

TRAPPED ION COOLING, HEATING, AND THERMOMETRY

A. J. Rasmusson

Submitted to the faculty of the University Graduate School

in partial fulfillment of the requirements

for the degree

Doctor of Philosophy

in the Department of Physics,

Indiana University

May 2024

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the
requirements for the degree of Doctor of Philosophy.

Doctoral Committee

Philip Richerme, Ph.D.

Herbert Fertig, Ph.D.

Gerardo Ortiz, Ph.D.

Babak Seradjeh, Ph.D.

May 7, 2024

Copyright © 2024

A. J. Rasmusson

Journey before destination

Acknowledgments

I am indebted to the many mentors, supporters, and colleagues I've had the pleasure to work with over the course of this work. My deepest appreciation goes to my advisor Phil Richerme for his unmatched mentorship, priceless advice, scientific example, and patience. It has been a joy to learn how to be a scientist from you. I hope that one day I can write an abstract as well and as quickly as you.

I will always look up to Marissa D'Onofrio and Yuanheng Xie—the first graduate students. Thank you for your time, patience, and friendship; for building the lab; and for mentoring and training me in how to do experimental physics. Jiafeng Cui and Antonis Kyprianidis, your insight and leadership as postdocs has been invaluable, and your jokes and laughter have made working in the lab delightful. Frank Schroer IV, Ilyoung Jung, and Wes Burkle, thank you for your patient listening and thoughtful critiquing of my crazy ideas over lunch. You've made the lab a fun, uplifting, and creative place to work. I look forward hearing of all your accomplishments as you complete your graduate studies. I've had the pleasure to get to know (and work with) a variety of undergraduate students and even high schoolers: Aaron McCann, Dane Cross, Evangeline Wolanski, Paula Madetzke, Nico Goldstein, Harvey Campos, Josh Apanavicius, Louis Farenci, Mathesen Skillman, Cass Egyhazi, Alex Alani, Advait Iyengar, and Brooke Liao. Each of you have left a special mark in the mosaic of our group—thank you. I have admired the common attributes you all share: curiosity, willingness, and hard work.

I thank the members of this committee: Herb Fertig, Gerardo Ortiz, and Babak Seradjeh for your patient instruction over the years and continued interest and support of this work. I owe special thanks to the undergraduate physics experience at BYU. Those early years of growth and struggle in physics rest fondly in my heart. Jean-Francois Van Huele, thank you for taking a chance on me, sending me to conferences, and growing my interest in quantum physics even though I had little experience. Thank you to David Allred and Dallin

Durfee for letting me play with lasers and vacuum systems originally sparking my interest in experimental trapped ion research (why I am here at IU). Lastly, I wish to thank Traci Neilsen for teaching me how to write. I still read your notes!

Mom and Dad, thank you for letting me do chemistry experiments in the garage and watching cheesy sci-fi shows and science documentaries. I believe these small and silly acts seeded my love of nature and physics. And, of course, thank you for everything else. I lovingly acknowledge my late Grandpa Chappell whose timely gift of an audio lecture on quantum physics and special relativity abruptly set my feet on the path of quantum physics.

My greatest debt is to my wife Marysa. Your support and sacrifice have been immeasurable. Neither of us understood what “doing a PhD” meant all those years ago. However, the trials have strengthened us, and the disappointments seasoned us. You are a pillar of strength; your smile, a buoy. At this point it should be obvious that this work would never have been completed without your encouragement.

To Grandpa Chappell, my parents, and Marysa

TRAPPED ION COOLING, HEATING, AND THERMOMETRY

Trapped ion experiments are a high-fidelity platform for studying quantum physics. Individual atoms are photoionized, held in free space by electric fields, and typically form a linear constellation. Lasers cool the ions to near their zero-point energy—the absolute ground state of motion—in preparation for driving high-fidelity quantum operations. This work presents a holistic study of the cooling, heating, and thermometry of trapped ions which are foundational to their use as an experimental quantum platform.

Ions are readily laser cooled to near their motional ground state. However, cooling can take 99% of experimental runtime; traditional methods do not work for extended ion wavepackets; and there has been no generalizable model of the process in the literature. We address these points by introducing a new framework which models the entire cooling process on a graph. The fastest possible cooling sequences are computed, and their improvement is experimentally verified compared to traditional methods.

Measuring the average motional energy of an ion is also a common trapped-ion protocol. However, we observe a fundamental assumption of standard thermometry techniques can be violated after ground state cooling. Consequentially, as much as an order of magnitude underestimate can be observed. We therefore develop and experimentally validate an improved thermometry method which accurately measures the average motional state of an ion after ground state cooling.

Finally, we study the heating of trapped ions during measurement. Heating reduces the fidelity of quantum operations, yet new protocols are proposing repeated measurements throughout an experiment which may cause heating. An ion’s quantum state is measured by state-dependent fluorescence. We observe that ions in the fluorescing state rapidly scatter thousands of photons which induces a large heating rate roughly 30 times faster than ions in non-fluorescing states. We introduce a new framework that accurately reproduces the experimental results and provides a unified description of ion heating.

Contents

List of Tables	xiii
List of Figures	xiv
1 Introduction	1
1.1 Trapped Atomic Ions for Quantum Science	2
1.2 Dissertation Outline	3
2 Theoretical Background	5
2.1 Paul Trap Theory	5
2.1.1 Confining Fields	6
2.1.2 Micromotion	7
2.2 Trapped-ion Motional States	8
2.2.1 Quantum Harmonic Oscillator	9
2.2.2 Thermal Distribution	10
2.3 Singly ionized Ytterbium 171	12
2.3.1 Energy Level Diagram	13
2.3.2 Atomic Transitions at 369.5 nm	14
2.4 Coherent Atom-laser Interactions	24
2.4.1 Rabi Flopping	25
2.4.2 Stimulated Raman Transitions	28
2.4.3 Spin-motion Coupling	30
2.4.4 Mølmer-Sørensen Interaction	44
3 Apparatus	50
3.1 Rod and Blade Traps	50
3.1.1 Electrode Geometries	51

3.1.2	Voltage Delivery	53
3.2	Vacuum System	56
3.3	Laser Systems	57
3.3.1	Frequency Reference	57
3.3.2	399 nm Laser System	58
3.3.3	935 nm Laser System	59
3.3.4	369.5 nm Laser System	60
3.3.5	355 nm Laser System	63
3.3.6	Laser Beam Delivery	71
3.3.7	Microwave Horn	72
3.4	Imaging System	73
3.4.1	Light Collection	73
3.4.2	Photomultiplier Tube Detection	75
3.4.3	EMCCD Camera	76
3.5	Real-Time Control	78
3.5.1	Advanced Real-Time Infrastructure for Quantum physics (ARTIQ)	79
3.5.2	ARTIQ Dashboard	84
3.5.3	ARTIQ Data Flow	87
3.5.4	Subartiq_lib	87
3.5.5	Kasli Crate	91
3.5.6	Advanced Use Cases	95
3.6	Typical Experiment Workflows	102
3.7	Troubleshooting	103
3.7.1	Trapping the First Ion	103
3.7.2	General ARTIQ Advice	104
3.7.3	Latencies and Odd Early Time Behavior	104
3.7.4	Motional Coherence	104

3.7.5	Lasers	105
3.7.6	RF Breakdown and DC Shorts	106
3.7.7	Miscellaneous	107
4	Optimized Pulse Sideband cooling	108
4.1	Resolved Sideband Cooling Theory	109
4.1.1	Alternative Cooling Techniques	111
4.2	Optimized Pulsed SBC Protocols	112
4.2.1	Graph-Theoretic Description of Pulsed Sideband Cooling	112
4.3	Summary and outlook	120
5	Trapped-ion Thermometry	121
5.1	Introduction	121
5.2	Existing Methods	122
5.2.1	The Ratio Method	122
5.2.2	The SVD Method	123
5.3	Modeling Post-SBC Distributions	124
5.4	Time-averaged Thermometry	125
5.5	Experimental Thermometry	127
5.5.1	Thermal Distribution	128
5.5.2	Sideband Cooled Distribution	130
5.6	Conclusion	134
6	Measurement-induced Heating	135
6.1	Introduction	135
6.2	Ambient Heating	136
6.2.1	QTT Model	139
6.3	Measurement-Induced Heating	141
6.4	Long-time Behavior	144

6.5 Discussion and Outlook	147
7 Conclusion and Outlook	148
A Unified heating formalism	151
A.1 Derivation of Continuous Heating Rate from QTT Model	151
A.2 Derivation of Discrete Heating Rate from QTT Model	152
A.3 Semi-classical Laser Heating	154
B Semi-classical Quantum Trajectory Numerics For Ion Heating	157
B.1 Fluctuating Fields	157
B.2 Discrete Photon-recoil	158
B.2.1 Emission Pattern	159
B.3 Determination of Double Thermal Distribution	161
Bibliography	164
Curriculum Vitae	

List of Tables

2.1	Phases needed to measure in the x , y , and z bases	36
3.1	Static code metrics for subaritm_lib and our ARTIQ repository	89

List of Figures

2.1	$^{171}\text{Yb}^+$ energy level diagram	15
2.2	Detection and optical pumping protocol schematics	16
2.3	PMT detection histogram	17
2.4	Optical pumping decay	18
2.5	Laser cooling schematic	19
2.6	Raman transition energy level diagram	29
2.7	Sideband scan	32
2.8	Carrier Rabi oscillation	34
2.9	Carrier Rabi oscillation measured at different bases	35
2.10	Ramsey oscillation	37
2.11	Sideband cooling protocol schematic	39
2.12	Ratio method thermometry	41
2.13	SVD method thermometry	43
2.14	MS energy level diagram schematic	45
2.15	MS oscillation	47
2.16	MS parity oscillation	49
3.1	Rod trap diagram	51
3.2	Blade trap diagram	53
3.3	Voltage Delivery Schematic	54
3.4	Diode laser optical layouts	58
3.5	369.5 nm optical layout diagram	61
3.6	355 nm optical layout diagram	64
3.7	355 nm warm up	66
3.8	Beanote lock circuit diagram	68
3.9	Beam delivery optical layout diagram	71

3.10	Horn circuit diagram	72
3.11	Imaging system schematic	74
3.12	EMCCD images with Doppler leakage	77
3.13	Minimal working example of an ARTIQ experiment	83
3.14	ARTIQ dashboard	86
3.15	Subartiq.lib folder structure	90
3.16	Kasli crate	91
3.17	Advanced use cases: negative delay	99
3.18	Advanced use cases: Phaser pulse shaping	101
4.1	Graph representation of sideband cooling	113
4.2	Optimized sideband cooling matrix elements	114
4.3	Sideband cooling protocol comparison	117
5.1	Comparison of different thermometry techniques at Doppler limit	129
5.2	Simulation of ratio method underestimation of $\bar{n}$	131
5.3	Comparison of different thermometry techniques after sideband cooling	133
6.1	Ambient heating	139
6.2	Measurement-induced heating	142
6.3	Measurement-induced heating at long times	146
B.1	QTT final trajectories for continuous heating source	158

Chapter 1

Introduction

The invention of ion trapping has revolutionized the study of atomic and quantum physics. Before the 20th century, atomic properties were discovered by probing ensembles of atoms such as a gas, liquid, or solid. However, one must assume an individual atom’s behavior from the ensemble statistics which limits precision and confuses competing natural processes. Starting in the 1950s, Wolfgang Paul and Hans Dehmelt started to “do something new in order to see something new” [Georg Christoph Lichtenberg [\[1\]](#)]. They developed the ion trap technique which can trap individual charged particles such as ionized atoms with unprecedented control. This provided direct access to individual atoms, as well as subatomic particles, for study and precision measurement. In 1989, Dehmelt and Paul shared the Nobel Prize with Norman Ramsey for their work [\[2\]](#).

By placing an ion trap in a vacuum chamber, a single ion could be held for hours if not days “so you could study [it] to your heart’s content” [Dehmelt [\[3\]](#)]. The atomic, molecular, and optical (AMO) physics community found this to be irresistible [\[4, 5, 6, 7\]](#). It sparked the interest of David Wineland, a graduate student of Norman Ramsey who was working on masers—a predecessor of the laser. After which, Wineland worked under Dehmelt as a post-doctoral researcher.

In 1975, Wineland joined the National Bureau of Standards (now the National Institute of Standards and Technology (NIST)) and, with his colleagues, started the ion storage group. They combined the extreme isolation of an ion trap in vacuum with the exquisite precision of lasers to study and manipulate individual ions [\[8, 5, 9\]](#). A key challenge for precision measurements was how to cool trapped ions. Wineland and colleagues first demonstrated a cooling technique using lasers in the late 1970s [\[10\]](#). (In fact, Peter Toschek’s group in Heidelberg, Germany submitted a manuscript for publication on the same topic a day

ahead of Wineland [9].) Combining laser cooling with ion trapping meant individual atoms could be trapped at temperatures near absolute zero, making trapped ions one of the most controllable and noiseless quantum platforms. Wineland continued on this path ultimately receiving a Nobel Prize in physics for “measuring and manipulation of individual quantum systems” [9] and sowing the seeds of what is now the modern trapped ion laboratory.

1.1 Trapped Atomic Ions for Quantum Science

Richard Feynman in his 1981 lecture “Simulating physics with computers” speculated that, “with a suitable class of quantum machines you could imitate any quantum system, including the physical world” [11]. These thoughts are often considering the starting point of modern quantum simulation: where quantum experiments (“machines”), such as trapped ions, are manipulated to imitate other quantum systems. Progress on this theoretical idea became fruitful in the late 1990s by pioneers such as David Deutsch [12], Seth Lloyd [13], Peter Shor [14] and many many others. The rise of trapped ion capabilities [5] and these new theoretical proposals [15] coincided, and Feynman’s original idea grew into a new field called quantum information science: the study of how information can be stored and processed using the principles of quantum physics.

In the last three decades, physicists have pushed trapped ion experiments to the quantum limit and repeatedly demonstrated that they make a powerful platform for studying quantum physics and quantum information science. Trapped ion experiments have led almost every major experimental breakthrough in quantum information science. The first entangling operations [16, 17] between two quantum bits—qubits—was performed by trapped ions, and today they still have the highest fidelities [18, 19, 20, 21, 22, 23]. As of this writing, and to the author’s knowledge, trapped ions hold the records for the lowest quantum state measurement errors [24, 25, 26], longest quantum coherence—a measure of how “quantum” something stays over time—[27, 28], and highest fidelity single qubit manipulations [25, 23]. Quantum error correction has also begun its first demonstrations [29, 30, 31]. These building

blocks amount to a potentially powerful *quantum* computer.

The future is looking even brighter for trapped ions. Hardware is shrinking in size through microfabricated electrode structures and integrated photonics [32, 33, 34, 35, 36, 37]. An increasing effort has been made to network trapped-ion qubits using photons [38, 39, 40]. New and potentially better ways of manipulating trapped-ion qubits such as laser-free gates [25, 41, 21] as well as new ion species and encodings [42, 24, 43, 44] are being explored. These points are meant to pique interest but are by no means a complete listing. There are many additional efforts by many other groups.

1.2 Dissertation Outline

Wolfgang Paul began his Nobel lecture describing the two quintessential elements of experiment [1]:

Experimental physics is the art of observing the structure of matter and of detecting the dynamic processes within it.

This work is the later, an uncovering of the dynamic processes of trapped ions. More specifically, this work will focus on the quantum motion of trapped ions: how to efficiently cool the motion down to its zero-point energy; new methods and considerations for measuring the motional energy of a trapped ion; and finally, the physics of how ion motion increases, also known as heating.

