ballance, and D. M. Lucas, "A short response time atomic source for trapped ion experiments", Review of Scientific Instruments 89, 053102 (2018) (Cited on p. 78).
- [194] W. Gifford, "The Gifford-McMahon Cycle", in Advances in cryogenic engineering, Vol. 11, edited by K. D. Timmerhaus (Engineering Research Center, 1966), pp. 152–159 (Cited on p. 81).

- [195] M. D. Atrey, Cryocoolers: Theory and Applications (Springer, 2020) (Cited on p. 81).
- [196] Stinger Cooling Engine Owner's Manual Rev. 03 (ColdEdge Technologies, 2019) (Cited on p. 81).
- [197] B. Merkel, K. Thirumalai, J. E. Tarlton, V. M. Schäfer, C. J. Ballance, T. P. Harty, and D. M. Lucas, "Magnetic field stabilization system for atomic physics experiments", Review of Scientific Instruments 90, 044702 (2019) (Cited on p. 83).
- [198] P. Huang and D. Leibfried, "Achromatic catadioptric microscope objective in deep ultraviolet with long working distance", in Novel optical systems design and optimization vii, Vol. 5524, edited by J. M. Sasian, R. J. Koshel, P. K. Manhart, and R. C. Juergens (International Society for Optics and Photonics, 2004), pp. 125–133 (Cited on p. 83).
- [199] A. C. Wilson, C. Ospelkaus, A. P. VanDevender, J. A. Mlynek, K. R. Brown, D. Leibfried, and D. J. Wineland, "A 750-mW, continuous-wave, solid-state laser source at 313 nm for cooling and manipulating trapped ⁹Be⁺ ions", English, Appl. Phys. B **105**, 741 (2011) (Cited on pp. 85, 88, 243).
- [200] N. Torelli, "Characterization of uv optical fibers analysis of the effect of hydrogen-loading and uv-curing on lma-10-uv photonic crystal fibers", PhD thesis (ETH Zurich, 2018) (Cited on p. 89).
- [201] E. R. Clements, "High-stability differential frequency comparisons of optical atomic clocks with correlated laser noise", PhD thesis (University of Colorado at Boulder, 2022) (Cited on p. 89).
- [202] M. Chwalla, J. Benhelm, K. Kim, G. Kirchmair, T. Monz, M. Riebe, P. Schindler, A. S. Villar, W. Hänsel, C. F. Roos, R. Blatt, M. Abgrall, G. Santarelli, G. D. Rovera, and P. Laurent, "Absolute frequency measurement of the 40 Ca $^+$ 4s $^2S_{1/2} 3d$ $^2D_{5/2}$ clock transition", Phys. Rev. Lett. 102, 023002 (2009) (Cited on p. 92).
- [203] M-Labs Limited, *PDQ documentation*, (2017) pdq.readthedocs.io (visited on 05/30/2024) (Cited on p. 98).
- [204] G. Kasprowicz, P. Kulik, M. Gaska, T. Przywozki, K. Pozniak, J. Jarosinski, J. W. Britton, T. Harty, C. Balance, W. Zhang, D. Nadlinger, D. Slichter, D. Allcock, S. Bourdeauducq, R. Jördens, and K. Pozniak, "Artiq and sinara: open software and hardware stacks for quantum physics", in Osa quantum 2.0 conference (2020), QTu8B.14 (Cited on p. 101).
- [205] D. Nadlinger, P. Drmota, T. Harty, C. Ballance, AVB25, J. Leppard, C. Baynham, M. Birtwell, H. McLaughlin, vmsch, and L. Stephenson, *Ndscan*, https://github.com/OxfordIonTrapGroup/ndscan, 2024 (Cited on p. 103).
- [206] D. T. C. Allcock, T. P. Harty, H. A. Janacek, N. M. Linke, C. J. Ballance, A. M. Steane, D. M. Lucas, R. L. Jarecki, S. D. Habermehl, M. G. Blain, D. Stick, and D. L. Moehring, "Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap", Applied Physics B 107, 913 (2011) (Cited on p. 105).
- [207] S. X. Wang, G. Hao Low, N. S. Lachenmyer, Y. Ge, P. F. Herskind, and I. L. Chuang, "Laser-induced charging of microfabricated ion traps", Journal of Applied Physics 110 (2011) (Cited on p. 105).
- [208] C. Jung, J. Jeong, S. Yoo, T. Kim, and D. "Cho, "Method for estimating locations of laser-induced stray charges on surface-electrode ion traps using secular frequency shift at multiple ion positions", Phys. Rev. Appl. 20, 014032 (2023) (Cited on p. 105).

- [209] U. Warring, C. Ospelkaus, Y. Colombe, K. R. Brown, J. M. Amini, M. Carsjens, D. Leibfried, and D. J. Wineland, "Techniques for microwave near-field quantum control of trapped ions", Phys. Rev. A 87, 013437 (2013) (Cited on pp. 108, 116).
- [210] Y. Colombe, D. H. Slichter, A. C. Wilson, D. Leibfried, and D. J. Wineland, "Single-mode optical fiber for high-power, low-loss uv transmission", Opt. Express 22, 19783 (2014) (Cited on p. 108).
- [211] H. Cummins and J. Jones, "Resonance offset tailored composite pulses", Journal of Magnetic Resonance 148, 338 (2001) (Cited on p. 134).
- [212] D. J. Wineland and H. G. Dehmelt, "Principles of the stored ion calorimeter", Journal of Applied Physics 46, 919 (1975) (Cited on pp. 138, 155).
- [213] J. Metzner, A. Quinn, S. Brudney, I. D. Moore, S. C. Burd, D. J. Wineland, and D. T. C. Allcock, "Two-mode squeezing and su(1,1) interferometry with trapped ions", arXiv, 2312.10847 (2024) (Cited on p. 138).
- [214] P.-Y. Hou, J. J. Wu, S. D. Erickson, D. C. Cole, G. Zarantonello, A. D. Brandt, A. C. Wilson, D. H. Slichter, and D. Leibfried, "Coherently coupled mechanical oscillators in the quantum regime", arXiv, 2205.14841 (2022) (Cited on p. 138).
- [215] S. C. Burd, R. Srinivas, J. J. Bollinger, A. C. Wilson, D. J. Wineland, D. Leibfried, D. H. Slichter, and D. T. C. Allcock, "Quantum amplification of mechanical oscillator motion", Science **364**, 1163 (2019) (Cited on p. 142).
- [216] S. C. Burd, R. Srinivas, H. M. Knaack, W. Ge, A. C. Wilson, D. J. Wineland, D. Leibfried, J. J. Bollinger, D. T. C. Allcock, and D. H. Slichter, "Quantum amplification of boson-mediated interactions", Nature Physics 17, 898 (2021) (Cited on p. 142).
- [217] S. C. Burd, H. M. Knaack, R. Srinivas, C. Arenz, A. L. Collopy, L. J. Stephenson, A. C. Wilson, D. J. Wineland, D. Leibfried, J. J. Bollinger, D. T. C. Allcock, and D. H. Slichter, "Experimental speedup of quantum dynamics through squeezing", PRX Quantum 5 (2024) (Cited on p. 142).
- [218] M. I. Fabrikant, P. Lauria, I. S. Madjarov, W. C. Burton, and R. T. Sutherland, "Cooling trapped ions with phonon rapid adiabatic passage", arXiv, 2403.02315 (2024) (Cited on p. 164).
- [219] K. R. Brown, C. Ospelkaus, Y. Colombe, A. C. Wilson, D. Leibfried, and D. J. Wineland, "Coupled quantized mechanical oscillators", Nature **471**, 196 (2011) (Cited on pp. 167, 237).
- [220] K. J. Weatherill, J. D. Pritchard, P. F. Griffin, U. Dammalapati, C. S. Adams, and E. Riis, "A versatile and reliably reusable ultrahigh vacuum viewport", Review of Scientific Instruments 80, 026105 (2009) (Cited on pp. 196, 197).
- [221] S. G. Cox, P. F. Griffin, C. S. Adams, D. DeMille, and E. Riis, "Reusable ultrahigh vacuum viewport bakeable to 240 c", Review of Scientific Instruments 74, 3185 (2003) (Cited on p. 197).
- [222] NASA, Outgassing data for selecting spacecraft materials, outgassing.nasa.gov (visited on 05/30/2024) (Cited on p. 197).
- [223] G. L. Weissler (1990) (Cited on p. 216).
- [224] S. D. S. Gordon and A. Osterwalder, "3d-printed beam splitter for polar neutral molecules", Phys. Rev. Appl. 7, 044022 (2017) (Cited on p. 220).

- [225] "Cm100- quartz crystal microbalance theory and calibration", (Cited on pp. 229, 230).
- [226] G. Sauerbrey, "The use of quartz oscillators for weighing thin layers and for microweighing", Zeitschrift für Physik **155**, 206 (1959) (Cited on p. 230).
- [227] S. C. Burd, J.-P. Penttinen, P.-Y. Hou, H. M. Knaack, S. Ranta, M. Mäki, E. Kantola, M. Guina, D. H. Slichter, D. Leibfried, and A. C. Wilson, "Vecsel systems for quantum information processing with trapped beryllium ions", J. Opt. Soc. Am. B 40, 773 (2023) (Cited on pp. 237, 242).
- [228] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: progress and challenges", Appl. Phys. Rev. 6, 021314 (2019) (Cited on p. 237).
- [229] 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 (2012) (Cited on p. 237).
- [230] A. C. Wilson, Y. Colombe, K. R. Brown, E. Knill, D. Leibfried, and D. J. Wineland, "Tunable spin-spin interactions and entanglement of ions in separate potential wells", Nature 512, 57 (2014) (Cited on p. 237).
- [231] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland, "Errors in trapped-ion quantum gates due to spontaneous photon scattering", Phys. Rev. A 75, 042329 (2007) (Cited on pp. 238, 245).
- [232] A. M. Steane, "How to build a 300 bit, 1 giga-operation quantum computer", Quantum Inf. Comput. 7, 171 (2004) (Cited on p. 238).
- [233] C. Monroe and J. Kim, "Scaling the ion trap quantum processor", Science **339**, 1164 (2013) (Cited on p. 238).
- [234] K. K. Mehta, C. D. Bruzewicz, R. McConnell, R. J. Ram, J. M. Sage, and J. Chiaverini, "Integrated optical addressing of an ion qubit", Nat. Nanotechnol. **11**, 1066 (2016) (Cited on p. 238).
- [235] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland, "Long-lived qubit memory using atomic ions", Phys. Rev. Lett. 95, 060502 (2005) (Cited on p. 238).
- [236] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland, "Single-qubit-gate error below 10^{-4} in a trapped ion", Phys. Rev. A **84**, 030303 (2011) (Cited on p. 238).
- [237] H.-Y. Lo, J. Alonso, D. Kienzler, B. C. Keitch, L. E. de Clercq, V. Negnevitsky, and J. P. Home, "All-solid-state continuous-wave laser systems for ionization, cooling and quantum state manipulation of beryllium ions", Appl. Phys. B 114, 17 (2014) (Cited on pp. 238, 243).
- [238] M. Kuznetsov, F. Hakimi, R. Sprague, and A. Mooradian, "High-power (> 0.5-w cw) diode-pumped vertical-external-cavity surface-emitting semiconductor lasers with circular TEM_{00} beams", IEEE Photon. Technol. Lett. **9**, 1063 (1997) (Cited on pp. 238, 242, 245).
- [239] M. Guina, A. Rantamäki, and A. Härkönen, "Optically pumped VECSELs: review of technology and progress", J. Phys. D **50**, 383001 (2017) (Cited on pp. 238, 242).
- [240] M. Myara, M. Sellahi, A. Laurain, A. Michon, I. Sagnes, and A. Garnache, "Noise properties of NIR and MIR VECSELs", Proc. of SPIE 8606, 86060Q (2013) (Cited on p. 238).