Trapped ions have “motion” which is in quotes because it is a somewhat abused term in the field. Trapped ions are not classical objects: an ion is not a ball oscillating back and forth in a bowl. Trapped ions are complex atomic systems with electronic orbital wavefunctions. The “motion” of an ion may mean the center of mass of the ionized atom which can follow classical physics trajectories. It could also mean the more quantum description of the valence electron’s orbital wavefunction—which at the zero-point energy is stationary despite its “temperature” not being at 0 Kelvin. This description will be used throughout the work. And, it is this wavefunction “motion” that enables the two-qubit quantum gate operations

necessary for universal quantum computing or quantum simulation [45, 46, 47]. Ion motion is a key and foundational quantum degree of freedom to its success. This work seeks to model and improve the cooling of ion motion, quantify ion motion accurately, and predict the increase of ion motion.

Chapter 2 starts things off with the theoretical background needed for the presented studies: Paul traps, ion motion, the ion of choice, and how lasers manipulate the quantum state and motion of the ion. Derivations are fairly parse with enough length to walk through the main points and mathematical steps without reinventing the literature. Experimental data accompanies the presented theory where possible. Chapter 3 details the experimental hardware side of this work. As this is primarily an experimental work, Ch. 3 is more technical and meticulous. Particular attention is given to the control software and hardware of our apparatus, which has not yet been thoroughly treated in a trapped ion thesis (to my knowledge) and which I spent a considerable effort developing. The remaining chapters are the experimental studies. Chapter 4 details an efficient computational model which we used to optimize laser protocols used to cool ions to their ground state. Chapter 5 reports our new technique for measuring the motional energy of an ion, and we experimentally observe how previous methods critically underestimate the motional energy. Finally, Ch. 6 rounds out the set of studies by measuring and modeling ion heating. We present the first unified model of ion heating adaptable to most any heating source, and we report the first measurement of how quickly an ion heats when its quantum state is measured. Concluding remarks are made in Ch. 7.

Chapter 2

Theoretical Background

When the quintessential physicist sees a challenge, they throw their favorite approximations at it and arrive at a fairly accurate back-of-the-envelope calculation. With enough luck, factors of 2 and π may come and go, but the right order of magnitude emerges at the end. This chapter will be a slightly more technical and mathematical than that since there is a lot more room to write than the back of an envelope.

The chapter begins in Sec. 2.1 with the mathematical basics of how to trap an atomic ion. The resulting quantum motional states are discussed in Sec. 2.2 with a focus on thermal states as this will come up again in all the experiment studies Chs. 4-6. The internal states of the ionized atom used in this work, Ytterbium, are introduced in Sec. 2.3 along with experimental procedures used to manipulate these states. Finally, how the quantum states of trapped ions can be manipulated including generating entangled states is outlined in Sec. 2.4 accompanied by measured data where possible.

2.1 Paul Trap Theory

In 1989, Wolfgang Paul and Hans Dehmelt shared the Nobel Prize with Normal Ramsey for developing “the ion trap technique” [2]. The RF, or Paul, trap can confine individual atoms for many days if not months in pristine isolation [1, 48]. Control at the atomic and subatomic level has found widespread use in physics from atomic clocks [49, 50] to quantum computing [7, 23] and even commercialized mass spectrometry applications [51]. Below, is an abbreviated treatment of how a charged particle, such as an atomic ion, is trapped in free space following Ref. [5].

2.1.1 Confining Fields

The Paul trap confines charged particles in an effective three-dimensional harmonic potential using electric fields. Earnshaw’s theorem, however, states that this is impossible for solely electrostatic fields. To get around this constraint, Paul traps use static and oscillating electric potentials. See Sec. 3.1 for diagrams and physical realizations.

For the specific case of the linear Paul trap, an oscillating potential of frequency Ω_{RF} and strength V_0 is driven on two long rods (taken to be infinite in the ideal limit) opposite each other in space by a distance $2r$. General wisdom suggests that the closer the rods follow a quadrupole the better. The oscillation potential near the geometric center of the electrodes, up to geometric factors, is approximately

$$V_{\text{RF}}(x, y, t) = \frac{V_0 \cos(\Omega_{\text{RF}}t)}{2} \left(1 + \frac{x^2 - y^2}{r^2} \right) \quad (2.1)$$

and will later be shown to drive a harmonic pondermotive potential in the x and y directions [5].

The static potential of strength U_0 and geometric factor κ is applied to endcap electrodes resulting in the potential

$$V_{\text{s}}(x, y, z, t) = \kappa U_0 \left[z^2 - \frac{1}{2}(x^2 + y^2) \right] . \quad (2.2)$$

Ideally, confinement along the z direction is solely controlled by the static potential on the endcaps, and that same potential will have an “anti-confining” effect on the x and y confinement. This is a fundamental to satisfying Maxwell’s equation for electrostatic fields and can be seen in Eq. (2.2) by the negative sign of the y and z terms.

To find the effective potential trapping a particle of mass m and charge Q , the classical equations of motions are solved for using the Lagrangian of the system. The equations of motion take the form of Mathieu differential equations and are solved in Ref. [5] Eq. (3).

The solution to the position of the ion shows harmonic motion in the radial directions

$$u_i(t) = A_i \cos(\omega_i t + \phi_i) \left[1 + \frac{q_i}{2} \cos(\Omega_{\text{RF}}) \right] + O(q_i^2) \quad (2.3)$$

where $u_i = x$ or y , A_i and ϕ_i depend on initial conditions, the secular harmonic frequency $\omega_i = \omega_x$ or ω_y , and the Mathieu “q”-parameter $q_i = q_x = 2QV_0/(m\Omega_{\text{RF}}^2 r^2) = -q_y$.

The primary radial motion in Eq. (2.3), called secular motion, is harmonic motion. The axial confinement is also harmonic resulting in a 3-dimension harmonic confinement for the ion with frequencies

$$\omega_x = \omega_y = \sqrt{\frac{Q}{m} \left(\frac{qV_0}{4d_0^2} - \frac{\kappa U_0}{z_0^2} \right)} \quad (2.4)$$

$$\omega_z = \sqrt{\frac{2Q\kappa U_0}{mz_0^2}} \quad (2.5)$$

typically on the order of 1 - 10 MHz [5, 52]. In practice, a small asymmetry is introduced between ω_x and ω_y to break the degeneracy, avoiding low-frequency rotations and allowing laser cooling to cool all three axes with one laser [53]. Additionally, the positions of the trapped ions can be manipulated by adjusting these secular frequencies. In previous studies, we used this feature to trapped ions in new two-dimensional (2D) geometries which may allow for more qubits in a given trap and perform quantum simulations of 2D systems [54, 55].

2.1.2 Micromotion

A secondary fast motion in Eq. (2.3), called micromotion, is driven at the RF frequency Ω_{RF} with amplitude $q_i/2$ to first order. It can affect atom-laser interactions by adding sidebands to atomic line shapes, and it is generally considered problematic to achieving high fidelities [56]. (Although it has found use in “hiding” ions to reduce crosstalk [57, 58, 59]). Typically, $q_i \sim 0.1$ or smaller, so this motion is low amplitude and can be negligibly small if the ion is near the trap center. The micromotion amplitude will increase linearly with distance from

the trap center [52], so most traps include additional electrodes with electrostatic fields to fine tune the ion position to minimize micromotion.

Experimentally, micromotion can be detected in many ways: ion position changes [56], atomic lineshapes [56, 60, 61], photo-correlation [60], parametric excitation [60], and many more. To first order, ion position changes are quickest. First, setup a monitor of the ion position continuously with a camera (laser cooling light is also on). Second, change the RF voltage V_0 . If the ion moves, then it must not be centered at the RF null. Adjust the ion position and repeat until the ion does not move. This method works quickly and efficiently for the two dimensions that camera can image but is not as precise as other methods. Measuring the micromotion sidebands is a common method for precisely estimating micromotion, though depending on what transitions are being used this can be limited to only certain axes. We will do micromotion sideband drives as a second order correction if we need a greater micromotion reduction. We briefly used photon-correlation measurements, but they were very lengthy compared to other methods.

2.2 Trapped-ion Motional States

Trapped ions are bound by an effective 3-dimensional quantum harmonic oscillator potential to good approximation. Since trapped ions are held in extreme isolation and at low motional energy—possibly at the zero-point energy—the quantum mechanical wave function of the ions is evident. Ions can be so isolated that superposition states of their spatial wavepackets are coherent for milliseconds as well as other nonclassical motional states such as squeezed states [62]. A quantum treatment of the motional states is necessary, and it underpins laser-ion interactions needed for generating entanglement (see Sec. 2.4.4).

2.2.1 Quantum Harmonic Oscillator

Trapped ions are confined in an effective harmonic oscillator to good approximation with the quantum harmonic oscillator Hamiltonian given by

$$H_{\text{qho}} = \hbar\omega(\hat{n} + 1/2) \quad (2.6)$$

with the fundamental constant $\hbar$, oscillation frequency ω , and number operator $\hat{n} = \hat{a}^\dagger \hat{a}$ with $\hat{a}^\dagger$ and $\hat{a}$ being the usual harmonic oscillator creation and annihilation operators [63]:

$$a = \sqrt{\frac{m\omega}{2\hbar}} \left(\hat{x} + i \frac{\hat{p}}{m\omega} \right) \quad (2.7)$$

$$a^\dagger = \sqrt{\frac{m\omega}{2\hbar}} \left(\hat{x} - i \frac{\hat{p}}{m\omega} \right) . \quad (2.8)$$

This definition assumes one spatial axis which is a common treatment as laser operations will typically only couple to one spatial direction. For trapped ion systems, the spatial and momentum eigenstates of the H_{qho} are not the focus of theoretical treatments. The Fock basis states $|n\rangle$ with the usual eigenvalues $H_{\text{qho}}|n\rangle = \hbar\omega(n + 1/2)|n\rangle$ are more often used. Additionally, it is also useful to track the complex phase space coordinate of the i th ion and the k th mode

$$\alpha_k(t) = \sqrt{\frac{m\omega_k}{2\hbar}} \hat{x}_i(t) + \frac{i}{\sqrt{2m\omega_k\hbar}} \hat{p}_i(t) \quad (2.9)$$

for two-qubit gate drives [64] (see Sec. 2.4.4) and heating rate models (see Ch. 6). The average motional state is related to the magnitude of α by $\bar{n} = |\alpha|^2$ [63].

For N trapped ions, there will be N modes which play a role in a given quantum simulation experiment (see Sec. 2.4.4). The motional Hamiltonian along one axis for N ions would be given by a sum over all motional modes

$$H_{\text{qho}} = \sum_k^N \hbar\omega_k(\hat{n}_k + 1/2) . \quad (2.10)$$

with different ions contributing to a given mode with different amplitudes [65]. For many ions, the vibrational frequencies of each mode ω_k and the participation of each ion in each mode can be computed numerically based on the ion positions and the Coulomb force between ions [66].

2.2.2 Thermal Distribution

There are a variety of quantum motional states, which have been and are currently being explored [62, 67] in a wide array of uses from quantum error correction [68, 69] to quantum metrology [70, 71] to speeding up two-qubit gate operations [72, 73]. This work focuses on thermal states and their foundational connection to ion cooling, heating, and thermometry.

The thermal state can be considered a probability distribution across all the motional Fock states $|n\rangle$ following Gibbs statistics. Starting from statistical mechanics, we can derive the thermal state probability distribution as a function of the motional state $|n\rangle$ and the average motional state $\bar{n}$. First, assume thermal equilibrium of the oscillator which has an average phonon occupation number $\bar{n}$. The probability of occupying the $|n\rangle$ th Fock state is given by the Boltzmann distribution, with Boltzmann constant k_B , temperature T , and variable change $\beta \equiv \hbar\omega/k_BT$,

$$\begin{aligned} p(n) &= \frac{e^{-\beta n}}{\sum_{m=0}^{\infty} e^{-\beta m}} \\ p(n) &= e^{-\beta n}(1 - e^{-\beta}) \end{aligned} \tag{2.11}$$

where the zero-point energy has been dropped without consequence. Weaving in $\bar{n}$, we

incorporate it is definition with Eq. (2.11)

$$\begin{aligned}
\bar{n} &= \sum_{n=0}^{\infty} np(n) \\
\bar{n} &= \sum_{n=0}^{\infty} ne^{-\beta n}(1 - e^{-\beta}) \\
\bar{n} &= 1/(e^{\beta} - 1) .
\end{aligned} \tag{2.12}$$

Rearranging to solve for the Boltzmann factor $e^{\beta} = (\bar{n} + 1) / \bar{n}$ and plugging back into the Boltzmann distribution (Eq. (2.11)), we find the probability for a given $\bar{n}$ to be

$$p(n) = \frac{\bar{n}^n}{(\bar{n} + 1)^{n+1}} \tag{2.13}$$

with the important Boltzmann distribution property that the ratio of any two neighboring probabilities is constant and therefore independent of the motional state n : $p(n + 1)/p(n) = \bar{n}/(\bar{n} + 1)$.

The thermal distribution has assumed the system (trapped ions) are in thermal equilibrium with an environment, but this is not generally true in practice. Although that may seem trivial at first glance, it's repeated for emphasis as it is crucial consideration in the case of trapped-ion systems. Trapped ions are extremely isolated meaning there is weak coupling to the surrounding environment. This leads to slow equilibration. Hardware systems are specifically designed to slow this equilibration. In fact, trapped ions never “truly” thermalize with their environment during typical experimental operations. Assuming ambient temperatures of 300 K, and a $2\pi \times 1$ MHz harmonic frequency, the environment would drive the ion to an $\bar{n}_{\text{eq}} \approx k_B T / (\hbar \omega)$ which computes to roughly 10^6 quanta. Such a high equilibrium energy will most likely lead to problematic experimental consequences. Experiments typically operate at the very lowest motional energy—the ground state of motion. This leads to the best quantum gate fidelities [64]. Even though ions are far from equilibrium with their environment, the motional state probabilities may follow a thermal distribution at any given

moment. They seem to reach a local equilibrium where entropy is maximized for that given moment and consequentially, the motional state population follow Boltzmann statistics to good approximation—with a few important exceptions which will be observed and discussed in Ch. 5 and briefly observed in Ch. 6. In this sense, an ion has thermalized to its own average motional state.

2.3 Singly ionized Ytterbium 171

The choice of what to trap in a Paul trap requires careful consideration. Atomic ions, meaning individual atoms each with a net charge, make an excellent choice for a variety of reasons not least of which is they are astronomically small giving us a window into quantum physics and can be extremely isolated due to the Paul trap and supporting systems.

The selection of elements to trap has been almost exclusively from elements with two valence electrons. This is a strategic choice for two reasons. First after photoionizing one electron, the atom is charged making it trap-able by the Paul trap’s electric fields. Second, since there is just one remaining valence electron, the internal energy level structure will resemble the Hydrogen atom to some degree which is exactly solvable and hence very well understood. Consequentially, as a rule of thumb, this also implies eased laser requirements for experiments. Although ions have been trapped for over 70 years, “new” ions are being explored for the first time such as the slightly radioactive $^{133}\text{Ba}^+$ ion [42, 24] and nonradioactive Barium isotopes [74, 75, 76] or even a proposal for trapping doubly ionized Lanthanum [43, 77]. The recent exploration for new ion species has been pushed, in part, for the need for longer wavelength transitions which better match fiber optic constraints and individual ion focusing technology. Building on the impressive history of “older” ions is left for highly anticipated future work. These are challenging and exciting efforts.