- [241] S. C. Burd, D. T. Allcock, T. Leinonen, J.-P. Penttinen, D. H. Slichter, R. Srinivas, A. C. Wilson, R. Jördens, M. Guina, D. Leibfried, and D. J. Wineland, "Vecsel systems for the generation and manipulation of trapped magnesium ions", Optica 3, 1294 (2016) (Cited on p. 240).
- [242] P. H. Moriya, R. Casula, G. A. Chappell, D. C. Parrotta, S. Ranta, H. Kahle, M. Guina, and J. E. Hastie, "InGaN-diode-pumped AlGaInP VECSEL with sub-kHz linewidth at 689 nm", Opt. Express 29, 3258 (2021) (Cited on p. 240).
- [243] J. N. Tinsley, S. Bandarupally, J.-P. Penttinen, S. Manzoor, S. Ranta, L. Salvi, M. Guina, and N. Poli, "Watt-level blue light for precision spectroscopy, laser cooling and trapping of strontium and cadmium atoms", Optics Express 29, 25462 (2021) (Cited on p. 240).
- [244] J. C. Hill, W. K. Holland, P. D. Kunz, K. C. Cox, J.-P. Penttinen, E. Kantola, and D. H. Meyer, "Intra-Cavity Frequency-Doubled VECSEL System for Narrow Linewidth Rydberg EIT Spectroscopy", arXiv, 2206.00096 (2022) (Cited on p. 240).
- [245] J. R. C. Woods, H. Kahle, A. C. Gray, J. Daykin, A. C. Tropper, C. Gawith, M. Guina, and V. Apostolopoulos, "High power 739 nm VECSELs for future Yb⁺ ion cooling", Appl. Opt. 60, 676 (2021) (Cited on p. 240).
- [246] A. Laurain, J. Hader, and J. V. Moloney, "Modeling and optimization of transverse modes in vertical-external-cavity surface-emitting lasers", J. Opt. Soc. Am. B 36, 847 (2019) (Cited on p. 241).
- [247] A. Tropper and S. Hoogland, "Extended cavity surface-emitting semiconductor lasers", Prog. Quantum. Electron. **30**, 1 (2006) (Cited on p. 242).
- [248] A. Castrillo, E. Fasci, G. Galzerano, G. Casa, P. Laporta, and L. Gianfrani, "Offset-frequency locking of extended-cavity diode lasers for precision spectroscopy of water at 1.38 μ m", Opt. Express 18, 21851 (2010) (Cited on p. 242).
- [249] R. Drever, J. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical resonator", Appl. Phys. B 31, 97 (1983) (Cited on p. 243).
- [250] R. Le Targat, J.-J. Zondy, and P. Lemonde, "75%-efficiency blue generation from an intracavity ppktp frequency doubler", Opt. Commun. **247**, 471 (2005) (Cited on p. **243**).
- [251] T. Hansch and B. Couillaud, "Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity", Optics communications **35**, 441 (1980) (Cited on p. **243**).
- [252] E. C. Cook, A. D. Vira, C. Patterson, E. Livernois, and W. D. Williams, "Testing quantum electrodynamics in the lowest singlet state of neutral beryllium-9", Phys. Rev. Lett. 121, 053001 (2018) (Cited on p. 244).
- [253] W. Demtröder, Spectroscopic instrumentation (Springer, 1981) (Cited on p. 244).
- [254] R. A. Carollo, D. A. Lane, E. K. Kleiner, P. A. Kyaw, C. C. Teng, C. Y. Ou, S. Qiao, and D. Hanneke, "Third-harmonic-generation of a diode laser for quantum control of beryllium ions", Opt. Express 25, 7220 (2017) (Cited on p. 245).
- [255] E. Kantola, T. Leinonen, S. Ranta, M. Tavast, and M. Guina, "High-efficiency 20 W yellow VECSEL", Opt. Express 22, 6372 (2014) (Cited on p. 245).



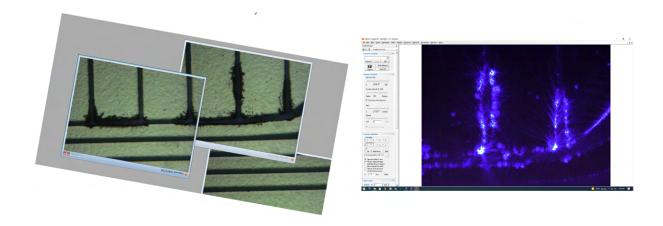
Rebuild

In January 2022, we discovered that our trap rf connection had cracked (shown in Fig. A.4), likely due to the extreme temperature cycling connected with our cryostat changeover (see Appendix B). The connection had been initially made with H21D, an electrically conductive epoxy, but one which is not rated for heavy thermal cycling. We took the opportunity to make several upgrades to the trap, pillbox, and vacuum chamber to improve experiment functionality and reduce the risk of future failures.

A.1 Trap

Opening up the vacuum chamber gave us the opportunity to replace our trap chip with a new trap of the same design. The old trap had come from a fabrication run which had some lumpiness in the gold plating which forms the trap electrodes, which created scatter spots on the trap surface. One notable scatter spot (pictured in Fig. A.2) looked exactly like an ion on the camera and was located just below and to the side of the load zone. This was actually convenient for aligning lasers and imaging systems, but ideally we would not have more background scatter than necessary, so we wanted to avoid having scatter spots in the trapping zone on a new trap. Another, more important

Figure A.1: White light- and focused laser pointer-illuminated images of the damaged area on the side of the old trap chip opposite to the loading zone. This damage was not present when the trap was originally installed, and its origins are unknown. We would suspect rf breakdown (arcing between electrodes that carry high trapping rf voltages) except that none of the involved electrodes carry our trap rf. This region made transport to this portion of the trap difficult and imaging of ions in that region impossible due to the amount of light scattered.



reason to replace the trap chip was the region of damage on the righthand portion of the trap. The origin of this damage is not known - it was not present when the trap was first placed under vacuum and its date of origin in unknown, but it made the right-hand portion of the trap unusable because the background scatter off of this damaged area was much higher than any ion fluorescence. Our mixed species loading and transport experiments ended up utilizing the majority of our trap length and so the work described in the main text would likely have been even more difficult or even impossible if we had not replaced the trap.

Figure A.2: The same scatter spot located in the load zone of the old trap chip, imaged in three ways: left, with our trap imaging camera, illuminated using a 397 laser; center, using a microscope-mounted camera, illuminated with white light; and right, using the same camera and microscope as the center photo but illuminated with a focused 405 nm laser pointer. This test was done to verify that the glancing laser pointer and microscope camera imaged potential scatter spots similarly to what we see on our trap imaging camera.



A.1.1 Scatter spot testing

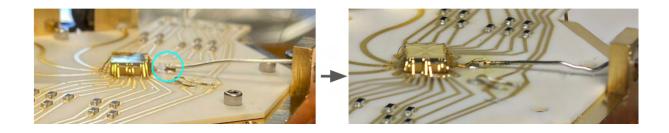
Daniel Slichter fabricated a new run of traps for us to use, but we still needed to select which trap we wanted to use. The traps can easily be inspected under a microscope while illuminated with white light, but visual roughness under the microscope does not necessarily translate into scatter spots when illuminated with glancing laser beams. In order to make sure that we did not install a new trap with similar or worse scattering to the old one, we rigged up a jig to illuminate the traps with a focused 405-nm laser pointer aligned to glance across the trap surface similar to our science

Figure A.3: The setup used to take the pictures in Fig. A.2. A 405 nm laser pointer was focused with a 150 mm lens and manually scanned over the trap surface. For the systematic evaluation of trap chips performed for the new trap candidates, the laser pointer was mounted to a stage for repeatability, and a clamp was used to hold the laser pointer button in the 'on' position.



lasers. To verify this procedure, we imaged the scatter spot on our old trap illuminated both with white light and this glancing 405 nm beam, as shown in Fig. A.2. We then examined all of our 12 candidate trap chips in the same manner to select a winner.¹

Figure A.4: On the left, the crack in the rf connection can be seen (circled in aqua). We chose to reconnect the rf using several gold ribbon bonds rather than using solder or epoxy. We had initially planned to replace all H21D connections with indium to provide resilience to thermal fluctuations, but the stainless steel wire that carries our rf was too stiff and the indium too gooey and the wire kept breaking free. We chose to instead use several ribbon bonds, for redundancy and flexibility with thermal fluctuations without loss of connection.



A.2 Pillbox

This trap was not intended to be operated at liquid helium temperatures and so has no heat shields between room temperature and our 8.5 K pillbox. We wanted to reduce line-of-sight from the trap to room temperature to improve performance. We did this by adding a lid to the top of the pillbox (where the old mesh was mounted) in the form of a 2 mm Edmunds UV-coat window, mounted with indium 'gaskets' on either side to prevent it from cracking under thermal stress.

A.2.1 Trap rf

The most essential repair we needed to perform was to reconnect our trap rf. The rf is carried into the trap on a solder-tinned stainless steel wire that is quite stiff, and we believe it is the stiffness combined with the thermal cycling which strained the joint until it cracked. We wanted an attachment method that would be robust to repeated, extreme thermal cycling without

¹Actually we ranked our top several candidates which was good as the hot-plate indium soldering had several casualties (see Appendix C).

compromising the quality of our rf signal. This meant that we did not want to replace the wire with a softer material - we just wanted a more resilient joint. Initially, we tested reattachment with indium, but when this joint was strained the stainless steel wire easily pulled free of the indium. Eventually we decided that a plurality of gold ribbon bonds connecting the stainless wire to the pad on the trap filter board was the best solution. Using several bonds gave us redundancy in case any single one failed, as well as reduced the strain on any single ribbon. The ribbons are also strongly attached at both ends but are themselves flexible to absorb any movement the stainless wire may experience while our trap temperature ventures wildly from 8 K to room temperature.

A.2.2 Filterboard

We also considered replacing all of the H21D joints on the trap filter board with indium solder, but this proved to be both unnecessary and infeasible. The filterboard (visible in Fig. A.4) holds the pins that our DC voltages connect to and filters the voltages on their way to the trap. There are 11 pins on each side of the board as well as 22 resistors and capacitors which filter these voltages on the way to the trap, all mounted with H21D. None of these connections had failed when we inspected them, likely because they do not experience any significant strain, so we chose to reuse the existing filterboard and keep the H21D connections there. Indeed, in the case of the pins, indium did not have sufficient strength to stand up to the attachment and reattachment of the pins, so we could not have used it for those joints.

A.3 Top viewport

The top viewport on the old chamber had four argon ion guns for surface milling, which we never ended up doing due to research presented in [192] which suggested that it might actually be harmful for cryogenic heating rates. These guns, visible in Fig. A.7, occupied a lot of space around the chamber and severely limited the space available for the imaging system. We replaced this top flange with a custom-welded reentrant viewport that did not have argon ion guns mounted on it. This allowed a great deal more space for the new imaging system, which has a very large diameter,

Figure A.5: On the left, our old pillbox had four nozzles for aligning the defunct Argon ion guns and no top window. Additionally, the optical access apertures into the pillbox were much larger than necessary, allowing greater line-of-sight from our cold stage to room temperature as well as providing little restriction for our laser beam paths to be level and aligned to the intended direction. By adding small copper nozzles that extend outwards from the pillbox on each of the optical access ports, we were able to reduce line-of-sight to room temperature by 70% as well as providing stricter boundaries for our beam paths. The top window also provides shielding from room temperature surfaces.

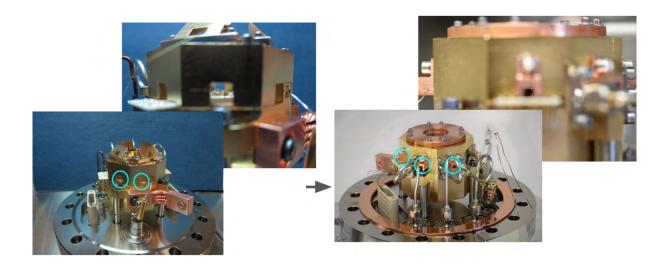


Figure A.6: On the left, our old trap mesh which sat on top of the pillbox (more than an inch from the trap) and was grounded to the pillbox. This protected us from charging on the top window but did little else. Right, our new mesh board is 4 mm from the trap surface and we can apply a variable voltage bias to it, which allows us to put a voltage 'lid' on our trap. This capability is crucial to successful ion trapping when our stray fields grow excessively large.

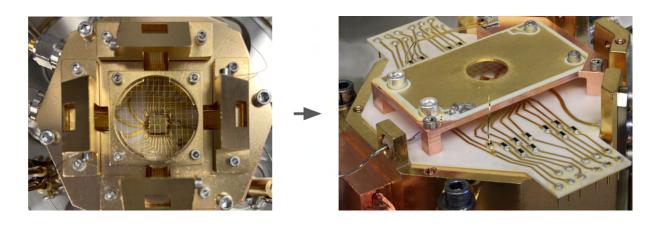
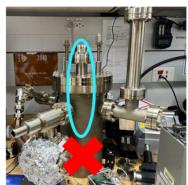
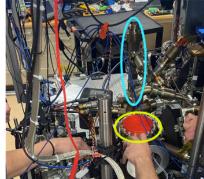


Figure A.7: The center photo shows our newly rebuilt chamber being installed back on the table in April 2022. On the left, the old TSP cryoshield is shown, and the right, the old top flange with its argon ion guns. Both of these pieces were never used for their intended purpose and occupied substantial space around the chamber. In the center photo, you can see the new, slender TSP housing as well as the standard-size top flange (the viewport was covered with a plastic lid for assembly, which was removed once we were ready to install the imaging system above). Red X's mark the eliminated pieces, while the colored circles show corresponding parts of the old and new chamber build.







as well as allowing the objective to be mounted close enough to the trap to accomadate its shorter working distance.

A.4 Homemade viewports

This chamber was originally constructed with custom-made low-stress fused silica viewports from MPF Products, coated for 280 nm light. When we started operating with calcium ions, we found that several of these viewports had higher-than-desirable reflection at 397 nm. However, to maintain good σ polarization for optical pumping, we wanted a high-quality low-stress viewport. Daniel Slichter proposed making our own viewports following [220] and this prototyping and testing work was performed by Alejandra Collopy.

The idea behind these viewports is that often when building an atomic physics chamber, one wants certain optical properties such as flatness or special coatings which may not be available in ready-made viewports, but are available as windows. One can mount the desired window to a machined conflat blank or half-nipple and thus create a homemade viewport with the desired optical properties.

Figure A.8: Left: The curing process for the homemade viewports: the flange is up on a small platform so it is not resting on its knife edge. The thermocouple is used to monitor the temperature during the curing process. Right: the homemade viewport in place on the assembled chamber.