For this work, we trap singly ionized Ytterbium (Yb^+). Ytterbium is a lanthanide with 70 protons and is the heaviest of the atoms typically trapped and used in quantum information science experiments. The ionized 174 isotope $^{174}\text{Yb}^+$ is trapped for troubleshooting trapping,

laser, and imaging systems as its lack of Zeeman levels means it scatters more brightly under resonant laser beams compared to the other isotope commonly trapped— $^{171}\text{Yb}^+$. The ionized 171 isotope $^{171}\text{Yb}^+$ is the workhorse of this work as well as many quantum simulation experiments [7] primarily due to the long-lived hyperfine states used to encode qubit states.

In terms of experimental performance, $^{171}\text{Yb}^+$ makes a powerful qubit. Ytterbium has long lived hyperfine states with demonstrations exceeding an hour of coherence [28]; fast cycling transitions for laser cooling, state preparation, and detection [78, 6]; and effective ion-ion interactions for generating high fidelity entanglement [45] making it an overall high quality qubit—or qudit [79]. There is another important aspect of all trapped ion species which is often glossed over in the literature but is foundational. Nature’s elements have the amazing property that one atom, such as $^{171}\text{Yb}^+$, is “exactly” identical to every other in the entire universe—to as exact as we can tell. They are no fabrication impurities; no difference in transition frequencies (up to external factors); no variations in charge to mass ratio. They are so identical that photons emitted from ions in separate vacuum chambers can be collected, interfered, and analyzed in such a way as to entangle the remote ions [80]!

2.3.1 Energy Level Diagram

The internal energy states of $^{171}\text{Yb}^+$ follow an approximately Hydrogenic (Hydrogen-like) level structure. In the $^2\text{S}_{1/2}$ manifold, the nuclear spin $I = 1/2$ and electronic spin $S = 1/2$ produce hyperfine and Zeeman states. The hyperfine states are robust states for storing quantum information because they are minimally perturbed by magnetic fields, exhibit extremely long quantum coherence times and lifetimes [27, 28], and are highly controllable with modern optical systems [5, 6].

For our apparatus and the experiments outlined in this work, the internal hyperfine states $^2\text{S}_{1/2} |F = 0 \text{ m}_F = 0\rangle$ and $^2\text{S}_{1/2} |F = 1 \text{ m}_F = 0\rangle$ will encode the qubit states $|0\rangle$ and $|1\rangle$ respectively. The states are split by $\sim 12.642 \text{ GHz}$ with a relatively small second order perturbation due to magnetic fields moving the states apart by $310.8B^2 \text{ Hz}$ where B is the

magnetic field in units of gauss at the ion. Magnetic fields on the order of 1 Gauss (G) would affect the transition by $\sim 10^{-7}$; this is one reason why hyperfine states such as these make such robust quantum states. In our lab, the longest maintained superposition state between hyperfine level has been measured to be ~ 0.2 *seconds*. An externally applied magnetic field of ~ 3.9 G sets the quantization axis and breaks the $^2S_{1/2}$ Zeeman state degeneracy without exceeding the $^2P_{1/2}$ linewidth of $2\pi \times 19.6$ MHz.

2.3.2 Atomic Transitions at 369.5 nm


The $^{171}\text{Yb}^+$ ion also lends itself to laser cooling, exponentially fast $|0\rangle$ state preparation, and high-fidelity state detection. The key ingredients for a qubit are realized by irradiating the 369.5 nm $^2S_{1/2} |F=0\rangle \rightarrow ^2P_{1/2} |F=1\rangle$ and $^2S_{1/2} |F=1\rangle \rightarrow ^2P_{1/2} |F=0\rangle$ transitions.

Detection

The qubit states can be detected through, what is commonly referred to as, state-dependent fluorescence readout [78, 81]. When irradiating $^2S_{1/2} |F=1\rangle$ to $^2P_{1/2} |F=0\rangle$, the ion will only scatter photons if it is in the $|1\rangle$ state. The $|0\rangle$ state is mostly unaffected since the nearest transition is detuned by the hyperfine splitting of ~ 12.6 GHz. An external magnetic field of ~ 3.9 G (up to 5 G is typical) destabilizes coherent state trapping among the Zeeman states within the linewidth of the transition (see Fig. 2.1) [82, 81]. Otherwise, fluorescence would be significantly more suppressed. The photons scattered by the ion are imaged onto a photomultiplier tube (pmt) or electron multiplying charged-coupled device (EMCCD) camera (hardware details are in Sec. 3.4). The collected photons are stochastically emitted by the ion in the $|1\rangle$, or bright state, and hence follow Poisson statistics

$$P(N) = \lambda^N e^{-\lambda} / N! \quad (2.14)$$

for N detected photons with an average number of photons λ as shown in Fig. 2.3.






Figure 2.2: Diagram of the detection and optical pumping 369.5 nm transitions. An EOM driven at 2.105 GHz shifts the optical pumping light frequency. Repumping out of the $^2D_{3/2}$ state is not shown. In the optical pumping diagram, only the spontaneous decay channels that accumulate population in the dark state $|0\rangle$ are shown.

The bright ($|1\rangle$) and dark ($|0\rangle$) states need to be distinguished from each other. The bright ion data will produce a Poisson distribution with average λ_{bright} (blue outlined distribution of Fig. 2.3). Although the bright state is the only state scattering photons for the PMT to detect, there is a non-zero chance that no photons are scattered during the detection period even though the ion was in the bright state $P(0) = e^{-\lambda_{\text{bright}}}$. Additionally, experiments are done in the real world so stray photons will hit the PMT resulting in another Poisson distribution of photons when measuring a dark ion λ_{dark} (red outlined distribution of Fig. 2.3). The standard approach to state discrimination is to measure the Poisson distributions for bright and dark ions and then set a threshold point between the two Poisson distributions as shown in Fig. 2.3. All detection periods that count more photons than the threshold are classified as a bright ion; all detection periods that count fewer photons than the threshold are classified as a dark ion. The detection infidelity is then the sum of the probability of the




Figure 2.3: Normalized histogram of PMT photon counts from a 1 ms detection period repeated 1000 times. When the ion is prepared in the bright state (blue outline), the mean photon counts are 11.3. For a dark ion (red outline), the mean background photon counts are 0.59. The solid bars indicate the respective overlap between the dark (solid red) and bright (solid blue) distributions, and the threshold (dashed black vertical line) is placed to minimize the overlap. The bright state detection fidelity is 96.5 %. The dark state fidelity is 99.0 %. The average fidelity is then 97.7 %.

ion being dark but collecting more photons than the threshold value and the reverse, the probability of the ion being bright but collecting fewer photons than the threshold value. We typically can detect dark states with high fidelity $\sim 99\%$, but detect bright ions with $\sim 96.5\%$ fidelity. A bright ion detection errors are commonly higher because a bright state is more likely to scatter into the dark state part way through measurement due to off-resonant coupling through the $^2P_{1/2}|F = 1\rangle$ (~ 2.1 GHz away) resulting in a tail on the bright ion Poisson statistics which overlaps with the dark state distribution [83, 84].

Optical Pumping

The ion qubit state can be prepared in $|0\rangle$, also known as the dark state, exponentially fast using optical pumping [78]. This is usually done at the start of experiment to help to ensure




Figure 2.4: Optical pumping from the bright qubit state to the dark qubit state. The y -axis is the probability of being in the bright qubit state. The x -axis is the length of time the optical pumping beam is on. In this example measurement, the characteristic decay time is $\tau_{\text{op}} = 1.12 \pm 0.08 \mu\text{s}$.

that repeated experiments are as identical as possible setting a reproducible starting point to then drive quantum gate operations.

State preparation is achieved by driving the triplet states in the $^2S_{1/2} |F = 1\rangle$ manifold to $^2P_{1/2} |F = 1\rangle$ or $^2P_{1/2} |F = 0\rangle$. If the ion is in the $|1\rangle$ state or any of the neighboring Zeeman levels, the laser light will excite the ion into one of the $^2P_{1/2}$ states. This is achieved experimentally with resonance 369.5 nm light and a 2.105 GHz EOM (see Sec. 3.3.4 for experimental details). There is a probability of spontaneously decaying from the $^2P_{1/2} |F = 1\rangle$ manifold into the $|0\rangle$ dark state and since there are no laser transitions out of that state, once the ion is in that state it stays there. If the ion decayed to the $^2S_{1/2} |F = 1\rangle$ manifold, incident light re-excites the ion until it is completely pumped into $|0\rangle$. This leads to the exponential likelihood of being pumped into the dark state. Typical time scales are $\sim 1 \mu\text{s}$ as shown in Fig. 2.4.

Optical pumping is fast and efficient. For the characteristic time of $\tau_{\text{op}} = 1.12 \mu\text{s}$ in

Fig. 2.4, a 9 μs optical pumping pulse would ideally have a state preparation fidelity of $1 - e^{(-9/1.12)} = 99.97\%$. We limit the optical pumping rate to about $\tau_{\text{op}} \simeq 1\ \mu\text{s}$ to avoid spontaneous emission into adjacent motional states [6]. This becomes important during ground state cooling procedures (see Sec. 2.4.3). Briefly, if optical pumping is too fast its effective linewidth will overlap with neighboring motional states and reduce the final achievable average motional state.

Doppler Cooling

Laser cooling of atoms was one of the landmark atomic physics discoveries of the 1970s and 80s [85]. The velocity of an atom will broaden its transitions reducing spectroscopic accuracy, and, in the context of modern quantum hardware atoms need to be cold to trap in the first place. Additionally, hot atoms will, in general, limit the experimental control needed for high-fidelity quantum gate operations. Today, laser cooling, also known as Doppler cooling, is the idle state for trapped ions continuously maintaining the ions at millikelvin temperatures.




Figure 2.5: Diagram of Doppler cooling split into its effective absorption and emission processes. The ion is a black circle in a harmonic well (grey u-shaped line). For absorption, the incident photon is red to indicate its frequency is below the atomic transition frequency. In the top row, the ion velocity is opposite the momentum of the incident photon. The Doppler effect up-shifts the photon frequency such that the atom absorbs the photon, and its momentum is reduced. (The blue ion indicates it has absorbed the photon.) In the lower panel, the Doppler shift reduces the likelihood of absorbing the photon as shown by the photon passing by and velocity unaffected. Emission is spontaneous decay and therefore in random directions and at random times as shown by the randomly distributed emitted photons.

Doppler cooling uses the geometric dependence of the Doppler frequency shift to selec-

tively control an atom’s absorption of photons (Fig. 2.5). By lowering the frequency of the incident photons below the transition frequency, but within the transition linewidth, atoms moving toward the incident photons will experience a Doppler shift which effectively moves the frequency of the light back toward resonance—in the ion’s frame of reference. Therefore, the likelihood of absorbing photons when the ion is moving anti-parallel with the photon wavevector increases, and the probability of absorbing photons when the ion is moving parallel with the photon wavevector decreases. The cooling comes in when we consider that the photon has momentum $\hbar k$. Since it is more likely that the atom absorbs photons when moving anti-parallel to the photon’s propagation, it is more likely that an absorption event will *reduce* the momentum of the atom. Indeed, it does so quite rapidly.

In $^{171}\text{Yb}^+$, Doppler cooling is achieved by irradiating all the states in the $^2\text{S}_{1/2}$ manifold with light that is nearly-resonant with the $^2\text{P}_{1/2}$ manifold. To do this, the scheme is similar to detection (Fig. 2.2), but an additional light frequency up-shifted by 14.7 GHz EOM (see Sec. 3.3.4 for experimental details) to drive the $^2\text{S}_{1/2} |F=0\rangle \rightarrow ^2\text{P}_{1/2} |F=1\rangle$ transition. The transition is fast with an average lifetime of 8 nanoseconds. Experimentally, we have estimated that millions of photons are scattered every second quickly cooling trapped ions to an average motional state around $\bar{n} \sim 10 - 20$ depending on the harmonic frequency. The stochastic nature of emission limits the lowest achievable motional energy to approximately $\bar{n}_{\text{dop}} = \gamma/2\omega$ where $\gamma = 2\pi \times 19.6$ MHz for $^{171}\text{Yb}^+$ is the natural linewidth and ω the harmonic motional frequency.

We now dive more deeply into Doppler cooling. We work out the semi-classical calculation of the motional energy of a trapped ion experiencing Doppler cooling following Refs. [8, 53] with some additional considerations and limits. We will also demonstrate how these equations can capture the physics of motional heating due to photon scatter. This derivation will focus on parameters specific to our current trap and $^{171}\text{Yb}^+$. We begin by assuming a resonant beam enters the trap with intensity I and wavevector $\vec{k}$ near 369.5 nm. When this light hits a trapped ion moving with velocity $\vec{v}$ and a photon is absorbed, a photon is later

emitted with wavevector $\vec{k}_s$.

The change in energy due to the absorption and emission of a single photon is

$$\begin{aligned}\Delta E &= \frac{(m\vec{v} + \hbar\vec{k} - \hbar\vec{k}_s)^2}{2m} - \frac{(m\vec{v})^2}{2m} \\ &= \frac{\hbar^2(\vec{k} - \vec{k}_s)^2}{2m} + \hbar(\vec{k} - \vec{k}_s) \cdot \vec{v}.\end{aligned}\tag{2.15}$$

To find the energy rate of change along a direction of interest (say, the x direction), we need to multiply ΔE by the scattering rate, and average over the absorption and emission scattering directions:

$$\frac{dE_x}{dt} = \left\langle \Gamma(\omega, \vec{v}) \left[\frac{\hbar^2 k^2}{2m} (f_x + f_{sx}) + \hbar k_x v_x \right] \right\rangle \tag{2.16}$$

with geometry factors f_x and f_{sx} and x -axis projection of the incidence photon wavevector k_x . In our experiment, $k_x = k/2$. The geometry factor $f_x = \hat{k}_x^2$ accounts for the fraction of energy transferred to the x -direction during a photon absorption, which is $f_x = (1/2)^2 = 1/4$ for our experiment. The geometric factor f_{sx} is likewise the fraction of energy transferred to the x -direction during spontaneous emission. Isotropic emission would give $f_{sx} = 1/3$. However, more detailed calculations argue for slight deviations from this value to account for the dipole emission pattern [86]. For $^{171}\text{Yb}^+$, when integrating over the scattering emission patterns for $\Delta m_F = 0$ and $\Delta m_F = \pm 1$ transitions [53] and weighting each one by 1/3, the isotropic result is recovered $f_{sx} = 1/3$.