The seal between the window and the stainless vacuum flange is created using indium, reinforced on both sides by epoxy. Indium has extraordinary resistance to stress and temperature cycling, but low tensile strength. Layering small amounts of epoxy on both sides of the indium provides extra strength as well as sealing any gaps between the indium and glass or metal. One of the primary motivations behind Weatherill's work was the low reusability of previous indium-sealed viewports [221] because strain on the stainless steel flange during attachment to and removal from the vacuum system can cause leaks in the seal. They added a separate clamping flange that holds the window to its steel pedestal and maintains substantial pressure on the joint at all times. Because resusability is not one of our primary concerns, we chose not to incorporate this separate clamping flange. It remains to be seen whether these viewports will be as reusable as those in [220].

Though this work was inspired by [220], the procedure we followed was quite different but ultimately successful. Though Weatherill used 0.7 mm diameter wire of Indalloy 165, which is comprised of 97.5% lead, 1.5% silver, and 1% tin, we chose to use 1 mm diameter 99.999% pure indium wire. Though they chose to carefully manipulate their wire to create a smooth, flat, and seamless ring, we chose to just smush the wire ends together and not trim any excess. While Weatherill used Epo-Tek 353ND epoxy, which passes the NASA low-outgassing standard ASTM E595 [222] when properly cured, we chose to use the thixotropic version, 353ND-T which we had on hand from building the calcium stabilization cavities, which does not. Though they found worse sealing with coated windows due to interaction of the seal with the coating surface and thus got windows coated only in the center, we used Edmunds UV-VIS coat windows which are coated to the edges. Both Weatherill and colleagues' design as well as previous work [221] utilized a separate clamping flange to apply force to the seal during the curing process and subsequent bakes. We chose instead to put a weight of unspecified mass on top of the viewport for an unrecorded portion of the curing time. We also chose to cure the epoxy at 100-110 C for one hour, rather than the specified 150 C.

²Though only barely! 353ND just squeaks under the threshold of 1% total mass lost (TML) after 24 hours at 398 K at 0.97%, while 353ND-T has a TML of 1.14%. I would still recommend using the lower-outgassing variety.

Perhaps shockingly, these viewports did hold vacuum. All three successfully assembled viewports passed helium leak tests at the 10^{-12} torr level. We installed one on our chamber on the north side, where our 397 σ beams enter the chamber, and consistently reach 7.7×10^{-11} torr at room temperature (as measured by our ion pump current), and have no measured ion pump current at 8.5 K. The primary lesson here may be that people are fussing too much about their vacuum viewports.

B

Stinger closed-cycle cryostat upgrade & operation

When this experiment first went cryogenic, it used a Janis ST-400 flow cryostat. This system provided good cooling power, sufficient to run the trap down to ~ 14 K (depending on helium flow setpoint), caused no noticeable mechanical vibrations which might affect experiments, and was relatively easy to use. However, it also required ~ 150 L of liquid helium per week [163], and even with NIST's helium recovery system's $\sim 90\%$ recovery rate [163], we had days when we could not run due to lack of helium. The cost was also becoming prohibitive, and the system requires the dewar be changed once or twice each week, which can be very difficult during government or pandemic shutdowns. For all of these reasons, we wanted to upgrade to a closed-cycle system. The installation was performed in the spring of 2021.

B.1 Old cryostat

The Janis ST-400 flow cryostat we used to use requires a helium dewar, typically 100 L, pressurized to 10 psi to drive liquid helium through the cryostat and out through the helium

recovery system. A heater prevents icing on the outside of the system, to avoid water dripping near the optics. Typically we used a flow rate of ~ 1 L/hour when running experiments, which gave a trap temperature near 15 K [163]. When not running experiments, we reduced the flow rate to ~ 0.3 L/hour to conserve helium [163]. At that flow rate, the trap temperature is ~ 55 K. Between these two temperatures minimal thermal shifts are observed and no realignment is required after temperature cycling. The system can also be run on liquid nitrogen to save money if there are periods of a week or more when we do not expect to run experiments, for instance during a government shutdown, which gives a trap temperature of ~ 77 K.

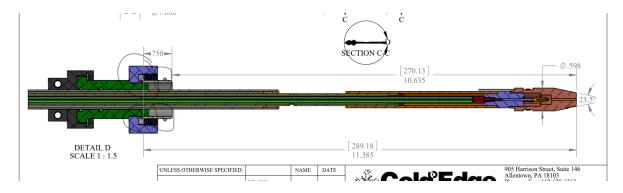


Figure B.1: Drawing of the custom bayonet insert Cold Edge made to fit into our existing cryostat. The ST-400's are not as uniform as one might hope, so to perform these retrofits the old cryostat needs to be disassembled and carefully measured to ensure that the bayonet will fit correctly, otherwise the cooling power will be compromised.

B.2 Stinger cryostat

In order to upgrade our cryostat to a closed-cycle system with minimal disturbance to the experiment, we purchased a Stinger cryocooler from Cold Edge with a custom bayonet made to fit into our system. This required disassembling the Janis system and carefully measuring the bullet that cradles the helium transfer line when it is in place, to ensure that the new system would fit snugly into the old system. While a retrofit of a Stinger into an ST-400 is becoming a common product from ColdEdge, they claim that the ST-400's are not all the same and thus this disassembly and measurement is necessary to ensure the new system will mate properly with the old one.

The technical operation of the cryostat is described in the main text, so this section will focus on the nitty gritty installation and every day operation of such a system.

B.2.1 Transitioning from the flow cryostat to the Stinger

The number one way to ensure a successful transition to a Stinger cryostat system is to avoid contaminating your system with air, whether via gross negligence or small leaks. Any non-helium gas in the system can freeze inside the narrow tubing inside the coldhead, building up like arterial plaques and eventually causing the system to clog. Additionally, loss of helium pressure will eventually cause loss of cooling power and runaway heating of the system. If you do not make sure that your system does not leak and take maximum precautions to avoid introducing air into the compressor, you will likely experience difficulties with clogging. In this section I will describe all the issues we had so that you can learn from our mistakes and hopefully have a better time than I did. While I will focus specifically on the flow cryostat retrofit, the majority of this information applies to any Stinger system.

The Stinger is so named because the cold tip is located at the end of a long, flexible bellows that rises from the coldhead like the stinger of a scorpion. To retrofit a flow cryostat with a Stinger, the bayonet on the end of this line is made to match the dimensions of the slot where the helium transfer line previously sat. The custom bayonet inserts into the ST-400 system and screws onto the existing threads via a brass collar, and the internal portion has an adjustable depth which is held at the correct length by several set screws. During our initial installation, a key set screw was not tightened, resulting in ~ 100 psi of helium gas blasting the internal bayonet out of the collar, snapping several delicate wires and, we believe, contaminating the adsorber in our HC-4E with air. This motivated us to tap the brass collar with two additional holes, so our system now has three set screws holding it in place. After this incident we continued to have near-weekly clogs, despite repeated cleaning cycles with the ColdEdge-provided liquid nitrogen cold trap, until the adsorber was replaced, after which the clogs were much reduced. We believe the contaminated adsorber was

¹We did both!

slowly leaching contaminants back into the system which would then freeze and cause clogs. We also had trouble with leaking around the interface between the ST-400 and the Stinger bayonet contributing to clogging, which we fixed by replacing all of the old o-rings on the Janis side along with a KF blank (on the previous helium recovery port) that was not rated for positive pressure. With these changes we reduced our helium gas loss to < 1 psi per day, and eventually we were able to completely eliminate the helium gas leaks by replacing a faulty aeroquip hose which had developed a pinhole leak, as well as tightening the Swagelok joints inside the ColdEdge manifold (new manifolds are welded inside, rather than Hylok, and thus should not leak). We believe the pinhole leak in our hose was due to constantly undoing and redoing the hose in order to take the cold trap in and out of line to clean the gas.

By replacing the press-fit tubing connections on the 'charge' and 'vacuum' ports on the manifold with Swagelok, we were able to pull a good enough vacuum on the system (being sure NOT to evacuate the HC-4E) to do a helium leak check with an RGA, just as we do for our ultra-high vacuum chambers. We highly recommend doing this in your system if possible, especially if you end up adding a lot of additional plumbing as we did (discussed later in this Appendix). However, even with 99.999% pure helium gas and no remaining leaks, we still experience periodic clogs. Scientists at Quantinuum believe that actually gas purities of up to 99.99999999 purity may be required to completely eliminate clogging, but this can be approximated by adding filters inline to the system as well as on the charging port,² and reasonable cold hold times are achievable using the various cleaning and recharging cycles detailed below. These have largely been adapted from procedures provided to us by ColdEdge, and refined by trial and error over the past few years.

B.2.2 Cleaning cycles

Due to the large amount of leaking and contamination present in our system at various times, we developed several different cleaning cycles for the Stinger system based on the severity and nature of the clogs. We would like to note here that others in the group who bought their system

²Personal communication with Riley Ancona at Quantinuum, via Daniel Slichter.

whole from Cold Edge did not experience clogging anywhere near the level we did, and we believe it was the difficulty of achieving a smooth interface between the old and new systems, as detailed above, and not any intrinsic fault of the system, that necessitated all of the cleaning. The two more serious procedures were given to us by Cold Edge, while the mini unclog was our own idea, inspired by the way our system would sometimes have a better base temperature after an unplanned power outage.³

B.2.2.1 Mini unclog

I will detail three different levels of cleaning cycle. The first is for a functioning system which has been cold but begins to clog, indicated by the cold tip temperature slowly creeping upwards while the cold head temperature creeps downward. This is caused by contaminants slowly freezing and clogging the narrow, cold parts of the Stinger system. We can attempt to remove those contaminants by isolating the experiment side of the recirculating system and pumping on it as we warm up, either partially or fully to room temperature. This is done by closing first the Supply and then the Return valves on the Stinger manifold, allowing enough time for the helium in the experiment side to be sucked back into the recirculating compressor before closing the Return valve (the three-way valve can be turned to Both Sides to evacuate the helium faster, before the clog starts to thaw), then turning off the F-70 to allow the system to warm up. Allow the clog to build up pressure inside the manifold as it thaws, using the Supply Side Gauge to tell how much gas is coming off at which temperatures. Typically the bulk of our clog thaws out by the time the cold head temperature gauge is around 50 K, but we usually allow the cold head to reach 90 K to ensure that any air has had a chance to unfreeze. Typical values for our clog volumes equate to 80-100 psi on the manifold gauge by 50 K, and an additional 10-20 psi between 50 and 90 K, although these have been reduced since adding the filters (see below). When you have reached the temperature above which you think your contaminants have thawed, you can pump out the system via the Vacuum valve (we just use a scroll pump for this), close the Vacuum valve, set the three-way

³Shout out to the terrible power grid stability in Boulder, CO.

valve to Return Side (Normal Position), open the Return valve, wait ~ 30 s, then open the Supply valve, and restart the F-70. While this technique works for minor clogs, if you have serious leaks or contamination this will likely only be a stopgap measure between more thorough cleaning cycles. Additionally, this method inevitably results in some loss of helium gas along with the clog, so it may be necessary to recharge the system periodically. We have found that it is fine to recharge the system when it is cold, although better results can be had if the new helium is added through a cold trap and/or filter to avoid introducing air into the system.

B.2.2.2 Cold trapping

This cleaning cycle, using a cold trap, is for lightly contaminated or leaky systems that are not responding to the above method. It requires the use of a cold trap (provided by Cold Edge with the Stinger) and a liquid nitrogen bath. For this cleaning cycle you want to start with your entire system at room temperature, and if your system leaks helium, it can be nice to charge it back to ~ 100 psi before beginning the cleaning cycle, just to be sure any contaminants potentially introduced during the charging process have a chance to be removed. Close the Supply and then Return valves on the Stinger manifold and insert the cold trap in the line between the Supply line from the HC-4E and the Supply Recirculator port on the Stinger manifold. We were doing this so often that we got sick of unscrewing the fittings and paranoid about how much air may or may not be getting in each time, and so designed a second, homemade manifold so the cold trap was always attached to the system but can still easily be bypassed. Submerge the cold trap in liquid nitrogen and be sure the valve on the cold trap is open, then open first the Return and then the Supply valves on the Stinger manifold. This will allow the gas to begin cycling through the system, including through the cold trap, and hopefully contaminants will be frozen out in the cold trap. Cold Edge recommends running this way for "as long as is convenient", at least a few hours. We typically did 3-5 hours. When you are finished cold trapping, close first the Supply and then the Return valves on the Stinger manifold, then the valve on the cold trap, and either remove the cold trap from the lines, reconnecting them as they were, or bypassing the cold trap if using a second

manifold. The cold trap will be at ~ 100 psi with helium gas as well as potentially large volumes of frozen gases that will expand quickly as they thaw, so it is good practice to start pumping out the cold trap before removing it from the liquid nitrogen, to avoid being surprised by the pop valve. Pump on the cold trap as it thaws, and additionally it can be a good idea to bake the cold trap to prepare for the next use immediately after using it. We do this with heater tape set to ~ 120 C to be sure all water is removed. We typically bake for a few hours or until the whole system has been above 100 C for half an hour or so. Once the cold trap has been removed from the line, you may reopen the valves on the Stinger manifold and let the helium recirculate freely. Cold Edge recommends letting it circulate at room temperature for 30 minutes before beginning to cool down, which is to allow the oil in the HC-4E to reach operating temperature and to let the adsorber in the HC-4E remove some contaminants, but we have found that while helpful this is not strictly necessary if you don't turn the HC-4E off.

While this cleaning procedure served us well for a while, it seemed as though the effectiveness of the cold trap degraded over time, and having to be at room temperature for "as long as is convenient" was not very convenient. We started going straight to total gas exchanges instead as they required less time at room temperature, no messing with liquid cryogens or baking filters, and had a higher success rate.

B.2.2.3 Total gas exchange

The final cleaning cycle is for when your system has been hopelessly contaminated and you want to burn it all down and start again. You cannot just vent all the gas from your system and totally replace it as the compressor is not built to be below atmospheric pressure and you risk sucking air in and seriously contaminating your adsorber. Instead this is the procedure we follow, heavily modified from the one provided by Cold Edge. We assume that you begin at room temperature with all valves in their normal operating orientations, the F-70 off and the HC-4E on.

(1) Ensure you have a cylinder of ultra-high purity helium connected to your Charge port and a vacuum pump (we use a scroll pump) connected to your Vacuum port.

- (2) Close Return valve. This is to prevent the compressor overpowering your (probably weaker) vacuum pump and sucking air in to itself.
- (3) Open the Vacuum valve and let it begin to pump out the helium, keeping an eye on the pressure via the Supply Side Gauge. Do not let the pressure drop below 20 psi. Pumping out the helium through the capillary can be very slow, this can be sped up by cracking the tee valve towards Both Sides, but it's best not to dump 100 psi of helium straight into a scroll pump so just open it a little bit. Once the pressure in the compressor, as read on the Supply Side Gauge, reaches roughly 22 psi, close the Vacuum valve and reopen the Return valve.
- (4) Slowly open the Charge valve. Charge the system until the recirculator reads ~ 100 psi, then close the Charge valve.
- (5) Repeat steps 2 through 4 3-5 times, on the last repeat charging the system fully to 100 psi (or a little over, I often go up to 115 psi or so).
- (6) Ensure your valves are back in their normal operating orientation and then you can immediately begin cooling down by turning the F-70 back on.

B.2.3 Supelco filter manifold

Perhaps the single most helpful upgrade we have made to the system is to install a pair of Supelco 27600-U helium filters in parallel between the valve manifold and the Stinger can itself. This tip came from Riley Ancona at Quantinuum⁴ and has helped to substantially increase our cold hold times and reduce downtime due to Stinger clogs. These filters, whose contents seem to be highly proprietary, claim to improve helium purity from 99.997% up to 99.99999%, removing a broad range of common contaminants. Riley told us that the filters can be regenerated by baking at 90 C for 24 hours, though we have experimented with shorter overnight bakes. He also noted

⁴Personal communication via Daniel Slichter.

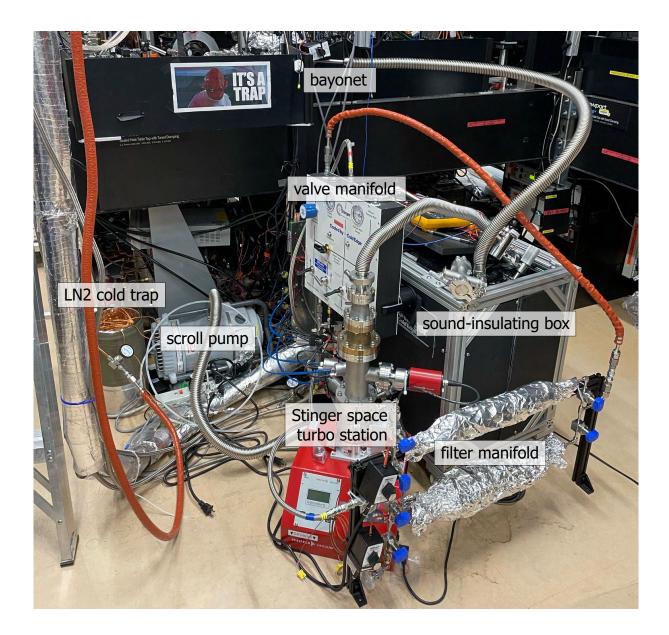


Figure B.2: A photograph of our Stinger in place the experiment, with all attendant accessories. The most important are the Supelco filters, which are mounted in parallel in a custom welded manifold fabricated by Swagelok. These filters greatly extend the time we can run between clogs, and having two in parallel means we do not suffer any downtime when the filters need to be baked. Both filters are left wrapped in heater tape and ready to bake at all times, but we are sure to unplug the heater tape when a given filter is inline to avoid heating up and releasing contaminants back into the circulating gas.

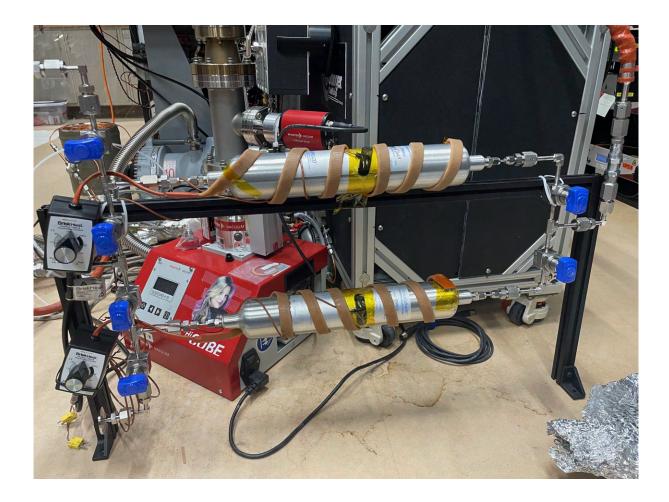


Figure B.3: A photograph of our custom-welded Supelco manifold without the foil wrapping, showing the heater tape and thermocouple for baking each filter separately. The separation was designed to be sufficient that one filter could be baked without heating up the other (and thus potentially releasing trapped contaminants back into the system).

that the max temperature recommendation of 100 C seems to be for the epoxy which holds the fittings on the ends of the filters and that it will melt if it gets too hot, and that the filters should be mounted horizontally for best performance - apparently if mounted vertically eventually oil⁵ will build up and clog the filters more quickly. They found that after six months or so the filters need to be totally replaced as the regeneration stops working well and this is roughly in line with what we have seen. We have a pair of filters in a custom welded manifold⁶ such that we can switch between them without having to fully warm up and bake for 24 hours, although we did operate that way for a few months while testing the filter performance.

Each filter has its own heater tape, thermocouple, and pump out port for baking and the manifold was designed to give enough space between that baking one filter does not disturb the other one's inline operation. Quantinuum says they had the Swagelok fittings that come on the filters replaced with VCR fittings, but we chose not to do this and do not seem to have leaks or other issues from the Swagelok. It is impossible to install the filters without exposing them to air as Swagelok fittings are not self-sealing, so we found that it is best to install the filters in the filter manifold while valved off from the main system, and bake them in place before letting the helium flow through them.

When we were operating with a single inline filter, we needed to warm up approximately every two weeks to bake the filter for 24 hours. When we switched to the two parallel filters in the manifold, the time between filter swaps was extended to roughly a month. It is possible that the single filter had small leaks around the fittings, or that the low-pressure helium system as a whole has leaks that open and close as it travels from 10 K to room temperature. I was able to separately leak-check the filter manifold before installing it in the system, and it allows us to swap filters without warming up, and I think these two factors contribute to the extended filter lifetime. After roughly six months of operation, the time between filter swaps started to drop dramatically

⁵We believe the oil comes from the HC-4E compressor, which is an oil-based compressor. One might wonder about the wisdom of using an oil-based compressor with UHP gas and indeed a second experiment in the group which has been testing an oil-free compressor has not had to do a single regeneration of their inline filter.

⁶Designed and manufactured by Swagelok, shout out to Bonnie and Johnny.

⁷I have long suspected this to be true but such leaks are extremely difficult to track down.

and so I replaced both filters. This coarsely matches anecdotal evidence from Quantinuum that their filters were regenerable for about six months before they needed to be replaced. Anecdotally, the Quantum Networks experiment which uses a KNF oil-free compressor instead of an HC-4E has never regenerated their filter, although they do experience periodic clogging.

B.2.4 Miscellaneous plumbing upgrades

We also made various minor plumbing upgrades during our efforts to find and eliminate leaks. Even very small leaks can cause lots of clogging and helium loss over time, but the system is very difficult to leak check as it consists of many meters of thin tubing with a lot of joints and relatively few pieces which can be isolated from each other. Earlier I mentioned that we replaced the press-fit 1/4" tube fittings that come attached to the Charge and Vacuum ports of the valve manifold with Swagelok ones, which we further adapted to KF for the Vacuum port. Having our Vacuum port adapted to KF25 allows for much better conductance than pumping through 1/4" nylon tubing. For charging the system, we now supply the helium from a cylinder in the galley through a 40 foot custom hose purchased from Swagelok (SS-FJ4RF4RF4-480-F), and then a homemade liquid nitrogen cold trap (a long coil of UHV-cleaned 1/8" copper tubing that can be submerged in liquid nitrogen) as well as an additional Supelco filter. This helps ensure that we are supplying the cleanest possible helium into our system and helps extend the lifetimes of our filters and compressors. We used to connect and disconnect a cylinder of UHP helium in the lab, purging the 1/4" nylon tubing we supplied the gas through to avoid putting air into the system, but we have found that this permanent, clean, leak-free, cold-trapped and filtered system provides better performance and is worth the hassle and expense.

In addition, we found replacing as many of the standard Swagelok fittings as possible with VCR fittings helped to reduce leaks and also allows for higher quality re-sealing of joints as you simply need a new copper gasket, rather than re-torquing on the same compression fittings. The filter manifold and as much of the clog catcher (see below) as possible was done with VCR fittings. To adapt the VCR pieces into the system, we purchased custom-welded VCR to Aeroquip adapters

from Clark Industries.

B.2.5 Automatic clog catcher

One of the handiest upgrades we made was to build the system we call our 'clog catcher.' This name comes from the fact that, when a clog starts to occur, we try to valve off the valve manifold and Stinger from the HC4E, to use the Stinger itself as a sort of cold trap which we can warm up and pump out as described in the mini unclog procedure above. However, by the time the Stinger temperature sensor reaches 25 K, a substantial amount of the clog has been thawed and thus released back into the recirculating helium where it could eventually clog again. In order for the mini unclog to be effective, it has to be implemented early in the clogging process, which can be tricky if you need to drive into campus and manually actuate a bunch of valves. The clog catcher, shown schematically in Fig. 6.4 and photographed in Fig. B.4, consists of four pneumatic valves (Swagelok 6LVV-DPFR4-P-CM), the actuation of each is done via a solenoid valve (Parker 71385SN2GNJ1N0C111C2) controlling the flow of compressed air tapped off of our optical table. These solenoid valves are controlled by a Labjack T4 with a PS12DC, which also controls the scroll pump connected to the Vacuum port. This is integrated into Artiq along with control of our F-70 compressor, enabling us to program in all of the cleaning and unclogging routines described above (with the exception of the total gas exchange, which relies on reading an analog pressure gauge⁸ and also is performed quite infrequently).

B.2.6 Daily operation

During normal operation, the Stinger keeps our trap near 8-10 K. The old cryostat operated more at the 12-15 K region, and we see much better ion lifetimes with this small temperature drop, which we attribute to increased cryopumping of hydrogen. In fact we seem to cryopump hydrogen so efficiently that it forms an ice on our trap surface and we find much improved trapping and lifetimes if we degas the trap by warming up to 35 K each morning to release the hydrogen and

⁸We did purchase a digital pressure and flow meter from Alicat (MC-50SLPM-D-DB9M-PCV50-25VCRM-IB) on the recommendation of Riley Ancona from Quantinuum but have not installed it as of this writing.

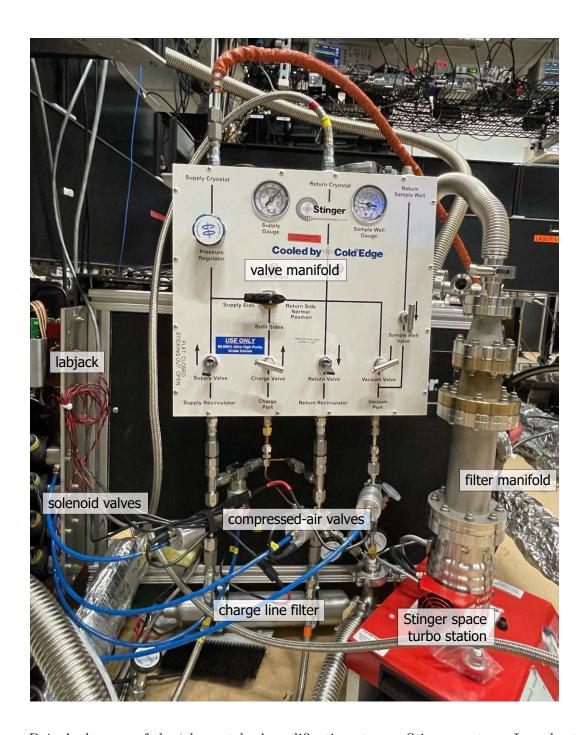


Figure B.4: A close-up of the 'clog catcher' modifications to our Stinger system. In order to be able to run our degas and mini unclog procedures remotely, we have added external to the manifold compressed-air actuated valves in analog to the Supply, Return, Vacuum, and Both Sides/ Return Normal valves. A Labjack controls the solenoid valves that pass air to these valves as well as a scroll pump, enabling us to perform many of our Stinger maintenance protocols remotely.