The scattering rate $\Gamma(\omega, \vec{v})$ for $^{171}\text{Yb}^+$ as defined in Ref. [81], at the optimal polarization angle $\theta = \arccos(1/\sqrt{3})$, is given by

$$\begin{aligned}\Gamma(\omega, \vec{v}) &= \frac{\gamma(s/18)}{1 + \frac{1}{216}(\frac{s\gamma}{\delta_B})^2 + \frac{8}{3}(\frac{\delta_B}{\gamma})^2 + 4(\frac{-\Delta + \vec{k} \cdot \vec{v}}{\gamma})^2} \\ \Gamma(\omega, \vec{v}) &= \frac{\gamma(s/18)}{1 + s' + 4(\frac{-\Delta + \vec{k} \cdot \vec{v}}{\gamma})^2}\end{aligned}\tag{2.17}$$

where $s' = \frac{1}{216}(\frac{s\gamma}{\delta_B})^2 + \frac{8}{3}(\frac{\delta_B}{\gamma})^2$, $\gamma = 2\pi \times 19.6$ MHz is the natural linewidth, δ_B is measured the Zeeman splitting, $s = I/I_{sat}$ is the saturation parameter, and Δ is the detuning from resonance ($\Delta > 0$ is above the transition—blue detuned). In Eq. (2.17), the Doppler shift $\vec{k} \cdot \vec{v}$ is small when starting at low energy, allowing the approximation:

$$\Gamma(\omega, \vec{v}) \approx \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}} \left[1 + \frac{8\Delta(\vec{k} \cdot \vec{v})}{\gamma^2(1 + s') + 4\Delta^2} \right]. \quad (2.18)$$

Putting this back into Eq. (2.16),

$$\frac{dE_x}{dt} = \left\langle \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}} \left[1 + \frac{8\Delta(k_x v_x + k_y v_y + k_z v_z)}{\gamma^2(1 + s') + 4\Delta^2} \right] \left[\frac{\hbar^2 k^2}{2m}(f_x + f_{sx}) + \frac{\hbar k}{2} v_x \right] \right\rangle. \quad (2.19)$$

A trapped ion has an zero average velocity $\langle v_i \rangle = 0$, so any term linearly-dependent on velocity terms in Eq. (2.19) can be discarded. Additionally, $\langle v_i v_j \rangle = 0$ for any two directions $i \neq j$ if the harmonic frequencies are unique which is typically the case experimentally. Therefore, the cross-terms in Eq. (2.19) can also be set to zero simplifying the expression to

$$\frac{dE_x}{dt} = \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}} \left[\frac{\hbar^2 k^2}{2m}(f_x + f_{sx}) + \frac{8\Delta\hbar f_x k^2 \langle v_x^2 \rangle}{\gamma^2(1 + s') + 4\Delta^2} \right]. \quad (2.20)$$

(Note the substitution of $k_x \rightarrow \sqrt{f_x}k$ in the 2nd term of Eq. (2.20)). Finally, we use the property that for a classical harmonic oscillator, the total energy $E_x = m\langle v_x^2 \rangle$ (with no factor of 1/2):

$$\frac{dE_x}{dt} = \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}} \left[\frac{\hbar^2 k^2}{2m}(f_x + f_{sx}) + \frac{8\Delta\hbar f_x k^2/m}{\gamma^2(1 + s') + 4\Delta^2} E_x \right]. \quad (2.21)$$

Eq. (2.21) is a first-order differential equation which can be solved analytically. For

convenience, we define:

$$\Gamma_0 \equiv \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}} \quad ; \quad R \equiv \frac{\hbar^2 k^2}{2m}(f_x + f_{sx}) \quad ; \quad D \equiv \frac{8\Delta\hbar f_x k^2/m}{\gamma^2(1 + s') + 4\Delta^2} . \quad (2.22)$$

then the differential equation (Eq. (2.21)) is recast as

$$\frac{dE_x}{dt} = \Gamma_0(R + DE_x) \quad (2.23)$$

which has the solution

$$E_x(t) = \left(E_x(0) + \frac{R}{D} \right) e^{\Gamma_0 D t} - \frac{R}{D} . \quad (2.24)$$

We now examine Eq. (2.24) in a few limits. First, in the long-time limit at negative detuning ($D < 0$), the final value of the energy converges to:

$$\begin{aligned} E_x(\infty) = -\frac{R}{D} &= -\frac{\hbar^2 k^2}{2m}(f_x + f_{sx}) \frac{\gamma^2(1 + s) + 4\Delta^2}{8\Delta\hbar f_x k^2/m} \\ &= -\frac{\hbar\gamma}{8} \left(1 + \frac{f_{sx}}{f_x} \right) \left[\frac{\gamma(1 + s)}{2\Delta} + \frac{2\Delta}{\gamma} \right] . \end{aligned} \quad (2.25)$$

This agrees with results in the literature [53, 6], which are typically derived by setting $dE_x/dt = 0$ in Eq. (2.21) and define detuning Δ with opposite sign.

Second is the special case of zero detuning ($\Delta = 0$ and therefore $D = 0$). This would not laser cool because it is not red-detuned, but this limit is important because it will define how resonant light—such as during detection or optical pumping—might affect the motional energy of a trapped ion. In the limit as $D \rightarrow 0$,

$$E_x(t) = E_x(0) + \Gamma_0 R t \quad (2.26)$$

With no damping force from cooling ($\Delta = 0$), the ion is heated linearly at the rate $\dot{n} = \Gamma_0 R$.

Third, if the beat is blue-detuned ($\Delta < 0$), the ion will heat exponentially. All of these

cases will be detailed and measured in Ch. 6.

As has been discussed, the 369 nm transitions $^2S_{1/2} |F = 0\rangle \rightarrow ^2P_{1/2} |F = 1\rangle$ and $^2S_{1/2} |F = 1\rangle \rightarrow ^2P_{1/2} |F = 0\rangle$ cover a substantial part of the needed operations for quantum science with trapped ions. These kinds of operations are just the starting point however. To do quantum information science such as quantum computing and simulation, coherent drives that manipulate the quantum state of the ion are needed and will be discussed in the next section.

2.4 Coherent Atom-laser Interactions

The theoretical background covered so far coincides with the historical develop of atomic, molecular, and optical physics. We will go backward in time a bit to follow a different scientific thread ending in the era of quantum information science in the 21st century.

Coherent atom-laser interactions make up the primary method for manipulating the quantum degrees of freedom of a trapped-ion qubits. Because lasers can be so finely tuned, quantum gate operations from atom-laser interactions can be precisely driven with fidelities steadily increasing [18, 19, 20, 21, 22, 23]. This section briefly outlines the essential Hamiltonians and state dynamics pertinent to ion trap experiments, which lays out the fine control available to experiment with.

Before diving in, I briefly define some notation. The Pauli matrices:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad , \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad , \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (2.27)$$

which in most quantum information science literature are denoted with capital letters X , Y , and Z respectively. I will stick to the σ_i notation for continuity with previous trapped ion work. For spin raising and lowering operators:

$$\sigma^+ \equiv (\sigma_x + i\sigma_y)/2 \quad , \quad \sigma^- \equiv (\sigma_x - i\sigma_y)/2 \quad . \quad (2.28)$$

This definition is often switched in quantum information contexts [87]. The reason has to do with the difference between spin-1/2 representation and qubit representation. The ground state in spin-1/2 representation is

$$|\downarrow\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (2.29)$$

to raise the energy to the excited state $|\uparrow\rangle$ the operator $(\sigma_x + i\sigma_y)/2$ is used. In direct opposition, the “ground” state in quantum information is zero

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (2.30)$$

which has the matrix form of the excited state of the spin-1/2 representation $|0\rangle = |\uparrow\rangle$. Due to this reversal, the raising and lowering operator definitions are also reversed between the two representations. In this work, I will stick with the historical spin-1/2 definitions for continuity with previous work with trapped ions, but I am of the opinion to migrate toward the colloquial representations emerging in the quantum information science community.

2.4.1 Rabi Flopping

The Rabi “flop” is arguably the first and most basic level of control of a quantum state [88]. Named after Isidor Rabi (1898 - 1988), Rabi flopping describes the oscillatory evolution of a quantum state between two states due to an external oscillating field and is shown in Fig. 2.8. The evolution is coherent meaning the intermediate states are quantum superposition states. In the emerging quantum information science field, these drives are called single-qubit gate operations. Single-qubit gates can be driven on groups of trapped ions simultaneously by irradiating them all with the same external drive (such as a laser beam). Individual ions can be driven using individual addressing techniques which, in essence, focus a laser beam onto the ions of interest [89, 32, 90, 91, 92]. Either way, the Hamiltonian relatively is the same for single-qubit drives.

The Hamiltonian of an effective two-level system separated by frequency ω_0 and inter-

acting with an external drive at frequency ω_{drive} (such as an optical or microwave photon), with a detuning $\Delta = \omega_0 - \omega$, may be written as

$$H_{\text{carrier}} = \frac{\hbar\Omega}{2} \left(\sigma^+ e^{-i(\Delta t + \phi)} + \sigma^- e^{i(\Delta t + \phi)} \right) \quad (2.31)$$

where Ω is the Rabi frequency, ϕ is the phase of the external drive. This description is useful when considering $^{171}\text{Yb}^+$ as a qubit encoded in the internal hyperfine levels and directly driving the transition with 12.6 GHz microwaves (see Sec. 3.3.7 for hardware details).

We can solve the unitary evolution driven by this Hamiltonian (Eq. (2.31)) using the Schrödinger equation. First, consider the on-resonance case $\Delta = 0$ making the Hamiltonian time-independent. Using the definitions $\sigma^+ \equiv (\sigma_x + i\sigma_y)/2$ and $\sigma^- \equiv (\sigma_x - i\sigma_y)/2$, the Hamiltonian simplifies

$$\begin{aligned} H_{\text{carrier}} &= \frac{\hbar\Omega}{2} \left(\sigma^+ e^{-i\phi} + \sigma^- e^{i\phi} \right) \\ &= \frac{\hbar\Omega}{2} \left[\sigma_x (e^{i\phi} + e^{-i\phi})/2 - i\sigma_y (e^{i\phi} - e^{-i\phi})/2 \right] \\ &= \frac{\hbar\Omega}{2} \left[\sigma_x \cos \phi + \sigma_y \sin \phi \right] . \end{aligned} \quad (2.32)$$

Some texts may further define $\sigma_\phi = \sigma_x \cos \phi + \sigma_y \sin \phi$.

The unitary time evolution operator of the Hamiltonian is then

$$\begin{aligned} U(t) &= e^{-iH_{\text{carrier}}t/\hbar} \\ U(t) &= e^{-i(\Omega/2)(\sigma_x \cos \phi + \sigma_y \sin \phi)t} \\ U(t) &= \mathbb{1} \cos(\Omega t/2) - i \sin(\Omega t/2) (\sigma_x \cos \phi + \sigma_y \sin \phi) . \end{aligned} \quad (2.33)$$

Consider the evolution of the z -basis state $|0\rangle$. An x -axis rotation can be achieved by setting $\phi = 0$ making the evolution $U(t)|0\rangle = \cos(\Omega t/2)|0\rangle - i \sin(\Omega t/2)|1\rangle$. Rotation around the y -axis is achieved by setting $\phi = \pi/2$ with the evolution $U(t)|0\rangle = \cos(\Omega t/2)|0\rangle - \sin(\Omega t/2)|1\rangle$.

The probability of measuring the $|1\rangle$ is then the commonly defined Rabi flopping equation

$$P_{|1\rangle}(t) = \sin^2(\Omega t/2) . \quad (2.34)$$

In quantum information science, these rotations are cast as quantum logic gates with the following time evolution definitions

$$\begin{aligned} R(\theta, \phi) &= \exp(-i\sigma_\phi\theta/2) = \mathbb{1} \cos(\theta/2) - i \sin(\theta/2)\sigma_\phi \\ &= \begin{pmatrix} \cos(\theta/2) & -i \exp(-i\phi) \sin(\theta/2) \\ -i \exp(i\phi) \sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \end{aligned} \quad (2.35)$$

$$\begin{aligned} R_z(\theta) &= \exp(-iZ\theta/2) = \mathbb{1} \cos(\theta/2) - i \sin(\theta/2)\sigma_z \\ &= \begin{pmatrix} \exp(-i\theta/2) & 0 \\ 0 & \exp(i\theta/2) \end{pmatrix} . \end{aligned} \quad (2.36)$$

For rotation around the x -axis, set $\phi = 0$: $R_x(\theta) = R(\theta, 0)$. For rotation around the y -axis, set $\phi = \pi/2$: $R_y(\theta) = R(\theta, \pi/2)$.

Going back to equation Eq. (2.31), if the detuning was nonzero, but still small relative to the transition frequency, Rabi flopping changes. The detuning becomes part of the generalized Rabi frequency $\tilde{\Omega} \equiv \sqrt{\Omega^2 + \Delta^2}$. Starting in the $|0\rangle$ initial state as before, the maximum likelihood of reaching $|1\rangle$ is now limited by the detuning as well as adjusted oscillation frequency. To be as general as possible, and physically relevant, an exponential decaying envelope due to decoherence or dephasing is also included with decay rate γ :

$$\begin{aligned} P_{|1\rangle}(t) &= \frac{\Omega^2}{\tilde{\Omega}^2} \sin^2(\tilde{\Omega}t/2) \\ P_{|1\rangle}(t) &= \frac{\Omega^2}{2\tilde{\Omega}^2} \left(1 - e^{-\gamma t} \cos \tilde{\Omega}t \right) . \end{aligned} \quad (2.37)$$

Experimentally, the Rabi frequency Ω , detuning Δ , time t , and rotation axis angle ϕ are well controlled. With these knobs, most any single qubit state can be realized up to

the precision of the hardware. In quantum computing, the continuous drives are defined in terms of fixed pulse times: π -pulse, $\pi/2$ -pulse, etc. The action is straight forward. For a π -pulse, assuming $\phi = 0$, the drive time is half a period $t = \pi/\Omega$ making the probability of flipping states $P_{|1\rangle}(t) = \sin^2[\Omega\pi/(2\Omega)] = 1$. This is equivalent to a bit flip in classical gate operations. A $\pi/2$ -pulse is more quantum. Starting from a basis state such as $|0\rangle$ and $\phi = 0$, a $\pi/2$ -pulse generates an even superposition state $U[\pi/(2\Omega)]|0\rangle = (|0\rangle + i|1\rangle)/\sqrt{2}$. If $\phi = \pi/2$, for a y -axis rotation, the final state would just have a different phase. There are a variety of single-qubit gate operation definitions in quantum information science (e.g. Y , $S^\dagger$, T), but they all ultimately boil down to a specified rotation axis ϕ and a fractional period drive time t .

There may be timing errors, a drifting Ω values, or a noisy phase ϕ , but these are typically higher order affects that will impact fidelities at the few percent level. They are readily suppressed by upgrading hardware to driving more “cleanly” or any one of the many quantum control techniques or composite pulse sequences [93].

2.4.2 Stimulated Raman Transitions

Stimulated Raman transitions are another experimental technique used to drive single-qubit rotations and two-body entangling interactions. By turning a 3-level atomic system into an effective two-level system, Raman transitions can have significant spin-motion coupling which is necessary for two-body interactions.

Raman transitions are two-photon transitions where one photon far-off-resonantly excites an internal state $|g_1\rangle$ of an ion to a virtual state—significantly detuned from an excited state—and the second photon stimulates emission out of the virtual state back into another internal state $|g_2\rangle$. Under the right conditions, namely the correct ratio of Rabi frequencies and detunings, the intermediate state is adiabatically eliminated thus minimizing spontaneous emission, but an ion still experiences recoil energies proportional to the incident photon wavevector (up to geometric factors). For example, we can excite the 12.6 GHz qubit transi-

tion with a two-photon stimulated Raman transitions at 355 nm, and a 355 nm photon has about 10,000,000 times the wavevector magnitude of a 12.6 GHz photon.




Figure 2.6: Schematic diagram of a Λ -configuration stimulated Raman transition. There are three internal states: $|g_1\rangle$, $|g_2\rangle$, and $|e\rangle$. The top dashed line can be thought of as a virtual internal state while the bottom dashed line is there to show the detuning δ from $|g_2\rangle$. The blue and orange lines with double arrows represent two incident beams at different frequencies and potentially different strengths.

Many trapped ions species are driven by a Λ -configuration stimulated Raman transition as shown in Fig. 2.6. Two laser beams are far detuned (Δ is large) from an excited state $|e\rangle$, but the frequency difference is nearly resonant between $|g_1\rangle$ and $|g_2\rangle$ ($\delta \ll \Delta$). In solving for the effective Hamiltonian H_{raman} , the excited state can be adiabatically eliminated leaving only coupling between the two ground states $|g_1\rangle$ and $|g_2\rangle$:

$$H_{\text{raman}} = \frac{\hbar\Omega}{2} \left[\sigma^+ e^{i(\Delta\vec{k}\cdot\vec{r}-\delta t+\phi)} + \sigma^- e^{-i(\Delta\vec{k}\cdot\vec{r}-\delta t+\phi)} \right] \quad (2.38)$$

where $\Delta\vec{k}$ is the wave vector difference of the two incident photons, $\vec{r}$ is the position vector of the ion, and ϕ is a constant phase difference of the two beams.