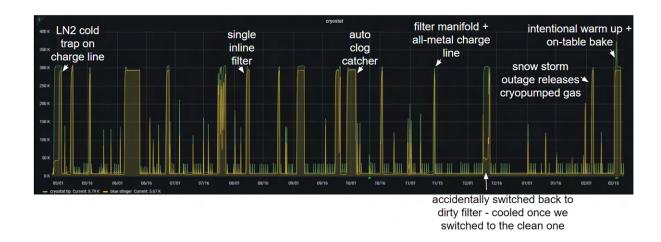


Figure B.5: Sample of nearly a year of temperature data from our Stinger. Room temperature warm-ups as well as daily degases can be seen, and significant events are annotated. The green trace shows the temperature at the tip of the bayonet while the orange trace shows the temperature within the 'blue can' exchange space.

allow it to pump into the ion pumps and TSP. This is done by closing the Supply valve for about five minutes, then reopening it. We do not turn off any compressors for this.

Over very long periods of being cold, we also seem to cryopump large amounts of something that thaws between 25 and 50 K. This is a gas that does not seem to be well pumped by any of our other methods (i.e. the NEG or TSP), and so any unplanned warmups above ≈ 50 K can liberate large amounts of gas which do not pump down as quickly as one might hope. We attempted to increase the temperature of our morning warmups to 100 K to allow this gas to be liberated and pumped back down via non-cryopumping methods but had issues with loading and lifetime for hours afterward, even though the ion pump currents showed that the vacuum pressure had returned to normal levels. We now believe this could be due to poor conductance out of the pillbox, and we save these larger warmups (we call them burps) for evenings or weekends when the slow pumping of the released gas does not cause issues.

Over the course of weeks of operation with this daily small warmups, we will typically see the trap temperature start to creep upwards, accompanied by the cold head temperature creeping downwards. This is indicative of reduced flow in the system, whether due to a clog starting to form or loss of helium pressure. We think the cold head gets colder because the reduced flow means less heat load for it to cool. Once the cold head temperature drops below 5 K, we almost always clog. Before the filter manifold, we would typically perform a mini-unclog to try and isolate and pump out the offending frozen gases. After the filter manifold, we combine this mini-unclog with changing over the filters. We see very good performance this way and can extend our time between room temperature excursions.

Every mini-unclog loses some amount of helium along with the clogged gases, so we need to add more helium periodically. In our system, we find that once the recirculating helium drops below roughly 80 psi, we start to lose cooling power and become much more susceptible to clogs. We add new helium through a liquid nitrogen cold trap as well as a Supelco filter, and flush the line sequentially through the outlet port in each section of the charge line to limit the contamination that comes in with the new helium.

B.2.7 Cooling power measurements

We performed cooling power measurements on the Stinger system using the heater built into the tip of the bayonet, operating it by way of a benchtop power supply. We applied various currents to the 50 Ohm heater and recorded the temperatures produced at the tip of the cold head to determine the cooling power of the system. Note that our base temperature here is below 10 K, and we have seen temperatures as low as 8 K, better than we could achieve with the ST-400 system.

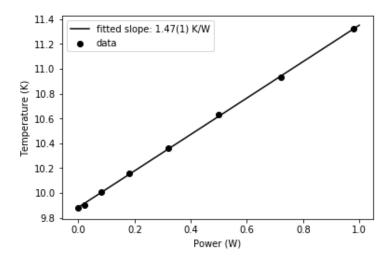


Figure B.6: Cooling power measurements were performed by driving a 50 ohm heater with a variable current supply. We measured that the tip of the coldhead warms up 1.5 degrees K for every W of heat applied.



Two methods for mounting microfabricated ion traps with indium solder

Microfabricated surface electrode traps [102, 103] have become very popular in the field of trapped ion quantum computing because the promise of using existing microfabrication techniques, easy reproduction of many identical traps, and recent advances in integrating many different technologies into the chips themselves are all big advantages for eventually scaling trapped ion quantum computers. However, these small and delicate chips need to be mounted to a larger trap assembly and, increasingly often, need to be able to withstand a huge amount of thermal stress (either in the form of cryogenic operation, dissipation of heat from applied control fields, or both). Historically our group has mounted traps with epoxies, often Epotek's H21D, a thermally and electrically conductive epoxy. However, given H21D's poor resistance to thermal cycling (see Fig. C.1, [137], and Appendix A) and relatively poor thermal conductivity, we wanted to investigate alternatives.

We looked to indium and its alloys for their high plasticity and low melting point; indium has been used in soft vacuum seals and cryogenic soldering for decades [223]. A low melting point is

Figure C.1: Left: Trap mounting tests using various epoxies, showing the high failure rates they exhibit when subjected to thermal stress. Right: An ion trap with an integrated detector which was mounted using Epotek H21D epoxy and shattered during the initial cooldown due to differential thermal expansion between the trap chip and mounting pedestal. Photo from [137].





desirable because microfabricated ion trap chips are very fragile and direct soldering at high heat may not be desirable, especially when more complicated technologies have been integrated into the chips. We tested both pure indium and a 97/3 indium/silver alloy and found no difference in bond strength or quality, but the alloy with silver was slightly less sticky and thus easier to maniupulate with tweezers. Due to the effective indistinguishability of the two alloys, we will generally refer to both as 'indium' in this appendix and won't specify which metal was used for which tests.

We investigated two methods for mounting trap chips using indium solder:

- (1) soldering with NanoFoil, a very thin layer of exothermic material, and
- (2) direct soldering using a hotplate.

The initial trap mounting with both method 1 and method 2 used pure indium out of paranoia, but subsequent tests by the Quantum Networks experiment have used 97/3 indium/silver with no adverse effects. Traps mounted using both of these methods were successfully used for cryogenic trapping of ions in systems that experience extreme thermal stress in [137] and the Quantum Networks experiment at NIST for method 1, and this work for method 2.

C.1 Indium soldering using NanoFoil

NanoFoil is a proprietary material produced by the Indium Corporation for soldering with indium. It consists of alternating thin layers of aluminum and nickel which, when given sufficient impulse, experience a cascading exothermic reaction that can very locally heat a pair of joining surfaces coated with indium and create a solder joint. The main things to be aware of when mounting a trap chip with indium is that indium will not bond well to silicon or sapphire, and though it will bond to copper it diffuses into the metal and creates brittle intermetallic compounds that compromise the joint. For this reason, the bottom of traps to be mounted with indium should be gold-plated, as should any mounting surfaces.¹

¹This also has the advantage of making your pillbox look shiny and gold! Indium Corp. claims that non-metallic surfaces can be bonded to indium if you heat the surface to 350 C, cool it to 200 C, and then rub it with indium-coated felt but we did not have success with this technique on silicon trap substrates, so we recommend gold-coating.

We performed tests where we mounted 8 sample trap chips to a piece of gold-coated copper (the same material our trap pedestals are made of) using indium and NanoFoil and then subjected the samples to extreme thermal stress by dunking them in liquid nitrogen, thawing them with a heat gun, then throwing them back in the liquid nitrogen, for eight total freeze-thaw cycles. None of the samples that were mounted properly experienced a single failure: there was no delamination of trap chips from the substrate or no cracking of trap chips due to differential thermal expansion, both of which can be seen in Fig. C.1. The only joints that failed were those that were mounted with insufficient pressure.

NanoFoil is very brittle and very flammable, so care should be taken when working with it. It is difficult to cut neatly, and often creates miniscule chips of NanoFoil that constitute a fire hazard and should be swept up and ignited intentionally so they cannot accidentally ignite on a bench surface. We found that scoring and snapping creates neater edges and fewer chips if square shapes are acceptable for your task. If you require a difficult-to-cut shape, Indium Corp. offers custom-cut NanoFoil sheets. NanoFoil can nominally be ignited just by applying pressure which was initially a big concern, but we did not experience a single accidental NanoFoil ignition during any of my testing. We still recommend wearing protective gloves, e.g. leather, and eye protection when working with NanoFoil.

The whole stack (trap chip, indium, NanoFoil, indium, substrate as shown in Fig. C.2) needs to be under a minimum of 50 psi of pressure when the NanoFoil is triggered in order to form a strong bond. Pressing down hard on the top surface of your expensive and delicate trap chip is something most people want to avoid. Additionally, because the NanoFoil is effectively a small explosion in a layer of very low melting point metal, there can be splatters which could feasibly make their way onto the chip surface. The method we used to prevent such occurrences for the trap used in [137] was to create a small jig that served the triple purpose of registering the whole stack, protecting the top trap surface from splatters, and allowing us to apply pressure only to the periphery of the trap chip. We considered using Teflon for these jigs but it cannot be machined with tight enough tolerances, so finally chose resin-based 3D printing using a FormLabs printer

owned and operated by the CIRES workshop at CU Boulder.² As opposed to the common plastic extrusion 3D printers, resin-based 3D printers have considerably higher resolution (\sim 25-100 μ m, depending on exact material and direction of features) and some resins can be solvent-cleaned and have been used in ultra-high vacuum [224]. While we never planned to place the jig itself under vacuum, this helped soothe our worries that contact with the trap surface could transfer any sort of material which could outgas or otherwise interfere with vacuum or trapping.

We chose to get four iterations of the soldering jig: one with tight tolerances and one with looser tolerances in each of two resins, the Rigid and the High Temp. Both of these materials can be cleaned with a variety of solvents so long as exposure time is limited. In practice we have used the Rigid material only due to its better solvent resistance. It is important to note that the jig must have sufficient access to ignite the NanoFoil, with e.g. a resistive soldering iron or 9 V battery. The Networks experiment has found it helpful to cut a notch into the indium foil to allow access to the NanoFoil for ignition. They have also found it helpful to pre-tin the mounting pedestal using the method described below in the hotplate soldering section to get a thin and even layer of well-adhered indium, and use indium foil only for the layer between the NanoFoil and trap chip. The neatest joins are created when the indium layers are thin and uniform as there is less material for the NanoFoil to expel out the edges. We applied the necessary bonding pressure to the jig using a Mark-10 M2-10 force gauge. Once the jig and everything is set up, the process is very quick and fun.

We were intially worried that the residual layer of NanoFoil might compromise the softness of the indium bond but that does not seem to be the case. We also worried about magnetic properties of this bond layer disrupting our quantization field but found the magnetic permeability of indium bonded with NanoFoil to gold-plated copper to be extremely good, causing less than a hundredth of a Gauss disturbance measured in a 16 G field when moving the sample around the Gaussmeter.

²Ken Smith and Jim Kastengren at CIRES also did a huge amount of the machining for our rebuild! I highly recommend the CIRES shop if you are in the Boulder area.

Figure C.2: Left: The CAD model for the 3D printed jig used to align the trap, indium foil, and NanoFoil layers to the pedestal as well as to protect the trap surface from indium splatter. The jig was designed to place no pressure on the actual trapping surface, to avoid deformation/scratching of the electrodes as well as potential residues, and to distribute force as evenly as possible over the ground plane or other less-sensitive areas. It is important to have even pressure distribution across your trap surface to get a strong and level bond. Right: a cartoon cutaway view of the bonding stack (not to scale). The recess in the jig over the trapping region ensures that pressure is applied only on the outer portion of the trap chip.

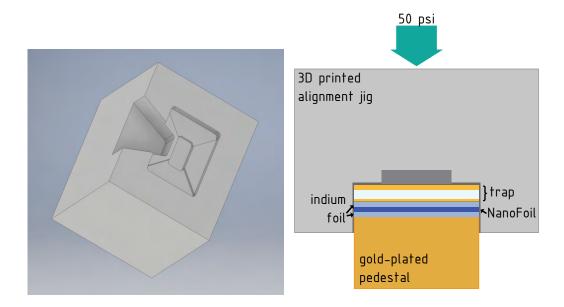
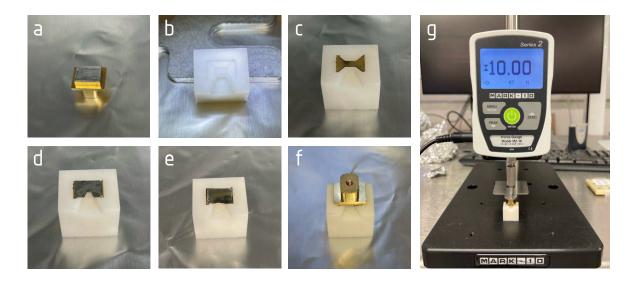


Figure C.3: Step by step photos of the NanoFoil trap mounting process. a. The pre-tinned trap pedestal. b. The 3D printed alignment jig. c. The trap placed face down in the alignment jig. d. The indium foil layer, with a notch cut for NanoFoil ignition access. e. The NanoFoil layer. f. The trap pedestal. g. The full assembly in the force gauge used to apply the pressure for bonding. All photographs kindly supplied by Katie David.



C.2 Indium soldering using a hotplate

If you do not want to purchase NanoFoil or design and manufacture alignment jigs to evenly distribute pressure across your trap, there is another way. When we went to mount our new trap, we knew that the NanoFoil would not work because we had areas on our trap pedestal which could not under any circumstances be coated with indium, as this would prevent us connecting the necessary ribbon bonds. While we did experiment with several masking materials to keep the indium off of these areas (Kapton tape, TechSpray Wondermask) we had concerns about both methods.

The WonderMask was easy to spread, adhered well to the gold-coated copper, and cured quickly in air but not so quickly that it was difficult to manipulate. It resisted being flooded with flux (Superior No. 30), vigorously smeared with gobs of molten indium, and prodded with the soldering iron. It also peeled off easily, so long as you were careful to apply a thick, even layer. Our concern with this method was with any residues left behind by the mask potentially causing problems with the ribbon bonds, as cleaning with solvents this close to the mounted trap would be very difficult.

If you had a larger area with straight edges you were trying to protect, Kapton tape might work well. It also resisted being flooded with flux, smeared with indium, and gently prodded with a soldering iron. But for the size of area we were trying to protect it was too small to get good adhesion. There was also a minor concern about residue left behind by the adhesive.

We also tested cold welding - indium is known to cold weld to certain metals as long as you remove the oxide layer. We were able to successfully produce cold welds between pieces of gold and indium by wiping the indium surfaces with a 5% solution of HCl and applying substantial pressure, but the bonds did not feel strong and the amount of pressure needed felt high to apply to a trap chip. We also later learned that HCl, when heated, creates a nasty green crust, so this method is not compatible with hot-plate soldering.

In the end we decided that hand-soldering using a hot plate as the heat source was the safest method. After much trial and error, below is the final procedure we arrived at for mounting a 550

Figure C.4: Left: The first hotplate soldering test. The gouges made when pressing the trap down with tweezers were deemed far too large. Right: The ion trap we installed in our system. The surface gouges on this attempt were minimal and the bond was made well and square to the trap pedestal.

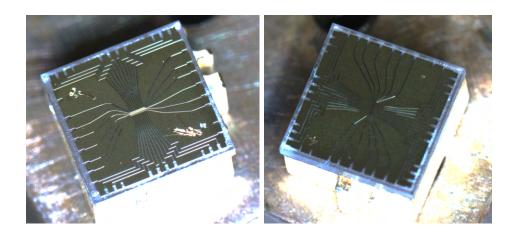
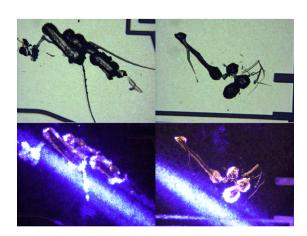
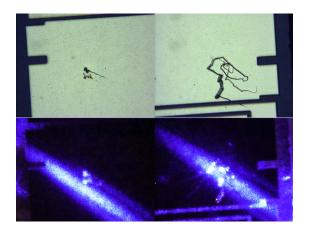


Figure C.5: Left: Close up of the damage to the trap surface from pressing the trap chip down with sharp tweezers. The size of these scratches is due to my hand shaking during placement and removal of the tweezer points. The top images are illuminated with white light, while the bottom row use the same scatter-assessment technique used to select trap chips described in Appendix A. Right: Close up of the gouges on the successful hotplate soldering test. It is difficult to compare the magnitude of scatter due to auto-exposure on the camera, but the damage to this trap was much less substantial and we did not think it would interfere with beam paths.





 μ m sapphire substrate, with gold on both sides,³ to a small gold-plated copper pedestal. It is easiest to get indium to adhere well to gold with the help of some flux and a heatgun, and so this portion is done outside the clean assembly area, then the pre-tinned pedestal is cleaned according to UHV cleaning procedures, and the final mounting is done in the clean tent.

(1) Tin trap pedestal:

- (a) Hold trap pedestal in a 'helping hands' type mount it will get very hot. Begin heating the pedestal with a heatgun (you can use a mounted heatgun or just your non-dominant hand). With your dominant hand, flux the surface with Superior No. 30 Flux, then pick up some indium directly on the soldering iron and spread it on the pedestal surface. The indium will not go where you don't push it with the iron, so guide it around with the tip, and it won't adhere well without the flux. You can check the adherence by trying to scrape the indium off with a razor blade- if it's properly adhered there will be a layer you won't be able to easily scrape off. We found that the best way to get a thin, even layer was to add sufficient indium and get it hot enough to form a neatly domed surface over the area you want covered, then flatten the doming by scraping it gently with a razor blade.
- (b) Clean the trap pedestal with soap and water to remove any flux residue, follow with the standard vacuum cleaning procedure. Some indium sometimes comes off during this step, especially if it's not well adhered, so make sure it looks like you still have good coverage.

(2) Mount trap:

(a) Put the tinned pedestal on a hotplate in the clean tent. We had to set our hotplate to 260 C to get the indium to melt easily, but it was on quite a large and high copper block, so experiment with the minimum temperature you need. It is visually obvious when the indium becomes molten.

³Electroplated on the top, evaporated with a Ti sticking layer on the bottom.

- (b) Take your trap chip and double check the orientation relative to your pedestal. We found it can help to tin the bottom of the trap as well: Take an old pedestal or chunk of gold-plated copper (it is best not to use bare copper, as indium and copper form brittle intermetallics) and melt some indium on it. Then dip the trap in can just set it down and pick it up or try to swipe. Be careful as the indium is quite gooey and may try to grab your trap (this is how I ended up soldering a trap to the control panel of the hot plate best to avoid it). Check the bottom surface to see if indium seems to be sticking, just a little is all you need most of it should be on the pedestal.
- (c) You can use tweezers to spread out the indium on the pedestal to make sure you have an even, molten layer. Then, using cleaned/different tweezers than you just stuck in the indium, carefully place your trap down. You can tap it gently from the sides to get it square with the pedestal.
- (d) Next is the absolute worst part: pressing the trap down. The indium will not seal properly without some downward pressure. You could use a jig, as for the NanoFoil method, but we just used sharp tweezers to press down on the edge of the trap chip. Under the microscope, carefully bring your tweezer points down over an unimportant piece of the trap. If there is a portion of the mounting surface which you are trying to keep free of indium, touch the tweezer down on that side first. As you press down, you should feel the surface tension wick the trap down, and a bead of excess indium will ooze out opposite of where pressure was first applied. This can be easily removed once everything has cooled down a bit. It is best to have some bad chips or something to practice this it took me several tries to get a clean press without making gouges in the trap which could scatter light.
- (e) DO NOT try to adjust the alignment after you've pressed it down it will need to be pressed down again after being poked from the side while the indium is molten.

Let it cool and then test it - poke it from the side and see if it pops off. Too little indium, too

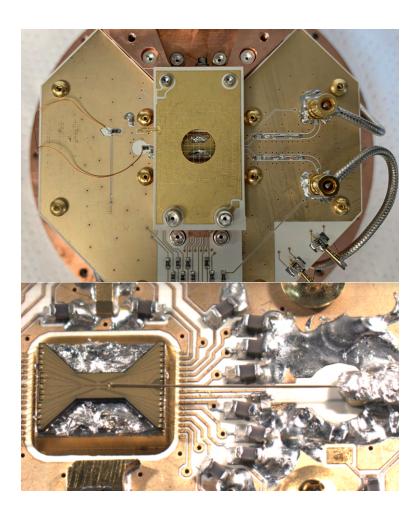
little heat, and too little pressure can all cause the bond to fail.

C.3 Use in experiments

Three experiments within the Ion Storage group have taken advantage of indium trap mounting, two using the NanoFoil technique and one using the hot-plate technique. The first NanoFoil-mounted trap was installed in the detector trap apparatus described in [137] after they suffered their second catastrophic trap delamination in 2019. We mounted one dummy trap as a test, then did the actual trap. It survived multiple thermal cycles, successfully trapped ions, and it was in this trap that the fast transport work described in [137] was performed.

More recently, our Quantum Network experiment, which is developing an on-trap fiber cavity, has used NanoFoil trap mounting with great success. They use a cryogenic system and a puck-based pillbox to enable rapid iteration through different prototype traps. The traps are fabricated in-house but it is still best to avoid exploding valuable prototype traps. To date they have mounted two traps using the NanoFoil procedure described above and have not experienced any failures associated with thermal cycling in their cryostat.

Figure C.6: Top: the first NanoFoil-mounted ion trap used in a real experimental system, this is the detector-less detector trap used for the fast transport results reported in [137]. Bottom: recently, the Quantum Networks project has taken advantage of NanoFoil-based trap mounting in their rapid-prototyping system.



D

Quartz crystal thickness monitors

This appendix describes the quartz crystal-based thickness monitors (QCTMs) developed in the summer of 2018 as a replacement for the old-style plating sensors used to measure the production of neutral flux by beryllium ovens. The beryllium ovens consist of a length of tungsten wire which is wrapped in beryllium wire and then coiled. To produce a beam of neutral beryllium, current is run through the tungsten wire, gradually evaporating the beryllium. This type of oven is much more delicate and prone to catastrophic failures than the pillow-style ovens used for calcium and magnesium, so the QCTMs are meant to help diagnose such failures, and ideally allow current adjustments to keep neutral atomic flux levels constant as the ovens age.

D.1 Background

The QCTMs are a miniaturization of a common industrial sensor known as a thickness monitor and often based on a quartz crystal microbalance [225]. Quartz crystals have electromechanical resonances, determined by the physical properties of the crystal, which can be piezo-electrically driven. When mass is deposited on the crystal, it shifts the resonance frequency linearly as described

by the Sauerbrey equation,

$$\Delta f = -\frac{2f_0^2}{A\sqrt{\rho_q \mu_q}} \Delta m,\tag{D.1}$$

where f_0 is the initial resonant frequency, Δm is the deposited mass, A is the (active) area of the crystal, ρ_q is the density of the crystal, and μ_q is the shear modulus [225, 226]. Unfortunately, in our situation these parameters are not always easy to know, but the important factor here is that the frequency changes linearly with the deposited mass. Thus by measuring the shifts in resonant frequency, we can determine the relative amounts of beryllium coming off of an oven. Furthermore, we can measure the shifts from typical loading currents and measure ion production using a Channeltron to assess the efficiency of our loading strategies. This work is described in Appendix E.

D.2 Board design

There are several reasons we could not just put commercial thickness monitors into our systems: first, they are too large to easily fit into our pillboxes without significant modifications, and second, they are not necessarily compatible with the ultra-high vacuum (UHV) environment of an ion trapping vacuum chamber. We also wanted high sensitivity to very small amounts of deposited beryllium, which is achieved by using smaller crystals. In order to satisfy all of these requirements, we designed the sensors to fit onto a small UHV-compatible circuit board the correct size and shape to replace the lids in the standard beryllium oven mounts used in the group, as seen in Fig. D.2. This means that no modifications to the pillbox are required and we can use small, UHV-compatible components.

There are two versions of the QCTM: active and passive, where the difference is whether the electronics used to track shifts in the resonance frequency are on the board itself (active) or external (passive). The active board requires a power input, and the passive does not. There are also different configurations depending on whether you want your outputs to be multiplexed together or not (all our boards have two crystals to fit over the double beryllium ovens typically used in the group).

The circuit diagrams for all of these combinations are shown in Fig. D.1, and there are two generations of circuit boards made from these schematics.¹ The first generation (QCTM1) was the non-miniaturized, non-UHV prototype, and is easier to solder and debug. The second generation (QCTM2) is miniaturized to fit over the ovens and was manufactured to be UHV-compatible.²

D.2.1 Passive

The passive boards need only the crystal and a suitable electrical line to track the resonance. The difference between the two versions is whether the signals from both crystals come out on the same coaxial cable or not, which will depend on your preference/available feedthroughs. If you choose to have both crystals on the same output cable, it is nice to select crystals with different resonance frequencies so they can be easily distinguished. There is room on the board for an optional inductor and capacitor, which were introduced to improve coupling, but it turns out the crystal resonance is quite strong on its own and the additional components just introduce additional sources of error.

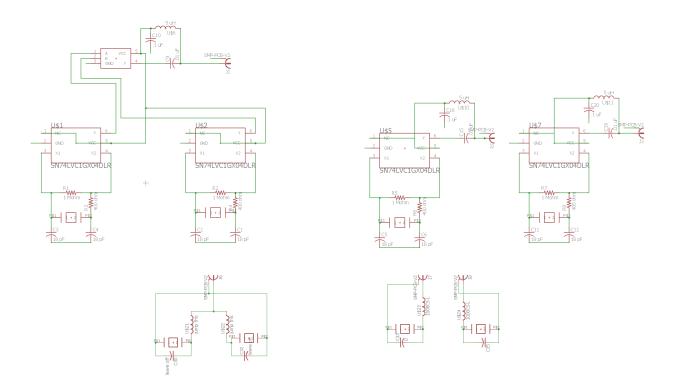
D.2.2 Active

Of the three active versions, only the top one requires separate external power, and thus more feedthroughs, making it less attractive. The other two use the same coaxial cable for their power and signal. Just like with the passive ones, the only difference is whether they share that line, and if you choose to share I recommend choosing crystals with different resonant frequencies.

¹For those in the Ion Storage group, the Eagle files can be found on jake in Ions/QuantumOne/trap_loading_test_setup/QCmonitor. For those outside the group, I can supply the files via email if you like.