In our lab, we drive a Raman transition in $^{171}\text{Yb}^+$ using 355 nm photons, and experimental details and motivations for this wavelength can be found in Sec. 3.3.5. There are several

theoretical considerations that make this wavelength a desirable choice [88]. The 355 nm light is very far detuned from internal states. Assuming the ion is in an $^2S_{1/2}$ state, 355 nm light is ~ 33 THz detuned from $^2P_{1/2}$ and ~ 66 THz detuned from $^2P_{3/2}$ (see Fig. 2.11). These are very large detunings which is advantageous because, in a Raman configuration, spontaneous decay is proportional to $1/\Delta^2$ where Δ is the detuning [94, 95]. In fact, spontaneous decay is estimated to contribute a negligible error of $< 10^{-5}$ per π -pulse [94]. Additionally, 355 nm is a nearly “magic” wavelength where the coupling to the two P states adds constructively to the Rabi frequency Ω while the differential AC stark shift $\sim 10^{-4}\Omega$ nearly cancels [94]. Both of these features are desirable. The 355 nm light also has a relative high momentum transfer with a Lamb-Dicke parameter $\eta \sim 0.1$ which can be a sweet spot for the motional coupling needed for trapped-ion two-qubit quantum gate operations.

2.4.3 Spin-motion Coupling

The motional coupling comes from the $\Delta\vec{k} \cdot \vec{r}$ phase. Assuming the Raman transition linewidth is narrow enough that the motional sidebands are resolved ($\Omega \ll \omega_z$), the motional phase term may be rewritten in terms of quantum harmonic ladder operators a and $a^\dagger$. Consider non-copropagating beams with a net wavevector $\Delta\vec{k}$ along the z -axis of motion. The motional phase, as shown in Ref. [5], is then

$$\exp\left(i\Delta\vec{k} \cdot \vec{r}\right) = \exp\left[i\eta(ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t})\right] \quad (2.39)$$

where the Lamb-Dicke parameter $\eta = \Delta k_z z_0$, $z_0 = (\hbar/2m\omega_z)^{1/2}$ is the center-of-mass (COM) mode zero-point energy wavefunction spread, and ω_z is the harmonic frequency along the z -axis. The Lamb-Dicke parameter compares the wavefunction spread of the ion to the gradient of the laser field. It is typically small $\eta \ll 1$, but not too small that the spin-motion Rabi frequencies derived later on are small. For experiments in this work, the Lamb-Dicke parameter is around $\eta \sim 0.12 \pm 0.2$.

The motional coupling to the internal qubit states becomes clearer by expanding the motional phase Eq. (2.39) in powers of η and plugging this back into Eq. (2.38). The expanded Hamiltonian is then

$$H_{\text{sm}} = \frac{\hbar\Omega}{2}\sigma^+ \left[\mathbb{1} + i\eta \left(ae^{-i\omega_z t} + a^\dagger e^{i\omega_z t} \right) e^{i(-\delta t + \phi)} \right] + h.c. . \quad (2.40)$$

Depending on the detuning δ , different terms in the Hamiltonian will dominate.

There are three unique transitions which are most often driven in trapped-ion experiments: a “red sideband” transition ($\delta = -\omega_z$); a “carrier” transition ($\delta = 0$); and a “blue sideband” transition ($\delta = \omega_z$). At the red detuning $\delta = -\omega_z$, the spin-motion Hamiltonian reduces to a Jaynes-Cummings Hamiltonian

$$H_{\text{red}} = \frac{\hbar\eta\Omega}{2} \left(\sigma^+ a e^{i\phi + \pi/2} + \sigma^- a^\dagger e^{-i\phi + \pi/2} \right) . \quad (2.41)$$

At the carrier detuning $\delta = 0$, it reduces to Eq. (2.32)

$$H_{\text{car}} = \frac{\hbar\Omega}{2} \left(\sigma^+ e^{i\phi + \pi/2} + \sigma^- e^{-i\phi + \pi/2} \right) . \quad (2.42)$$

At the blue detuning $\delta = \omega_z$, it reduces to an anti-Jaynes-Cummings Hamiltonian

$$H_{\text{blue}} = \frac{\hbar\eta\Omega}{2} \left(\sigma^+ a^\dagger e^{i\phi + \pi/2} + \sigma^- a e^{-i\phi + \pi/2} \right) . \quad (2.43)$$

In the three equations above, a slight detuning may exist, but this phase term has been omitted for brevity. The red and blue sideband Hamiltonians couple spin and motion as a spin flip will also raise or lower the motional state. The carrier is motional state independent.

The detuning δ controls which Hamiltonian is driving an ion, and δ is controlled by acousto-optic modulators at with high accuracy (see Sec. 3.5.5). Figure 2.7 shows a typical sweep of δ to measure the transitions frequencies. For example, for an ion to experience the




Figure 2.7: Frequency scan centered about the carrier transition. The red and blue COM modes are labeled and colored accordingly.

red sideband Hamiltonian H_{red} , we would experimentally set the laser frequency to the “-x COM” peak in Fig. 2.7. For H_{blue} , set the laser frequency at the “+x COM” peak. Same for the carrier transition. This laser frequency dependence gives experimenters control over what Hamiltonian trapped ions experience. This will also be the case for two-qubit gate operations shown in Sec. 2.4.4.

At times, higher order corrections to the red, carrier, and blue sideband Hamiltonians are useful, particularly when computing the Rabi frequency of transitions between different high energy motional states. The three equations are a good approximation for low motional energy states and low η which satisfy the Lamb-Dicke regime condition $\eta^2(2n+1) \ll 1$. The

exact Rabi frequencies can be computed following Ref. [5]:

$$\Omega_{n',n} \equiv \Omega \langle n' | e^{i\eta(a+a^\dagger)} | n \rangle \quad (2.44)$$

$$= \Omega e^{-\eta^2/2} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta^{|n-n'|} \mathcal{L}_{n_{<}}^{|n-n'|}(\eta^2), \quad (2.45)$$

where $n_{<}$ ($n_{>}$) is the lesser (greater) of n and n' and

$$\mathcal{L}_n^{(\alpha)}(X) = \sum_{i=0}^n (-1)^i \binom{n+\alpha}{n-i} \frac{X^i}{i!} \quad (2.46)$$

is the generalized Laguerre polynomial. This definition of the Rabi frequency will be needed for the SVD thermometry method outlined in Sec. 2.4.3.

Carrier Experiments

The previous calculations established the basis for single-qubit manipulations. This subsection highlights their experimental realization and techniques which rely on them such as rotation gate operations, changing measurement bases, and Ramsey interferometry.

Carrier Drive Starting with the most straight forward drive, a typical carrier drive (Eq. (2.42)) is shown in Fig. 2.8. The dynamics following a two-level system where the Rabi oscillations are between the two hyperfine qubit states. The measured Rabi frequency is $\Omega = 149.0 \pm 0.14$ kHz with no measurable decay within the time window. Higher and lower Rabi frequencies are readily achieved adjusting the laser power (see Sec. 3.3.5). Error bars on data points are smallest at near a probability of 0 and 1 and largest near a probability of 0.5 due to shot noise. The maximum oscillation contrast is approximately 97.5 % without correcting for detection errors. Detection errors—about 3.5 % for bright state detection Sec. 2.3.2—conflates the estimate of the π -pulse error, but it can be removed by either standard measurement error mitigation post analysis [96] or fitting a decay to the oscillation. The decay rate τ is then inversely proportional to the fidelity $F = t_\pi/\tau$. The estimated π -pulse




Figure 2.8: Carrier oscillation from a stimulated Raman drive. Rabi frequency is $\Omega = 149.0 \pm 0.14$ kHz and approximately 99.2 ± 0.3 % π -pulse fidelity. The maximum oscillation contrast is 97.5 % mostly due to bright state detection errors.

fidelity for this example data is 99.2 ± 0.3 %.

The measurement basis of a carrier drive (or red and blue sideband drives) can be precisely controlled. This is particularly useful for witnessing GHZ state entanglement as will be discussed in Sec. 2.4.4. Consider a single qubit oscillating around the x basis axis due to a carrier drive as shown in Fig.(c) 2.9. The ion is measured along the three axes as shown in Fig. 2.9: $\langle \sigma_x \rangle$, $\langle \sigma_y \rangle$, and $\langle \sigma_z \rangle$. Experimentally, this is realized in three steps. First, the qubit state is initialized in the z -basis $|0\rangle_z$. Second, a carrier oscillation ($\phi = 0$) drives the time evolution $|\phi(t)\rangle_z = U(t) |0\rangle_z$ where $U(t)$ is from Eq. (2.33). Lastly, a $\pi/2$ -pulse rotation swaps the measurement axis. Since trapped ion experiments detect the qubit state along only one axis, that final $\pi/2$ -pulse rotates the desired axis for detection.

For a z -basis measurement, no $\pi/2$ -pulse is needed. For a y -basis measurement $\langle \sigma_y \rangle$, a $\pi/2$ -pulse with $\phi = 0$ rotates the y components of the state vector $|\psi(t)\rangle$ to be along the z




Figure 2.9: Rabi flopping round the x -axis ($\phi = 0$) measured along different bases: (a) x -basis, (b) y -basis, (c) z -basis.

axis. The experimental sequence written as an expectation value is as follows

$$\langle 0|_z U(t)^\dagger R^\dagger(\pi/2, 0) \sigma_z R(\pi/2, 0) U(t) |0\rangle_z . \quad (2.47)$$

The inner rotation operators transform the σ_z in to a σ_y as shown below

$$\begin{aligned}
R^\dagger(\pi/2, 0)\sigma_z R(\pi/2, 0) &= \frac{1}{2}(\mathbb{1} + i\sigma_x)\sigma_z(\mathbb{1} - i\sigma_x) \\
&= \frac{1}{2}(\sigma_z + i[\sigma_x, \sigma_z] + \sigma_x\sigma_z\sigma_x) \\
&= \sigma_y
\end{aligned} \tag{2.48}$$

where $i[\sigma_x, \sigma_z] = 2\sigma_y$ and $\sigma_x\sigma_z\sigma_x = -\sigma_z$. Plugging this back into Eq. (2.47), we recover the expectation value of σ_y

$$\begin{aligned}
\langle 0|_z U(t)^\dagger R^\dagger(\pi/2, 0)\sigma_z R(\pi/2, 0)U(t)|0\rangle_z &= \langle 0|_z U(t)^\dagger \sigma_y U(t)|0\rangle_z \\
&= \langle \sigma_y \rangle .
\end{aligned} \tag{2.49}$$

Similarly, for an x -basis measurement, a $\pi/2$ -pulse with $\phi = 3\pi/4$ rotates the x components of the state vector $|\psi(t)\rangle$ to be along the z axis. These results are summarized in Table 2.1.

Measurement Basis	Rotation
X	$\pi/2$ -pulse with $\phi = 0$
Y	$\pi/2$ -pulse with $\phi = 3\pi/4$
Z	No rotation

Table 2.1: Summary table of what rotations are needed to measure in what basis.

Ramsey Interferometry Ramsey interferometry is a technique which measures, to great precision, the detuning of a driving field from the transition frequency. In quantum information science, this is also known as phase kick back where the phase accumulation is due to the field detuning and the Ramsey $\pi/2$ -pulses are the Hadamard gates. It is the primary technique for accuracy measuring transition frequencies as well as detecting AC Stark shifts which shift transition frequencies. In fact, we will use Ramsey oscillations in Sec. 3.3.5 to

align high power 355 nm laser beams to the ion.




Figure 2.10: Ramsey oscillation with 0.13 *seconds* of coherence time. The drive frequency is purposefully detuned to see an oscillation. The fit frequency is 109.0 ± 0.1 Hz. The $\pi/2$ -pulses are driven by a microwave horn.

The experimental sequence is as follows. An ion is Doppler cooled (possibly ground state cooled as well) and optically pumped into the dark state. A $\pi/2$ -pulse, typically around the y -axis in our lab (i.e. $R(\theta/2, \pi/2)$), rotates the quantum state into an equal superposition. We wait a time T and apply the same rotation. The rotation can be around the other axis, just not z , (i.e. $R(\theta/2, 0)$) which will change the phase of the oscillation. This is useful for long drives near resonance because it helps to avoid confusing the decay of the oscillation with the oscillation itself [97]. Finally, the quantum state is measured. If the drive was exactly on resonance and there are no sources of $R_z(\phi)$ rotations, the ion will have completed a full π -pulse ending in the bright state. However, a detuning in the driving field, or equivalently an uncompensated AC Stark shift, δ will effectively rotate the quantum state during the wait time T by $R_z(\delta T)$. This rotation becomes evident in the z -basis when the second rotation

occurs as the $\pi/2$ -pulse will map the phase rotation to the z -basis amplitudes (identical to the phase kick back protocol).

In our lab, the longest maintained superposition state between hyperfine level has been measured to be at least 0.13 *seconds* as shown in Fig. 2.10. This extremely long coherence is achieved almost for free due to the atomic properties of the hyperfine state. In fact, the limiting factor is primarily magnetic field fluctuations. Over an hour of coherence has been observed when such noise sources have been carefully mitigated [28]. Additionally, we can easily measure the transition frequency between the two hyperfine qubits levels to Hz precision (< 1 ppb). In fact, a two-story NMR experiment down the hall went offline. Once they finished, we remeasured the hyperfine qubit transition and found it shifted by ~ 657 Hz.

Ground State Cooling

A trapped ion, or group of trapped ions, may be cooled down to its motional ground state—the zero-point energy state. This is desirable because then the beginning of each experiment may have the same known motional pure state $|n = 0\rangle$ with high probability. A common technique to ground state cool is called resolved sideband cooling, or just sideband cooling (SBC), [5, 98]. Electromagnetically-induced-transparency (EIT) cooling [99] or continuous sideband cooling [100] are also efficient strategies with EIT cooling working as much as $\times 10$ faster across a wider bandwidth though with a slightly higher final $\bar{n}$ than SBC [101, 102].

There are two principal steps in SBC and assuming the ion(s) has been prepared in the qubit state $|0\rangle$ by optical pumping (see. Fig. 2.11). The first step is to drive a red sideband π -pulse. This will flip the qubit state and, more importantly, lower the motional state quanta by 1: $|0\rangle |n\rangle \leftrightarrow |1\rangle |n - 1\rangle$. Second, the qubit state is reset but the motional state is approximately preserved by optical pumping: $|1\rangle |n - 1\rangle \leftrightarrow |0\rangle |n - 1\rangle$. The net effect is an ion back in the original qubit state $|0\rangle$, but one lower motional state $|n - 1\rangle$. Cycling these two steps many times will drive the ion(s) to the motional ground state.




Figure 2.11: Diagram of a sideband cooling cycle using a stimulated Raman transition for $^{171}\text{Yb}^+$. (a) A red sideband stimulated Raman transition lowers the motional state by 1 quantum with Rabi frequency $\Omega_{n,n-1}$. (b) Optical pumping on a fast transition (with a decay rate of Γ) approximately preserves the motional state while resetting the qubit state in preparation for another cycle of cooling.

The spontaneous decay leg of optical pumping sets the lowest achievable $\bar{n}$ for resolved sideband cooling [6]. The characteristic decay rate γ_{op} sets an effective linewidth which can lead to off-resonant coupling to other motional states. Consider when an ion is near, or in, the motional ground state. If the decay rate is fast enough, there could be a significant probability of spontaneously decaying into the first excited state. Trapped ion systems typically have harmonic frequencies $\gtrsim 2\pi \times 1$ MHz so optical pumping rates of ~ 1 μs are typical. We currently operate with an optical pumping decay time somewhere between 0.9 - 1.2 μs .

In the limit of mostly occupying the motional ground state $|n=0\rangle$, the probability of remaining in $|n=0\rangle$ after spontaneous decay (as part of the optical pumping step) is

$$p_0 \approx 1 - \left(\frac{\gamma_{\text{op}}}{2\omega}\right)^2 \quad (2.50)$$

where ω is the mode frequency being cooled. For $\omega = 2\pi \times 1$ MHz, and $\gamma_{\text{op}} = 1/(1 \mu\text{s})$,

the final ground state population be approximately $p_0 = 0.9937$. Similarly, the steady state $\bar{n}_{\min} \approx (\gamma_{\text{op}}/2\omega)^2$ [5], which is roughly $\bar{n}_{\min} = 0.006$ for this example calculation. In practice, imperfect red sideband π -pulses and other experimental conditions will raise $\bar{n}_{\min}$ above its theoretical limit.

Ion Thermometry

Ion thermometry is the measurement of the average motional state of a trapped ion $\bar{n} \equiv \langle a^\dagger a \rangle$, or the motional state probabilities p_n [98]. Since errors of gate operations and other critical quantum manipulations often depend on $\bar{n}$, accurate thermometry is needed to ensure the ion(s) are cold enough for the desired experiment or fidelity constraints. There are two historical methods for measuring the ion motion (with our newly developed method in chapter 5): the “ratio” method and the “SVD” method.

The ratio method [103] estimates the average harmonic state $\bar{n}$ of a thermal distribution $p_{\text{th}}(n) = \bar{n}^n/(\bar{n} + 1)^{n+1}$ by using the unique property $p_{\text{th}}(n + 1) = p_{\text{th}}(n)\bar{n}/(\bar{n} + 1)$, which property comes from the underlying Boltzmann distribution used to compute $p(n)$ (see Sec. 2.2.2).

Given a red sideband Rabi oscillation

$$\begin{aligned} P_{|1\rangle}^{\text{RSB}}(t, \Delta) &= \sum_{n=1}^{\infty} p_{\text{th}}(n) \frac{\Omega_{n,n-1}}{2\tilde{\Omega}_{n,n-1}} \sin^2 \left(\frac{\tilde{\Omega}_{n,n-1} t}{2} \right) \\ &= \frac{\bar{n}}{\bar{n} + 1} \sum_{n=0}^{\infty} p_{\text{th}}(n) \frac{\Omega_{n,n-1}}{2\tilde{\Omega}_{n,n-1}} \sin^2 \left(\frac{\tilde{\Omega}_{n+1,n} t}{2} \right) \end{aligned} \quad (2.51)$$

and a blue sideband Rabi oscillation

$$P_{|1\rangle}^{\text{BSB}}(t, \Delta) = \sum_{n=0}^{\infty} p_{\text{th}}(n) \frac{\Omega_{n+1,n}}{2\tilde{\Omega}_{n+1,n}} \sin^2 \left(\frac{\tilde{\Omega}_{n+1,n} t}{2} \right) \quad (2.52)$$




Figure 2.12: Ratio method. (a) Red and blue sideband drives immediately after sideband cooling. (b) The computed ratio. The ratio is not a constant value implying the motional population is nonthermal.

their ratio is a function of $\bar{n}$ for any time t or frequency detuning Δ

$$r \equiv \frac{P_{\uparrow}^{\text{RSB}}(t, \Delta)}{P_{\uparrow}^{\text{BSB}}(t, \Delta)} = \frac{\bar{n}}{\bar{n} + 1}. \quad (2.53)$$

Solving for $\bar{n}$ in terms of r ,

$$\bar{n} = \frac{r}{1 - r} \quad (2.54)$$

In the time or frequency domain, the same information about $\bar{n}$ can be measured and show be constant for both.

Figure 2.12 shows an example measurement with $\Delta = 0$ and varying the time of the Rabi drive. In Fig.(a) 2.12, a red and blue sideband oscillation is shown after ground state cooling. The blue sideband dynamics are strong, and the red sideband is not oscillating at

all as expected for the motional ground state. This leads to a nonconstant ratio however, shown in Fig.(b) 2.12. If the ratio is not constant, then the underlying distribution does not follow a thermal distribution implying the ratio method would give an erroneous estimate of $\bar{n}$. Indeed, in chapter 5, I will dive into how it leads to *underestimates* of $\bar{n}$ and other consequences in detail [104, 98, 100].

The SVD method [62] is a frequency-domain analysis of a red or blue sideband Rabi oscillation. In this method, $\Omega_{n,n'}$ is independently calculated, and its contribution to the overall Rabi oscillation is constructed into a rectangular matrix (dimension $M \times N$) with M time steps taken in the experiment, considering N harmonic states of interest, and elements $b_n(t) = \sin^2(\Omega_{n,n-1}t/2)$. This matrix acts on the harmonic distribution vector ($N \times 1$) to produce a vector representing the measured fluorescence at each experimental time step. For example, a blue sideband oscillation would be constructed as follows

$$\begin{pmatrix} b_1(t_0) & b_2(t_0) & \dots \\ b_1(t_1) & b_2(t_1) & \dots \\ b_1(t_2) & b_2(t_2) & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} p(0) \\ p(1) \\ p(2) \\ \vdots \end{pmatrix} = \begin{pmatrix} P_{\uparrow}^{\text{blue}}(t_0) \\ P_{\uparrow}^{\text{blue}}(t_1) \\ P_{\uparrow}^{\text{blue}}(t_2) \\ \vdots \end{pmatrix}. \quad (2.55)$$

Using singular value decomposition (SVD), the rectangular matrix is pseudo-inverted to solve for the harmonic distribution vector. Once this vector of $p(n)$ is known, the average occupation is found by calculating $\bar{n} = \sum np(n)$.

Figure 2.13 shows the results of an SVD analysis. An ion is cooled near the motional ground state and then left in the dark for 1 ms. Ambient conditions, such as fluctuating electric fields heat the motional state of the ion. The blue sideband oscillation of Fig. 2.13(a) has an irregular oscillation to it, and SVD captures this with excellent agreement. The resulting motional state probabilities are shown in Fig. 2.13(b) with probabilities $p_{n=0} = 0.4906$, $p_{n=1} = 0.3103$, $p_{n=2} = 0.0895$, $p_{n=3} = 0.0350$, $p_{n=4} = 0.0645$, and $p_{n=5} = -0.0203$.

The SVD method may give unphysical results. Probability constraints such as $0 \leq p_n \leq 1$ are not incorporated into the analysis as evidenced by the negative $p_{n=5} = -0.0203$ value. It




Figure 2.13: SVD method. (a) Blue sideband data points used in the SVD method. The blue sideband was driven after 1 ms of ambient heating increased the motional state probabilities. The SVD fit is the black solid line and shows good agreement. (b) Histogram of motional state probabilities as determined by the SVD method.

is simply a pseudo-matrix inversion. Additionally, in cases with poor blue sideband coherence or low contrast oscillations, the motional state probabilities may oscillate and blow up. In these cases, lower the number of phonon states considered in the analysis until reasonable probabilities are estimated. Finally, there seems to be a noise floor in the probabilities estimates. Anecdotally, we have been suspicious of probabilities roughly $\leq 4\%$. Shot noise, detection errors, and the experimental repeatability of the blue sideband drives are most likely at play. How to analytically estimate the influence these factors have on SVD is unclear and remains for future work; numerically we can use Monte Carlo methods as described below.

The SVD method does not estimate uncertainties. To estimate the error, we use Monte

Carlo error propagation with the follow steps. First, the blue sideband data points are randomly sampled assuming each point is Gaussian distributed about the measured mean with the standard deviation set by the measured error per point. Approximately 5000 samples or more may be taken to get clear results. Second, other experimental factors that play into the SVD analysis such as Rabi frequency and the Lamb-Dicke parameter are randomly sampled based on measured data as well—again assumed to be Gaussian distributed. Additionally, the number of phonon states considered in the SVD analysis is varied across samples. Each sample is run through the SVD analysis. Over the many samples, each predicted motional state probability typically forms a skewed histogram. Finally, the median value is numerically estimated along with the asymmetric 68 % confidence interval. These are the reported probabilities and associated error bar values, respectively.

2.4.4 Mølmer-Sørensen Interaction

The Mølmer-Sørensen (MS) scheme for entangling trapped ions operates by virtually coupling through the vibrational modes of the ion crystal effectively driving two-body Ising interactions $\sigma_x \otimes \sigma_x$ (or $\sigma_y \otimes \sigma_y$) between the qubit states of separate ions [46, 45, 47]. This is the most common method for driving many-body interactions in trapped ion systems. In the last two decades, MS gates have steady improved and now boast fidelities from 0.9989(7) to $1_{-0.0017}^{+0}$ with and without lasers [18, 19, 20, 21].

The operating principal beyond the MS gate is actually quite similar to a stimulated Raman transition in that they both virtually couple through intermediate states to drive an effective two-level system. To see this in action, consider two trapped ions. They will have two shared motional modes along one axis: the center-of-mass (COM) and tilt modes. For clarity, we will only focus on the COM mode, though many experiments have coupling predominately through the tilt mode since it heats much more slowly [17]. Next, apply symmetrically detuned red and blue sideband drives as shown in Fig. 2.14. The detuned tones will virtually couple through intermediate qubit and motional states shown in Fig.




Figure 2.14: Schematic energy level diagram of the MS scheme. Red lines represent detuned red sideband; blue lines represent detuned blue sideband. Only one mode is shown. Dashed lines are virtual states.

2.14. With the right choice of detuning δ and red and blue sideband Rabi frequency (similar to a stimulated Raman transition), the coupling to motion is adiabatically eliminated and the dynamics are effective two-level system which is a *joint* two-qubit system $|00\rangle \leftrightarrow |11\rangle$. The incredible result of this is that a $\pi/2$ -pulse drive (entangling pulse) will leave the two ions in one of the entangled Bell states!

The MS Hamiltonian driving these entangling dynamics falls typically falls two categories: the “slow” gate and “fast” gate. We operate in the slow regime, so that will be the focus here. I will give a brief description of the mathematics as derived many times in other references [46, 47, 45, 97, 105] and quote the result. By adding the red and blue sideband Hamiltonians (Eqs. (2.41) and (2.43)) together with a symmetric detuning such that $\delta_{\text{red}} = -\delta_{\text{blue}}$, the motion and spin operators become separable. The time evolution of this Hamiltonian is then computed using a Magnus expansion. Only the first two terms are nonzero. Higher order terms go to zero due to commutations relations between ladder operators. Alternatively,

second-order perturbation theory will also give the same result as was original done in the first proposals [46, 47, 45]. Either way, the effective Hamiltonian experienced by the trapped ions is

$$H_{\text{MS}} = \Omega^2 \frac{\hbar \Delta k^2}{2m} \sum_{i < j} \sum_k \frac{b_k^i b_k^j}{\mu^2 - \omega_k^2} \sigma_\phi^i \sigma_\phi^j \quad (2.56)$$

where m is the mass of an ion, Δk is the net wave vector of the Raman transition, ω_k is the harmonic frequency of the k th mode, μ is the beatnote detuning from the carrier transition, $\sigma_\phi = \cos \phi \sigma_x + \sin \phi \sigma_y$, and b_k^i is the i th ion's participation in the k th mode (i.e. the i th element of the k th eigenmode vector).

Equation (2.56) is often defined in quantum simulation or computation contexts as

$$H_{\text{MS}} = \sum_{i < j} J_{ij} \sigma_\phi^i \sigma_\phi^j \quad (2.57)$$

with

$$J_{ij} = \Omega^2 \frac{\hbar \Delta k^2}{2m} \sum_k \frac{b_k^i b_k^j}{\mu^2 - \omega_k^2}. \quad (2.58)$$

By manipulating the detuning—or adding multiple detunings—a variety of qubit-qubit couplings can be realized as detailed in Ref. [65]. Experiments have typically operated with a single detuning μ outside of the modes leading to an approximate power law dependence in the J_{ij} couplings [7].

In the lab, we are in the early days of working with these drives. We have driven the MS gate on 2 in a linear chain and 3 ions in a 2D triangle. Figure 2.15 shows the state dynamics for 2 ions. The ions were prepared in the $|00\rangle$ state after ground state cooling. Driving the MS gate using bichromatic, red and blue sideband, tones symmetrically detuned—as prescribed—the two-qubit system coherently evolves between $|00\rangle$ and $|11\rangle$. For details on all the calibration experiments leading up to driving an MS gate, begin by reading the theses of Rajibal Islam [97] and Crystal Senko [106].




Figure 2.15: MS oscillation of 2 ions. MS oscillation frequency 1202 ± 5 Hz. Decay rate is 3609 ± 480 μ s.

Witnessed Entanglement

This data alone is not enough to confirm the two ions are entangled at the $\pi/2$ -pulse time of the MS drive. An entanglement witness, akin to the Bell state tests, can verify whether a state is entangled or not [107]. If the measured witness value is above a particular threshold, the state is provably entangled. If the measured witness is below the threshold, nothing can be concluded. A convenient witness for trapped ion systems is to pair MS oscillation data with what is known as a parity oscillation. This was first demonstrated with 2 and 4 ions, but it will also work with an odd number of ions if an additional single-qubit rotation is applied [17].

A parity oscillation, or parity check, observes the coherent interference of the $|0 \dots 0\rangle$ and $|1 \dots 1\rangle$ parts of the quantum state $|\psi\rangle$. Hence, it can give information about the off-diagonal elements of the density matrix $\rho_{00,11}$. This is needed to estimate the fidelity that the state of the ions $|\psi\rangle$ is a GHZ state $|\text{GHZ}\rangle = (|0 \dots 0\rangle + |1 \dots 1\rangle)/\sqrt{2}$. The experimental procedure is the prepare an entangled state by driving a $\pi/2$ MS pulse then varying the measurement

basis (see Sec. 2.4.3). Repeat this experiment for various measurement basis angles as shown in Fig. 2.16(b). Coherent interference will show itself as a $A \cos(N\phi)$ oscillation where N is the number of ions, ϕ is the measurement basis angle, and A is the amplitude of the oscillation and is twice the off-diagonal density matrix $\rho_{00,11} = A/2$.

From Ref. [17], the fidelity of a state $|\psi\rangle$ being a genuine entangled state GHZ state $|\text{GHZ}\rangle$ is given by

$$F = \frac{1}{2} (p_{00} + p_{11}) + \rho_{00,11} \quad (2.59)$$

where p_{00} (p_{11}) is the probability of measuring $|00\rangle$ ($|11\rangle$) and $\rho_{00,11}$ is the off-diagonal density matrix element. From Fig. 2.15(a-b), we can determine that $p_{00} = .504$ and $p_{11} = 0.369$ at the $\pi/2$ -pulse time of $208 \mu\text{s}$. The half amplitude of the parity oscillation (Fig.(c) 2.16), $\rho_{00,11} = 0.292$. In total, $F = 0.729$! To be provably entangled $F > 0.5$ [17], which this data exceeds. By following the MS scheme, the two ions have been genuinely entangled.




Figure 2.16: Data used to estimate GHZ state fidelity of 2 entangled ions. (a) MS oscillations. (b) Count data has been labeled as p_{00} or p_{11} based on thresholding method. Fits help estimate p_{00} and p_{11} exactly at the $\pi/2$ -pulse time. (c) Parity oscillation fit to cosine. Half amplitude is used in fidelity estimate Eq. (2.59).

Chapter 3

Apparatus

Today’s modern laboratories are marvels of unusual apparatus and electronic sophistication while simultaneously jury-rigged with duct tape, tin foil, and solder. That is not meant as a degrading comment, but rather one that hints at a key art of theoretical and experimental physics: to know what is important and what can be approximated.