²This means gold traces on Rogers board, then assembled using SAC lead-free solder.

Figure D.1: Top: Active QCTM schematics; Bottom: Passive QCTM schematics; Left: Combined outputs for two crystals; Right: Separate outputs for two crystals.



D.3 Soldering instructions & tips

D.3.1 Supplies

For those in the Ion Storage group, these supplies are located in a box in the assembly room labeled 'Hannah's Crystal Thiccness [sic] Monitors.'

D.3.1.1 Passive

- 1 board (QCTM 2)
- 2 crystals (Statek CX11 24 and/or 26.5 MHz, lidless)
- 1 or 2 SMP connectors (Molex 85305-0232)
- small inductor or jumper

D.3.1.2 Active

- 2 crystals (Statek CX11 24 and/or 26.5 MHz)
- 1 or 2 SMP connectors (Molex 85305-0232)
- \bullet 2x SN74LVC1GX04 crystal oscillator driver, DRL package
- 0402 capacitors: 4x 18 pF, 2x 0.01 uH, 2x 0.1 uH
- 0402 resistors: $2x 1 M\Omega$, $2x 350 \Omega$
- 1008 inductors: $2x + 4.7 \mu H$
- for shared line boards only: NAND

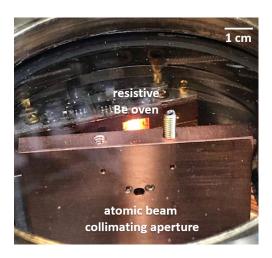
D.3.2 Crystal soldering tips

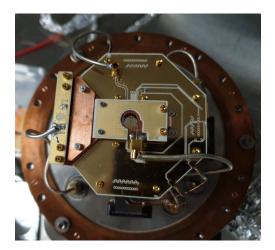
• I recommend doing the crystals last because they are delicate and on the opposite side of the board from everything else.

- The electrodes of the crystal should point towards the nearest edge of the board to be centered over the oven. For most of the boards this means putting the electrodes on the small pad and the other end of the crystal on the ground plane.
- It's difficult to solder the crystals with an iron due to the tiny size of the pads. I find it simplest to use a heatgun.
- The crystals are not as fragile as one might think. I've had good luck detaching and re-attaching them with a heatgun multiple times, and the only times I've truly broken one were when I accidentally grabbed the crystal itself with the tweezers, instead of the housing.

D.4 Testing Data

Figure D.2: Left: an active QCTM in place in the test chamber used to take the test data in this appendix as well as the beryllium oven investigations described in Appendix E. Right: a passive QCTM installed in the Triangle trap apparatus at NIST. Note that they have connected their output lines to require only one coax line out of the vacuum chamber.





D.5 Guidelines for use

The active versions need 3-5 V teed in with a bias tee (e.g., Mini-circuits ZFBT-4R2GW+) to the signal out SMA cable. The resonant frequency can be measured with a frequency counter.

Power dissipation for a single active QCTM is 60 mW, and for a double active is 80 mW. While some may worry about the additional heat load, especially in cryogenic systems, you only need to run these while you are running your oven, which is probably dissipating quite a bit more energy. Assembled boards can be tested by gently pressing a gloved thumb to the crystal surface - the warmth of your thumb should shift the resonance. In order to distinguish shifts from the heat of an oven from shifts due to plating, turn off the oven and check if the resonance comes back. This effect can be seen in Fig. D.4.

The passive QCTMs are easier to assemble but more complicated to set up: you either need to use a network analyzer or a phase-locked loop to monitor the shifts in the crystal frequency. However, the savings in vacuum feedthrough lines, power dissipation, and tiny lead-free soldering may be more than worth it.

D.6 Implementation in experiments

The Triangle trap apparatus used an active plating sensor to monitor their beryllium ovens. Over time, the sensors eventually stopped working. When the pillbox was opened they discovered that (due to unclear instructions on my part) the electrodes had been placed over the oven aperture and had been plated with beryllium to the point where they electrically shorted. However, they were able to revive both sensors by dunking the crystals into a 5% HCl solution 3-7 times for 15 seconds at a time and then rinsing with water to remove the plated beryllium and any residues.

We have also used one of the non-UHV plating sensors to determine what current was necessary to produce neutral flux after making a few modifications to past oven designs.

Figure D.3: The active QCTM from Fig. D.2 after running both ovens. The beryllium plating is the circular shadow visible on both crystals. It is round because that is the shape of the aperture above each oven in the mounting block meant to protect the electrodes (silver spots at bottom of each crystal) from becoming so plated that they short to each other. The faint rectangular outline of each crystal within its housing can be seen relatively clearly on the left crystal. The large components at the bottom edge are the inductors which hang off the back edge of the board behind the oven mount.

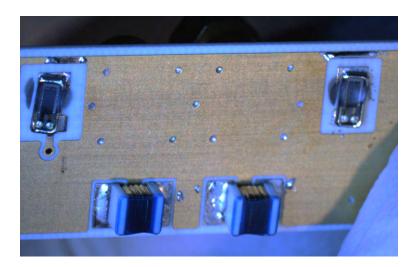
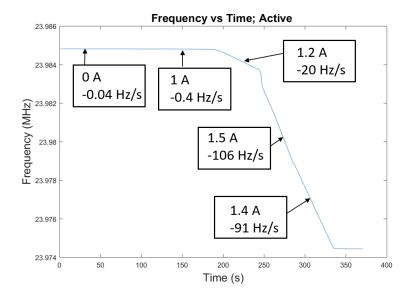


Figure D.4: Data collected from the QCTM shown in Fig. D.3 while running the beryllium oven at different currents. For reference, 1.2 A is a typical loading current, and even that relatively low flux is easily distinguishable from the baseline sub-Hz frequency drift. By comparing the resonance before and after the oven is run, you can confirm that the sensor was plated and that the measured frequency change was not due to heating of the crystal.





VECSEL source for beryllium photoionization

Portions of this have appeared in [227], which focused on demonstrating the basic elements of quantum information processing on beryllium ions using two VECSEL-based laser systems. I worked only on the photoionization system, so only the relevant portions of the paper are reproduced here.

E.1Introduction

Quantum information processing (QIP) technology based on atomic physics is steadily emerging from research laboratories and moving into commercial development [228]. Often the selection of the atomic system for a particular QIP application is determined more by the availability of suitable lasers than by atomic properties. This is especially apparent with atomic species requiring ultraviolet (UV) laser sources, where cost and reliability present obstacles. One important example is the beryllium ion (⁹Be⁺), which has several properties advantageous for QIP. The low ion mass helps to achieve high secular trapping frequencies [23], allowing for faster quantum gates and ion transport operations [136, 229], along with stronger Coulomb-mediated coupling between ions in separated-trap arrays [219, 230]. As a result, comparable high-fidelity two-quantum bit (qubit) gates can typically be implemented with less laser intensity than is needed for heavier ion species [231]. Reducing laser requirements may be critical for scaling to larger processors [23, 24, 232–234] and eventual fault-tolerant operation. Another attractive feature of ${}^{9}\text{Be}^{+}$ ions is the $\sim 1.3\,\text{GHz}$ ground state hyperfine splitting, accessible with low-cost microwave electronics and relatively simple antennas. Qubits stored in ${}^{9}\text{Be}^{+}$ hyperfine states have exhibited coherence times of several seconds [235] and single-qubit gate errors of $2.0(2) \times 10^{-5}$ [236]. Furthermore, two-qubit gates between ${}^{9}\text{Be}^{+}$ hyperfine qubits have been demonstrated with an error of $8(4) \times 10^{-4}$ [6], one of the lowest reported in any physical system to date.

Beryllium QIP experiments typically use laser light at 235 nm to photoionize neutral atoms, and at 313 nm for laser cooling, state preparation, quantum gates, and measurement. The relevant features of the ⁹Be⁺ energy level structure are shown in Figure E.1. Light at 235 nm is often generated by nonlinear frequency conversion of lasers operating at 940 nm, including titanium-sapphire lasers (both pulsed and CW) and semiconductor diode lasers [237].

Vertical-external-cavity surface-emitting lasers (VECSELs [238]) offer a promising alternative to the approaches listed above. VECSELs combine the advantages of external-cavity solid-state disk lasers with those of quantum-well semiconductor lasers, and have made considerable progress over the last decade [239]. The external cavity geometry of these lasers enables high-power and single-frequency operation with near-diffraction-limited beam quality, in a relatively compact package. The semiconductor gain material permits a wide tuning range and can be engineered for emission over a broad range of wavelengths. Thanks to the long (> 10 cm) high-Q cavity and the short carrier lifetime of the semiconductor gain medium, VECSELs operate in the so-called class A regime having cavity photon lifetimes much longer than the upper-state lifetime of the gain medium. As a result, unlike in solid-state lasers and many diode lasers, the laser dynamics are dominated by the photon lifetime in the external cavity. As a result, VECSELs produce very low intensity noise light [240] and do not suffer from relaxation oscillation. Owing to the large intracavity power and strong gain saturation, the broad spectral pedestal due to amplified spontaneous emission that is generally present in fiber lasers, external-cavity diode lasers, and tapered-amplifier systems is largely absent in VECSELs. These advantageous features have been validated by the demonstration of

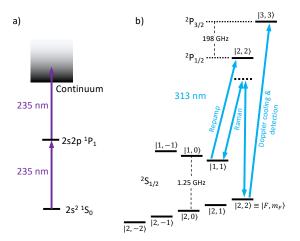


Figure E.1: a. Energy levels of neutral $^9\mathrm{Be}$ for photoionization with 235 nm light. b. Energy levels of $^9\mathrm{Be}^+$. Light near 313 nm is used for laser cooling, state preparation and detection, repumping, and multi-qubit quantum gates based on stimulated Raman transitions. Additional repump light resonant with $^2S_{1/2} |2,1\rangle \leftrightarrow ^2P_{1/2} |2,2\rangle$ is omitted for clarity.

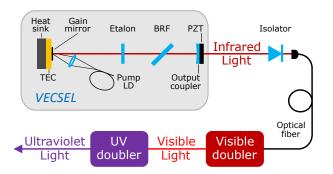


Figure E.2: Diagram of the VECSEL cavity configuration and two stages of external frequency doubling.

VECSEL-based systems for the generation and manipulation of trapped magnesium ions [241], with measured performance fulfilling the requirements of typical ion-trapping experiments. Although the benefits of using VECSELs in quantum-technology applications are gaining more attention (see for example recent publications reporting single-frequency VECSELs with sub-kHz linewidth and 140 mW single-mode output power [242], or linewidths of tens of kHz and watt-level single-mode output powers [243, 244], targeting transitions in neutral Sr [242, 243], neutral Cd [243], and neutral Rb [244]; a VECSEL with 22 mW single-mode output power before frequency-doubling to Yb⁺ ion wavelengths has also been reported [245]), single-frequency operation at wavelengths matching Be⁺ ion transitions has not been demonstrated.

Here, we demonstrate a VECSEL-based system for the generating ${}^{9}\mathrm{Be^{+}}$ ion based on a new implementation for a gain mirror operating at a fundamental wavelength of 940 nm, and external frequency conversion stages to generate laser light at 235 nm. The functionality of the UV laser is assessed in terms of its ability to create ions from neutral atoms. The work demonstrates the versatility of VECSEL technology in fulfilling the needs of ${}^{9}\mathrm{Be^{+}}$ and other ion systems, and its promise to provide practical solutions in a number of quantum-technology applications.

E.2 Laser setup

A schematic diagram of the laser system is shown in Figure E.2. The VECSEL cavity consists of a 940 nm gain mirror and an output-coupler mirror (OC, Layertec, $\sim 2\%$ transmission, 200 mm

radius of curvature, 12.7 mm diameter) spaced approximately 125 mm apart, giving a ~1.2 GHz free spectral range (FSR). To provide single-frequency operation and coarse tuning, we used the following intra-cavity elements: a Brewster-angled birefringent filter (BRF, single quartz plate, 3 mm thickness) and an etalon (yttrium aluminium garnet, 1 mm thickness). Both these elements are actively temperature-stabilized. For fine tuning and stabilization of the VECSEL output frequency, the small OC mirror is mounted on a ring-shaped piezo-electric transducer (PZT), which allows the cavity length to be adjusted. The gain mirror is optically pumped by a high-power, multi-mode diode laser emitting near 808 nm, which is fiber coupled and focused to produce a $\sim 200 \,\mu\mathrm{m}$ beam waist (radius) on the gain-mirror surface. Gaussian or super-Gaussian pump intensity profiles with a laser-mode-to-pump-spot ratio larger than 0.8 are typically preferred for single-mode operation [246]. Although the pump absorption band of the gain material is very broad, high-power diode lasers emitting near 808 nm were selected, as these are readily available and relatively inexpensive, and the gain mirror pump absorption length was optimized accordingly. The gain mirror copper heatsink is mounted on a thermo-electric cooler (TEC) for temperature stabilization and control. The TEC is water-cooled using a micro-channel heat exchanger and a low-vibration chiller (important for narrow-linewidth laser operation). The VECSEL cavity components are mounted on an Invar steel baseplate that is housed in an O-ring-sealed enclosure for stable operation. We have found that removing intra-cavity water vapor, by including a desiccant inside the laser enclosure and purging with dry nitrogen, improves frequency stability and output power at wavelengths where water absorption lines are present (for example at 940 nm).

In single-mode operation of the VECSEL, coarse-range tuning (in $\sim 80\,\mathrm{GHz}$ steps over $\sim 10\,\mathrm{THz}$) is achieved by BRF rotation, as well as by adjusting the BRF set-point temperature (0.1 THz–1 THz range with $\sim 37\,\mathrm{GHz/K}$ rate). Intermediate-range tuning is achieved by tilting the etalon and adjusting its temperature (1 GHz–100 GHz range with $\sim -3\,\mathrm{GHz/K}$ rate, with $\sim 1.2\,\mathrm{GHz}$ steps). Fine tuning is achieved by adjusting the laser cavity length using the output-coupler PZT (0.01 GHz–1 GHz with $\sim 30\,\mathrm{MHz/V}$ rate). By simultaneously tuning the etalon temperature and the laser cavity length, it is possible to achieve mode-hop-free tuning ranges over several GHz, limited

by the maximum travel of the cavity PZT.

The VECSEL gain mirror is comprised of a distributed Bragg reflector (DBR) and an active region with strain compensation layers, quantum wells (QWs), barrier/spacer layers, and a window layer, arranged for resonant periodic gain (RPG). The mirror was fabricated using solid-source molecular beam epitaxy (MBE). A general description of VECSEL gain mirror technology and design constraints related to material systems can be found for example in [239, 247], and for specific design details of these systems we refer the reader to the full paper this appendix is sourced from [227].

For proper VECSEL operation, the heat resulting from pump laser absorption must be removed and the temperature of the gain mirror must be stabilized. To this end, for the 940 nm gain mirror we used a standard flip-chip cooling method [238] where the heat flows from the active region through the DBR to the heat spreader. For this purpose, the as-grown gain mirror was first diced into $2.5 \times 2.5 \,\mathrm{mm^2}$ chips and the DBR back surface of each chip was bonded to a $3 \times 3 \times 0.3 \,\mathrm{mm^3}$ synthetic multi-crystal diamond heat-spreader. After bonding, the GaAs substrate was removed with a combination of mechanical lapping and wet etching. An ion-beam-sputtered (IBS) anti-reflection (AR) coating was applied to the top surface of the approximately $5 \,\mu$ m-thick semiconductor structure for reduced reflection of pump and laser light and for protection of the gain chip surface. Finally, the diamond back surface was soldered to a temperature-stabilized copper heat sink for efficient heat extraction.

E.2.1 940 nm VECSEL

Performance characteristics of the 940 nm VECSEL are shown in Figure E.3. With 14.5 W of pump power, the tuning range is ~ 30 nm, and the slope efficiency at 940 nm is 27(1)% for single-frequency operation. To estimate the VECSEL linewidth, we analyze the beat note signal between the 940 nm VECSEL output and that of a free-running titanium-sapphire (TiS) laser (nominal linewidth $<50\,\mathrm{kHz}$), with the VECSEL frequency-offset locked to the TiS laser [248]. From the spectral width of the beat signal, we determine the linewidth of the VECSEL to be $<100\,\mathrm{kHz}$.

This is considerably less than the linewidth of relevant atomic transitions and sufficiently narrow that frequency fluctuations will not be converted to significant amplitude fluctuations by subsequent resonant frequency doubling stages.

E.2.2 Doubling to 470 nm and 235 nm

The output of the 940 nm VECSEL is coupled into a high-power polarization-maintaining (PM) optical fiber that delivers light to the first of two second-harmonic generation (SHG) enhancement cavities. This 940 nm to 470 nm frequency-doubling cavity is locked to the single-frequency output of the VECSEL using the Pound-Drever-Hall (PDH) method [249], with the sidebands generated via electrico-optic modulation of the infrared light. Following the design by Lo and colleagues [237], frequency doubling to 470 nm is implemented using a periodically-poled potassium titanyl phosphate (PPKTP) nonlinear optical crystal (Raicol Crystals, 20 mm length). The performance of the VECSEL-driven 470 nm source is shown in Figure E.5 a. With 0.48(2) W at 940 nm we obtain 0.27(1) W at 470 nm. While this is more than sufficient for our application, we note that the 940 nm power to the doubling cavity could be further increased. However, cavity-locking instabilities due to apparent thermal lensing in the crystal prevent reliable locking above ~ 0.50 W at 940 nm [237, 250].

The second stage of frequency doubling, from 470 nm to 235 nm, is implemented using an enhancement cavity design [237], adapted from an earlier 313 nm design [199], that uses a Brewster-angled beta barium borate (BBO) crystal. This doubling cavity is locked using the technique developed by Hansch and Couillaud [251]. The performance of the 235 nm frequency doubling stage is shown in Figure E.5 b. With 0.27(1) W at 470 nm we obtain 54(3) mW at 235 nm, and a maximum power ratio of approximately 20 %. For context, this UV power is approximately 50 times greater than the power we typically use to photoionize beryllium atoms to load them into an ion trap.

E.3 Photoionization testing

Photoionization with the 235 nm VECSEL system was tested on a thermal beam of neutral beryllium atoms. The thermal beam was generated by resistively heating a length of beryllium

wire spiral-wound onto tungsten support wire. The beam was weakly collimated using a 2.5 mm diameter aperture. Immediately downstream from the aperture, the atomic beam intersected with a perpendicular 235 nm laser beam near-resonant with the 1S_0 to 1P_1 transition and focused to an intensity of $\sim 80(20) \,\mathrm{kW/m^2}$ at the center of the atomic beam. For comparison, the saturation intensity of this transition is $\sim 8.7 \,\mathrm{kW/m^2}$ [252]. Ions are produced by a two-photon process (see Figure E.1 a). On resonance, the first photon excites the neutral atom to the 1P_1 state and a second photon excites the electron to the continuum. These ions are counted using a Channeltron electron multiplier (CEM) (Photonis Magnum 5901 Electron Multiplier) with a bias potential of -1.7 kV.

With the neutral atomic beam flux held constant, we record the ion count rate as a function of the VECSEL frequency to obtain a photoionization lineshape (Figure E.6). Maximum photoionization rates are measured at a VECSEL frequency of 319.0200(6) THz, corresponding to 1,276.080(2) THz in the UV, consistent with a recent precision measurement of the 1S_0 to 1P_1 transition [252]. The central feature of the photoionization lineshape includes contributions from the natural linewidth of 87(5) MHz [252], power broadening by a factor of $\sim 2.4(3)$, and 1.3(2) GHz of residual Doppler broadening from imperfect collimation of the atomic beam [253]. A Voigt model [253] including only these mechanisms shows good agreement with the central feature. The origin of the weak off-resonant photoionization (the broad pedestal feature of the lineshape) has not been investigated.

E.4 Doppler spectroscopy

The same spectroscopy was performed with the 235 nm light propagating counter to the atomic flux rather than perpendicular, in order to measure the approximate velocity distribution of the atomic beam. Fitting the resulting curve with a thermal distribution estimates a temperature of 1290(180) K corresponding to an approximate peak velocity of 1540(110) m/s. By correlating the Doppler-free spectroscopy with the count rates from the plating sensors, we can estimate that < 10% of the neutral beryllium beam is ionized in the Doppler-free configuration, and cross-referencing with the thermal distribution and knowing the depth of our ion traps we can estimate that < 15% of the

ionized atoms are within the trappable velocity range. These measurements confirm our suspicions that loading surface traps from thermal sources is inefficient and produces many untrappable charged particles, which can collide with surfaces and contribute to stray fields. Alternative strategies, such as MOT loading [129] could potentially be helpful for species such as beryllium which feature deep-UV photoionization light as well as high thermal velocities at the temperatures necessary to create atomic vapor beams.

E.5 Conclusion

In summary, we have presented two widely-tunable VECSEL-based laser sources capable of implementing tasks for quantum information processing experiments involving trapped ⁹Be⁺ ions. The first system generated up to 2.4 W single-frequency light at 940 nm and was frequency doubled twice to generate 235 nm light for photoionization of neutral Be. We used this light to perform spectroscopy on the thermal atomic beams produced by our beryllium ovens to help improve understanding of our beryllium loading strategies.

Over more than eight months of continuous use at an output power of about 1 W, we observed long-term power drifts at 940 nm of less than 10%. In May 2021 the M Squared Ti:Sapph that was previously used to generated photionization light for all four beryllium experiments in the group died and was replaced with this VECSEL. The 940 nm light from the VECSEL was threaded into the M Squared doubler and the resulting 470 nm light is distributed to multiple UV doublers for the different trapped-ion setups; each doubler produces a few mW at 235 nm, sufficient for producing $^9\mathrm{Be}^+$ ions in our experiments. The frequency of the VECSEL is locked to a wavemeter.

We note that the large tuning range of the 940 nm VECSEL system presented in this work could allow for generation of light near 313 nm by third harmonic generation [254]. The inherent power scalability of the VECSEL design should allow the generation of higher output power [238], which is desirable for mitigating spontaneous emission errors associated with quantum-logic gates using far-detuned stimulated Raman transitions[231]. Further improvements to overall system efficiency could be achieved using intra-cavity second harmonic generation [255].

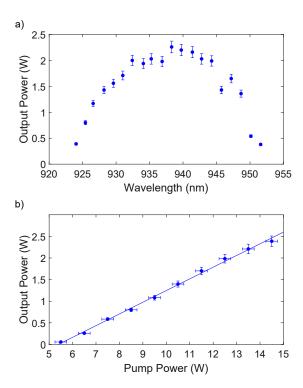


Figure E.3: a. Single-mode output power at $940\,\mathrm{nm}$ as a function of wavelength with $14.5\,\mathrm{W}$ pump power, obtained by rotating the BRF and tuning etalon temperature at each wavelength. b. Output power at $940\,\mathrm{nm}$ versus $808\,\mathrm{nm}$ pump power. A linear fit to the data (solid line) gives a slope efficiency of $27(1)\,\%$. Power measurement uncertainty is $5\,\%$ and the wavelength accuracy uncertainty is $0.002\,\mathrm{nm}$.

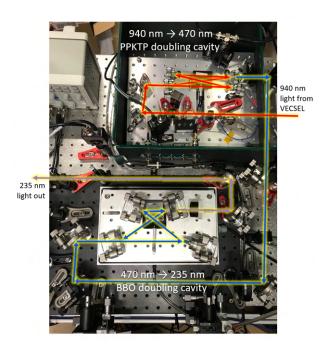


Figure E.4: Annotated photograph of the two doubling cavities used to convert the 940 nm light to the 235 nm light needed for beryllium photoionization. The first cavity uses periodically-poled potassium titanyl phosphate (PPKTP) doubling crystal, which must be temperature-controlled for ideal doubling performance. The second cavity uses a beta barium borate (BBO) doubling crystal which is phase-matched to the incoming light by careful adjustment of the crystal angle in the cavity.

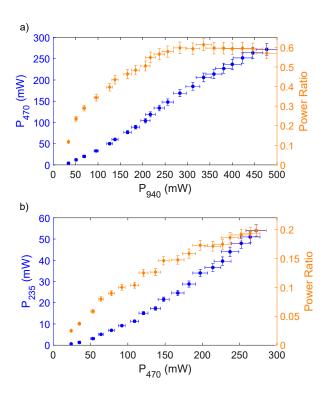


Figure E.5: a. The $470\,\mathrm{nm}$ power (blue circles), and the $470\,\mathrm{nm}$ to $940\,\mathrm{nm}$ power ratio (orange diamonds), versus the input power at $940\,\mathrm{nm}$. b. The $235\,\mathrm{nm}$ power (blue circles) and the $235\,\mathrm{nm}$ to $470\,\mathrm{nm}$ power ratio (orange diamonds) versus the input power at $470\,\mathrm{nm}$. Uncertainty in all power measurements is $5\,\%$.

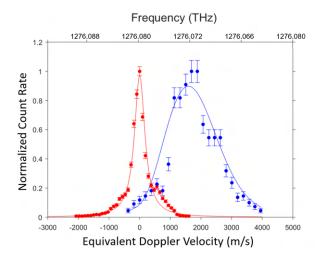


Figure E.6: Count rate measured by the CEM as a function of the quadrupled VECSEL (UV) frequency detuning from 1,276.080(2) THz. The central peak is overlaid with a Voigt model. The red points show the Doppler-free spectroscopy, which exhibits a Doppler-broadened linewidth of 1.3(2) GHz of Doppler broadening due to the imperfect atomic beam collimation. The red line is a Voigt model with the expected power and Doppler broadening and demonstrates good agreement with the data. The blue points show the data from spectroscopy with the 235 nm light counterpropagating with the atomic beam. The peak is shifted Count rate measurement uncertainty is estimated to be 4%.

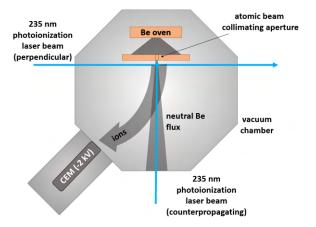


Figure E.7: Diagram of the chamber used to test the beryllium photoionization efficiency of the 235 nm light, featuring a resistive beryllium oven, CEM, and two possible configurations for the laser light: one that is nominally Doppler-free (modulo imperfect collimation of the atomic beam) and a second with maximum Doppler sensitivity. The comparison of the ionization peaks from these two configurations allows us to estimate the thermal velocity of the atomic beam.