The chapter layout is as follows. First, the ion traps we have operated at outlined followed by a brief mention of the vacuum system which houses them. Laser systems are detailed as well as a microwave horn used for single qubit rotations. Imaging the ion is next following by the real-time control system. Examples of typical operations in the lab are described with some troubleshooting suggestions.

3.1 Rod and Blade Traps

We have historically built rod and blade style ion traps [54, 108, 55]. Rod style traps are robust; easier to design, construct, and assemble; and have a proven track record. This was the first trap architecture we successfully trapped an ion. As of this writing, we are operating a rod trap, but I partially fried it in early 2023. Although, I brought it back to life supporting the notion that rod style traps are fairly robust. We have assembled a backup rod trap with improvements to the compensation electrode positions just in case, but it has not been needed yet. The blade style traps are more sophisticated in design, construction, and assembly [55] with the goal of trapping and imaging 2D radial ion crystals—in our case. Our first blade trap was manufactured in pieces by Translume, but the gold layer on the RF electrodes unexpectedly boiled off making trapping impossible. See Ref. [108] for more details. We designed and constructed a new blade style trap using in-house wire EDM machining, electro-polishing, and diamond sandpaper resulting in an impressively small 7

nanometer average roughness [55]. After successfully trapping up to 29 ions in a 2D crystal [55], we have retired this blade trap and are now assembling a new trap manufactured by Translume. The main advantage in the Translume micro-structured ion traps is micron level tolerances in alignment, critical dimensions, and electrode symmetries. We anticipate the high precision alignment will result in less axial micromotion, lower heating rates, and higher secular frequencies than previous traps.

3.1.1 Electrode Geometries




Figure 3.1: Digitally rendered diagram of the rod trap from Ref. [54]. Panel (a) shows a far field image. Panel (b) shares the same viewing plane as the imaging system and labels the ion to DC electrode distance $z_0 = 355$ micron. Panel (c) labels the ion to RF electrode distance $r_0 = 460$ micron.

Rod style traps typically have at four parallel “rod” electrodes as shown in Fig. 3.1. One pair of opposing rods carry the RF voltages, and the other pair are grounded relative to the RF signal. These electrodes create the oscillation quadrupole potential detailed in Sec. 2.1 which effectively results in harmonic confinement. The DC endcap electrodes are far from the ion typically with a sharp needle-like geometry. These “needle” electrodes confine harmonically along the axial direction of the trap (z -axis of Fig. 3.1). Typically voltages for the studies in this work operate anywhere between 200 – 400 V of RF and 0.02 – 30 V DC.

The ion to DC electrode distance $z_0 = 355$ micron. The ion to RF electrode distance $r_0 = 460$ micron. From secular frequency measurements using the ion as a detector, we have estimated the geometric factor to be approximately $\kappa \simeq 0.12$. The view of the trap with respect to the imaging system is shown in panel Fig. 3.1(b), but the ion size has been greatly exaggerated in the figure and the imaging system has a much smaller field of view (see Sec. 3.4 for imaging system details). Not shown in Fig. 3.1 are compensation electrodes. These electrodes, sometimes called shim electrodes, are needed to fine tune the ion position. We currently have 3 compensation electrodes which run parallel with the four RF and ground rods.

Blade, or open-endcap, style traps open up the endcaps such that ions can be viewed where the needle electrodes would be for a rod trap; see Fig. 3.2. For comparison to the rod trap geometry, consider the view of Fig. 3.2(a). An ion chain would be confined into the page, but a 2D radial crystal would be in the plane of the page exactly in position for site-resolved imaging (see Ref. [55]). The DC blades are segmented into three parts. The outer two sections play the same role as the DC needles from the rod trap. The middle segment can also be used to manipulate the ion positions. Two rod compensation electrodes (seen in Fig. 3.2(a-b)) add additional control of the aspect ratio of a 2D radial ion crystal.




Figure 3.2: Images of blade style trap from Ref. [55]. Panel (a) is a front view of the trap as viewed from the imaging system perspective. Panel (b) is a zoomed in schematic showing the RF, DC, and compensation electrodes. Panel (c) shows the solid RF blade, segmented DC blade and on-board capacitors for RF pick filtering.

3.1.2 Voltage Delivery

To trap an ion, static and oscillating voltage fields need to be delivered to the trapping electrodes with minimal noise—particularly at the secular frequency of the ion. I will first detail the static, or DC, voltage channels and associated noise filtering then the radio frequency (RF) channel as shown in Fig. 3.3. Lastly, I briefly outline the magnetic field coil and atomic oven source electronic components.

There are two sets of electrodes with static (DC) voltage. The first set are the distant and pointed endcap electrodes which produce a harmonic potential to good approximation. The second set are compensation, or shim, electrodes which fine-tune the ion position or curvature of the trapping potential. Most of the DC signals for all these electrodes begin with the ARTIQ Zotino DAC as shown in Fig. 3.3 (see Sec. 3.5.5 for more info on the Zotino). The Zotino can programmatically output -10 to 10 V signals. We amplify four different DC




Figure 3.3: Part circuit diagram, part schematic of voltage delivery components and filtering stages.

channels with R4G120-0.3 Matsusada power supplies. In remote control mode, a 0 - 10V input signal from Zotino is amplified $\times 12$. These signals are then routed to a breakout box. Some Zotino channels are directly connected to the breakout box. Currently, a Rigol DP832 controlling two compensation electrodes at 1.72 V and 25 V is also connected to the breakout box. All these trap electrode signals are fed into a breakout box with Mini-Circuits BLP-2.5+ MHz low-pass filters. Capacitors (68 nF) soldered to ground initially filter out high frequency noise. A d-sub cable carries the DC channels from the breakout box to an array of inductors (1 Henry Tridac C-24X) for additional filtering. The d-sub cable then goes down to the vacuum chamber where a custom LP filter on a PCB does one last round of filtering. These LPs were originally placed there to prevent the high voltage RF signal used for trapping from transmitted out of the vacuum chamber on the DC lines. The LP PCB is wired to a 9-pin subminiature-c Accu-Glass cable and ultimately the Accu-Glass vacuum port and trap electrodes.

The current rod trap was experiencing up to 20 V Vpp of RF pickup on the DC channels. It was disrupting electronics across the lab including the wavemeter and PID locks. The RF signal is most likely from capacitive coupling between the RF electrodes and the DC electrodes. The RF electrodes and the DC electrodes are in close proximity and share a parallel geometry. The blade trap did not have this issue since there were LPs in the vacuum chamber. The rod trap does not have LPs in vacuum however. I constructed a custom LP

PCB and copper tube housing that mounts directly to the air side of the vacuum port to address this issue. The 10 nF SMD capacitors soldered onto the PCB reduced the pickup down to an acceptable 500-350 mV V_{pp}.

The RF channel begins with an HP 8648A signal generator near 23.82100 MHz. The exact output power varies depending on the attenuation in the line and power locking circuitry. The signal generator output is immediately LP filtered with an SLP-30+ (DC - 32 MHz) to clean up any harmonics. Although, this does not seem to be a significant source of noise. Next, the signal is amplified by a high power ZHL-5W-1+ amplifier. The RF is routed into a ZFBDC20-61HP-S+ bi-directional coupler oriented with the output port going toward the trap electrodes and the out-CPL going to an oscilloscope for monitoring signals reflected from the trap. The output signal going to the trap passes through a custom built biased quarter-wave helical resonator [109, 55, 108] the output of which enters the vacuum and attaches to the RF electrodes. For a future helical resonator, we will add a capacitive voltage pick-off so the locking circuitry can be rearranged to stabilize the resonator output. Locking the output of the helical resonator is critical to stable motional mode frequencies since the motional frequency is linearly proportional to RF voltage. Most voltage drifts happen due to temperature and mechanic related drifts of the resonator, therefore a pick off after the resonator is needed to counteract these noise sources [110, 55].

We drive a ~ 3.9 Gauss magnetic field at the ion which is oriented perpendicular to the optical table. The magnetic field coil is nominally driven at 1.4 A (16.21 V) by a WON ODP3033 power supply. The coil holder is 3D printed from a temperature resistant ABS and wrapped with 231 windings. The coil is geometrically registered to the ion position by a 3D printed cap secured to the top of the vacuum chamber head. The large coil diameter of 11.6 cm provides protection against orientation perturbations. The expected magnetic field strength at the ion is much higher than what the ion measures. We believe this is due to the metal of the vacuum chamber. A future project will be to increase the number of windings of the coil so we can reach ~ 5 Gauss. Increasing the current is not an option as the coil is

already running warm and could soften the 3D printed material to the point of mechanical failure.

We can drive two Yb ovens inside the vacuum chamber: one with ^{174}Yb and the other with ^{171}Yb [108]. We use an HP E3614A at 0.57 V and 2.59 A to heat our ^{171}Yb oven for daily loading. A BNC cable connects the power supply to RF filters (to reduce RF pick-up happening inside the vacuum chamber) and then to the vacuum leads attached to the oven.

3.2 Vacuum System

The vacuum system is a heavy stainless-steel housing wherein the ion trap lives and operators. Inside the vacuum chamber, the air pressure is in the ultra-high vacuum (UHV) regime $\sim 10^{-9}$ - 10^{-11} Torr or lower.

Despite this low pressure, collisions from background gas particles can occur. Stray atoms and molecules in the vacuum system may collide or even react with ions in the trap disrupting experiments, heating ions, and potentially breaking trapped ions free of the trap altogether. The Langevin rate can be used to estimate this background collision rate. It gives an upper limit on whether a background neutral atom will collide with an ion, and it is discussed in greater detail in Sec. 4.1.9 of Ref. [5] and references therein. An abbreviated calculation is computed here. Assuming the background gas is mostly H_2 molecules, the vacuum system walls are at room temperature 300 K, and the pressure is 10^{-11} Torr, the collision rate is $\gamma_{\text{Langevin}} \approx 5 \times 10^{-4} \text{ s}^{-1}$, or once every 30 minutes. We periodically notice background collisions when an ion spontaneously goes dark, however we have anecdotally observed such events on the time scale of hours implying the chamber might be at a lower pressure than the assumed 10^{-11} Torr.

To achieve such low pressures, the vacuum components use ConFlat (CF) flanges for joints and are rated for UHV. Four different pumping technologies are employed at various stages. First, an Agilent Technologies TPS Bench pumps down the chamber using an IDP-3 scroll pump and TwisTorr 84 FS turbo pump capable of reaching 0.75×10^{-8} Torr. A

starcell Agilent VacIon Plus 20 ion pump and non-evaporative getter (CapaciTorr Cf 35 D400/B200) reduce the pressure into the UHV regime or lower. The specifics of the pump down and baking process greatly impact the final achievable pressure, and these details are carefully outlined in Ref. [108].

3.3 Laser Systems

Lasers (light amplification of stimulated emission of radiation) have been a revolution technology for the 20th and 21st centuries with applications ranging from precision measurement to scanning a barcode at the grocery store. A laser’s precision in frequency, time, and space can make them difficult to work with, but is also what gives an experimenter exquisite control of atomic systems with quantum degrees of freedom.

3.3.1 Frequency Reference

A device that can measure a 100s of THz laser frequency with MHz precision (10 ppb) is crucial to any atomic physics laboratory due to the narrow linewidth of many useful atomic transitions. Atomic transitions can have as narrow a linewidth as 1 ppt or smaller. A less extreme, but vitally import example, is a 369.5nm photon which is approximately 810 THz. It is resonant with the $^{171}\text{Yb}^+$ transition $|^2\text{S}_{1/2}\rangle \rightarrow |^2\text{P}_{1/2}\rangle$ transition, but that transition linewidth is only 20 MHz! And, some in the atomic physics community might consider that a large linewidth.

We measure the various laser frequencies with a HighFines WS Ultimate 10 wavemeter. An 8 channel mechanical switch samples the various light sources in the lab for the wavemeter to measure. All 8 channels can be monitored in software with wavelength accuracy gaurenteed up to 10 MHz. However, based on the response of trapped $^{171}\text{Yb}^+$ we have anecdotally observed ~ 2 MHz resolution over the course of months. A LabView VI (virtual instrument) collects the measured values from the wavemeter and coordinates locking the frequencies across multiple PIDs. This is typically done with a circuit PID board interfacing

with the associated laser’s controller. In the future, we may replace this LabView VI with an ARTIQ RPC [111].

The wavemeter requires a frequency reference to determine the absolute value of a measured wavelength. Our reference is a frequency stabilized Helium-Neon (HeNe) laser from REO operating at a wavelength of 632.990950 nm. The quoted stability is ± 1 MHz over 1 hour and ± 2 MHz over 8 hours with 1.5 - 3 mW of output power. In all my time in the lab, I have never touched this laser. See the optical layout in Fig. 3.4.




Figure 3.4: Optical layout of the HeNe, 399 nm, and 935 nm laser systems created using the component library [112]. The HeNe is primarily directed toward the wavemeter via fiber optic cable and acts as a frequency reference. The 399 nm and 935 nm diode lasers are directed toward the trap and wavemeter via fiber optic cables.

3.3.2 399 nm Laser System

The “399” laser, as we call it, provides the first of two photons used to photoionize ^{171}Yb or ^{174}Yb . The 399 nm photons excite neutral ^{171}Yb or ^{174}Yb from their respective $^2\text{S}_{1/2}$ to $^2\text{P}_{1/2}$ states. Once excited, a second incident photon at 369.5 nm or 355 nm will free the excited electron producing $^{171}\text{Yb}^+$ or $^{174}\text{Yb}^+$.

We operate a MOGLabs 398.91 nm diode laser in a Littrow configuration at a wavelength of 398.91045 nm (751.528250 THz). The laser can output ~ 25 mW at the max diode injection current of 60 mA. (Newer models which we operate on a different experiment can output as

much as 120 mW when driven at 190 mA of current.) The first optical element outside the laser is an optical isolator (ThorLabs IO-3D-405-PBS) with $\sim 80\%$ transmission. The beam is then split, fiber coupled, and routed to the wavemeter and ion through SQS FC/APC Nufren SM 405nm optical fibers (see Fig. 3.4). In the past we have locked this laser’s frequency, however, not locking it has not impacted our loading rate. The neutral atoms exits the oven at high temperature and therefore a broad velocity distribution. This leads to a Doppler broadened transitions on the order of 100s of MHz and removes the need for MHz accuracy. The beam going to the ion meets a mechanical Edmund Optics c-mount shutter powered by a custom built PNP transistor circuit with 5 V power supply input and Zotino channel logic input. Applying 5 V closes the shutter; 3.5 V opens the shutter. If the shutter does not open properly, try lowering the 3.5 V voltage by 5 %. Repeat until success. Be careful to not go any lower than need be. Since the circuit uses a PNP transistor, a small current going *into* the Zotino when the control voltage is lower than 5 V, and this can damage Zotino. A 1 k Ω resistor is in place to protect Zotino but still be careful. The light then couples into the optical fiber that leads to the ion.

A future improvement to trapped ion systems, in general, would be to sweep the 399 nm laser frequency over a few 100 MHz during ion loading. Sweeping the frequency compared to a fixed frequency would excite neutral atoms across a wider range of velocities and hence excite more atoms increasing the likelihood of photoionization.

3.3.3 935 nm Laser System

The “935” laser, as we call it, supports all $^2S_{1/2} \rightarrow ^2P_{1/2}$ transitions such as Doppler cooling, optical pumping, and detection. There is a small branching ratio of 0.5% from $^2P_{1/2}$ to $^2D_{3/2}$ (see Fig. 2.1). Like a leaky hole in a dike, this branching ratio needs to be “plugged” so the fast $^2S_{1/2} \rightarrow ^2P_{1/2}$ cycling transition can continue cycling without an ion getting stuck in the meta-stable $^2D_{3/2} |F = 1\rangle$ state for ~ 53 ms. The 935 nm laser, and 3.06952 GHz EOM, resolve this issue by pumping from $^2D_{3/2} |F = 1\rangle \rightarrow ^3[3/2]_{1/2} |F = 0\rangle$. The transition will

either succeed and spontaneously decay back to $^2S_{1/2}$ or be reattempted until success. These events are relatively rare compared to the fast $^2S_{1/2} \rightarrow ^2P_{1/2}$ cycling transition and hardly impact the ion's scattering rates [113]. However, without this laser the ion will be quickly pumped into the $^2D_{3/2} |F = 1\rangle$. In fact, we often temporarily block the 935 nm laser beam, pumping the ion into the dark $^2D_{3/2} |F = 1\rangle$ state, to measure the photon background of the photon detector.

We operate a MOGLabs 935 nm diode laser in a Littrow configuration at a wavelength of 935.187687 nm (320.569295 THz). The laser can output 80 mW at the max diode injection current of 200 mA. (Newer models which we operate on a different experiment can output as much as 95 mW when driven at 200 mA of current.) The first optical element outside the laser is an EOT 04-850-00090 optical isolator with > 30 dB of isolation, $> 88\%$ transmission, and bandwidth from 840 - 960 nm. The beam is then split and routed to the wavemeter and ion through SQS FC/APC S Nufren SM 780 nm optical fibers (see Fig. 3.4). Using the measured frequency output from the wavemeter, a LabView program and electronic PID circuit lock the laser frequency through the diode laser controller. The beam going to the ion first meets a mechanical Edmund Optics c-mount shutter with an identical drive system as the 399 nm laser. When the shutter is open, the light is fiber coupled to a polarization maintain fiber and propagates through a PM-0S5-10-PF EOSpace EOM (electro-opto modulator). Driven with an RF frequency of 3.06952 GHz at 7.2 dBm, the EOM produces sidebands with 2% relative intensity. The light is then coupled back into a yellow FC/APC optical fiber and routed to the trap.

3.3.4 369.5 nm Laser System

The 369 nm laser system fulfills multiple roles: laser cooling, state preparation, and state detection as theoretically discussed in Sec. 2.3.2. It is also the most intricate laser system in the lab (see Fig. 3.5). The end goal for this system is to produce at least ~ 10 s of mW of 369.5 nm light, which in the early days of the lab could not be achieved with diode laser

technology, but now may be feasible. The high power is needed because the beam will be split at least four times to accomplish the separate tasks of laser cooling, optical pumping, and detection. The overall optical layout is shown in Fig. 3.5 and discussed below.




Figure 3.5: Diagram of the 369 nm optical layout created using the component library [112].

Like many laser systems, to get high power the system starts with a Nd:YAG crystal lasing 1064 nm. In our case, this is done by a Lighthouse Photonics Sprout-H laser. The 1064 nm light is frequency doubled in the Sprout-H laser head ultimately outputting 10 W of brilliant green 532 nm light. The 532 nm light enters an M Squared SolsTis which consists of a titanium-sapphire (Ti:saph) crystal in a bow-tie cavity. For 10 W of green 532 nm light pumping the Ti:saph crystal, 2.75 W of red 740 nm is output. After the SolsTis, a small pick-off sends ~ 15 mW of ~ 740 nm light to the wavemeter to be measured and subsequently frequency locked to 739.052409 nm (405.644381 THz). The final leg of the journey is an M Squared ECD-X which consists of a frequency doubling crystal in a bow-tie cavity. With 2.75 W of 740 nm light entering the ECD-X, 1 W of 369.5 nm light can be produced. Under typical conditions, we operate the Sprout as low as possible 5 - 8 W. This leads to ~ 80 mW of 369.5 nm light which is plenty of power with room to spare for our current requirements.

We have seen the ECD-X double efficiency drop in recent years. Carefully cleaning the front face of the doubling crystal (often with a single wipe) has dramatically improved the doubling efficiency.

With the 369 nm light source ready to go, the light is routed to four sections of the optical table to complete four specific tasks: Doppler cooling, detection, more Doppler cooling, and optical pumping. The amount of light power going into each section is controlled by half-wave plate, which rotates the light polarization, and a polarizing beamsplitting (PBC) cube which reflects only the vertical light by 90° (see Fig. 3.5). Acousto-optic modulators (AOMs) are used in each section for MHz-level frequency control and fast shutter control, as will be described below.

The first section is the main Doppler cooling beam. The light polarization is rotated and then passes through a 14.747811 GHz EOM from QUBIG GmbH (EO-WG14.7M2). The EOM is driven with 28 dBm of RF power producing sidebands of 1 % relative intensity. This is 1/5 the quoted EOM performance, but we do not believe this is limiting Doppler cooling. Next, the light diffracts inside a Brimrose CQF-420-100 AOM with the first order beam up-shifted by 420 MHz, which will be 10 MHz red-detuned from the $^2S_{1/2} |F = 1\rangle \rightarrow ^2P_{1/2} |F = 0\rangle$ transition. The first order beam is the main Doppler cooling beam for the experiment. The zeroth order is labeled projection plus and is used to cool very hot ions. The zeroth order beam has a ThorLabs SHCP025 shutter which we typically only open during ion loading. Both the first and zeroth orders are fiber coupled and routed to the ion.

The next section is detection. Light passes through another Brimrose CQF-420-100 AOM this time driven at 430 MHz, which will be exactly on resonance with the $^2S_{1/2} |F = 1\rangle \rightarrow ^2P_{1/2} |F = 0\rangle$ transition. Only the first order beam is fiber coupled and routed to the ion. Polarization is controlled with a half-wave plate.

The third section is protection—another Doppler cooling beam. The AOM here is the same model as the previous ones but driven at 390 MHz to cool mildly hot ions. It is on whenever the main Doppler beam is on. It has been quite useful during experiments if kept

at low enough power.

The last section is optical pumping. The light first passes through a 2.105 GHz EOM (QUBIG GmbH E0-T2100M3). The 25.8 dBm RF drive produces 1st order sidebands with $> 30\%$ relative intensity. Another AOM driven at 430 MHz provides the necessary frequency and shutter control. Only the first order beam is fiber coupled and routed to the ion. There is an additional optical pumping AOM for future projects relating to individual addressing.

3.3.5 355 nm Laser System

The 355 nm laser system is used to drive single and two-qubit quantum gate operations on the qubit states of the trapped ions. These are essential quantum operations needed by any quantum simulator. There are theoretical and experiment reasons for choosing 355 nm light for quantum stimulated Raman transitions. The theoretical reasons were covered in Sec. 2.4.2. Here we focus on the experimental side.

Experimentally, high power 355 nm light is readily accessible. The company Coherent produces a mode-locked laser series called the Paladin which produces Watts of 355 nm light from a frequency-tripled Nb:YVO4 laser in a bowtie cavity. We operate an 8 W version with a repetition rate of 80.9 MHz and 12 ps wide pulses. The train of laser pulses makes a frequency comb centered at 355 nm with a bandwidth of ~ 100 GHz which can be used for nearly full power drives of the 12.6 GHz hyperfine qubit splitting without the use of an EOM [97]. Because the 355 nm light is so far detuned, a significant amount of power is needed, but the coincidence of the Coherent Paladin offering Watts of power overcomes this hurdle as well. Although it is commercially available, the laser is expensive and dangerous. The high-power UV light can melt cables, ignite note cards, and even burn skin. Among other safety procedures, we enclose the entire beam light in dark foam board for our protection and beam stability. UV light is notoriously difficult to couple through fibers, and it is typically out of spec for spatial light modulator technology useful for individual addressing. 355 nm light has been a workhorse for $^{171}\text{Yb}^+$ trapped ion experiments, but alternative wavelengths

are being explored [114, 115].




Figure 3.6: The 355 nm optical layout created using the component library [112].

The optical beam path for our 355 nm laser is shown in Fig. 3.6 and is laid out as follows. A 99:1 BSF10-UV beam sampler immediately picks off a few mW of the 8 W beam and directs it to a fast Alphas UPD-30-VSG-P photodiode for locking the beatnote—which will be discussed in Sec. 3.3.5. The beam power is then managed with a half-wave plate, polarizing beamsplitting (PBS) cube, and Thorlabs BT600 high power beam dump. We currently discard as much as ~ 7.8 W of power while maintaining a 180 kHz Rabi frequency. An Isomet M1136B-FS80L-3 with 80 MHz center frequency, 3mm wide aperture, 20 MHz bandwidth and 91 % efficiency sends laser power fluctuations into the first order beam in a “noise-eater” configuration such that the zeroth order beam is power stabilized. A downstream beam pick off monitors the zeroth order beam and outputs to a PID controlling the Isomet AOM.

Following the power stabilization, the beam is split into two arms by a 50-50 beamsplitter: the “inner” and “outer” arms. The arms take nearly identical paths. Both arms pass through Brimrose CQF-210-40-355 AOMs with a center frequency of 210 MHz, 40 MHz bandwidth, maximum RF drive of 2 W, and operate with a measured 68 % efficiency. The 13 mrad separation angle produces two distinct beams with the 0th order reflected by a D-mirror and dumped. We are suspicious that thermal effects from driving the AOMs at 2 W of

RF are causing a drifting Rabi frequency at the ion. We currently drive the AOMs at 1 W to mitigate the thermal effect, but it is still present. A future upgrade such as water cooled AOMs or always-on driving may resolve this. The inner arm AOM conditions the light frequency based on the beatnote lock effectively stabilizing the frequency comb of the two arms when they interfere at the ion (see Sec. 3.3.5 for more details). The outer arm AOM is free to vary its frequency and amplitude. We use this AOM to control the effective frequency of the stimulated Raman transition whether it be a red, carrier, or blue transition as well as pulse shaping (see Sec. 3.5.6).

After the AOMs, an intermediate focus with 125 mm best form lenses adds fine alignment control for both arms. The polarization is then manipulated with half-wave and quarter-wave plates in both arms. The light pulses traveling the two arms must reach the ion at the same time. A custom-built timing delay, or optical delay, stage is in the inner arm beam path. It consists of a 25 mm right triangle knife-edge prism mirror MRAK25-F01 which reflects light into a transparent right-angle prism (CVI P180-100-355-UV) which, due to total internal reflection, reflects the light back to the other side of the knife-edge prism mirror and continuing its way to the ion. By translating the transparent wedge relative to the knife-edge prism on a ThorLabs linear translation stage, the travel time of the inner arm light can be finely adjusted. We only use this for fine tuning the timing delay when initially aligning the 355 nm beams.

The last optical elements are best form lenses which focus the light onto the ions. The outer arm has a 150 mm best form lens mounted on a LineTool model A right-handed translation stage focusing the light down to an estimated 30.8 micron $1/e^2$ beam diameter. The inner arm has a 125 mm lens mounted on a left-handed version of the same translation stage focusing the light down to a 25.7 micron $1/e^2$ beam diameter. Both stages have been upgraded with differential micrometers (Newport 110-502NPT) for greater alignment control.




Figure 3.7: Frequency stability measurement of the 156th comb tooth (~ 12.626 GHz) during warm up. Laser was turned in current mode. Frequency was measured after the bandpass filter and high frequency amplifier (see Fig. 3.8

Beatnote Stabilization

To perform a Raman transition with the 355 nm laser, there needs to be, among other considerations, a 12.642812 GHz (plus the magnetic field shift) frequency difference in the two laser beams. Second order driving of an EOM could achieve this for a CW laser, but at the cost of poor efficiency and phase related problems [97]. The 355 nm laser is effectively a frequency comb with a ~ 100 GHz bandwidth making it capable of bridging the 12.6 GHz transition with nearly full power. The interference of the two combs at the ion produces another frequency comb often called the beatnote in the literature [116], and it is that beatnote which needs to be frequency stabilized to drive clean resonant transitions. On warm up, the 156th comb tooth drift as much as 200 kHz as shown in Fig. 3.7. After warming up, the repetition rate will continue to drift ~ 100 s of Hz per hour [116], which needs to be corrected for successful quantum gate operations.

The general principle of stimulated Raman transitions with two frequency combs is as

follows. The two 355 nm arms are each a frequency comb. Consider the interference of the m th outer arm comb tooth and $(m+n)$ th inner arm comb tooth where m and n are positive integers. This produces a beatnote at the frequency $n\nu_{\text{rep}}$ where ν_{rep} is the repetition rate of the laser, or comb tooth spacing (80.9 MHz in our case). The AOMs in each arm can then shift the optical frequency by $|\Delta\nu_{\text{AOM}}| = |\nu_{\text{AOM1}} - \nu_{\text{AOM2}}|$ such that the desired level splitting (the hyperfine levels in our case ν_{HF}) is resonantly bridged in the typical lambda configuration $\nu_{\text{HF}} = n\nu_{\text{rep}} - |\Delta\nu_{\text{AOM}}|$ [116]. Said another way, the atom absorbs a photon from the $(m+n)$ th inner arm comb tooth and emits into the m th tooth of the outer arm comb to drive simulated Raman transitions [116] with an overall wave vector $\Delta k_{355} = \sqrt{2}k_{355}$.

To stabilize the beatnote, we follow the theory and setup from Ref. [116], though newer, more digital, approaches have been implemented since then [117]. The first pick off of the 355 nm light (see Fig. 3.6) continuously measures the frequency of the 156th comb tooth near 12.626 GHz. Other teeth are suppressed by a Fairview band-pass filter, and the 156th tooth is amplified (Mini-Circuits ZVA-183-S+). The tooth is mixed with a local oscillation $\nu_{\text{LO}} = 12.420822$ GHz to produce a signal within the AOM bandwidth ~ 207 MHz. A low-pass filter with cutoff frequency 250 MHz (Mini-Circuits SLP-250+) further suppresses any unwanted teeth and unwanted mixer harmonics. The signal is quite weak, so it is amplified by three Mini-Circuits ZFL-500LN+ amplifiers and then passed into a custom-built phase-locked loop. See Fig. 3.8 for the circuit diagram.

Before going into the phase-locked loop details, I will brief walk through why this scheme stabilizes the beat note following the arguments made in Ref. [116]. After the mixer, the signal that would be used to drive one of the AOMs is now $\nu_{\text{AOM1}}(t) = n\nu_{\text{rep}}(t) - \nu_{\text{LO}}$. Since the AOM frequency is shifted relative to the fixed local oscillator, any fluctuations in the tooth frequency are countered by the frequency driving the AOM. This can be seen in the

calculation of the overall beatnote frequency

$$\nu_{\text{beatnote}} = n\nu_{\text{rep}}(t) - \nu_{\text{AOM1}} + \nu_{\text{AOM2}} \quad (3.1)$$

$$\nu_{\text{beatnote}} = n\nu_{\text{rep}}(t) - n\nu_{\text{rep}}(t) + \nu_{\text{LO}} + \nu_{\text{AOM2}} \quad (3.2)$$

$$\nu_{\text{beatnote}} = \nu_{\text{LO}} + \nu_{\text{AOM2}} . \quad (3.3)$$

The local oscillator ν_{LO} replaces the noisy comb tooth $n\nu_{\text{rep}}(t)$. In practice, the local oscillator will have some noise, but this is an engineering problem and easily reduced by purchasing a high-quality signal generator such as an HP 8672A.




Figure 3.8: Circuit schematic of the beatnote lock as outlined in Ref. [116] and updated from Ref. [105]. Blue lines represent high-frequency cables.

The phase-locked loop is technically not needed, but in practice it is essential. The original photodiode signal is noisy due to laser amplitude noise and the heavy amplification. This would carry over to the AOM drive and consequentially imprint onto the laser leading to noisy, exaggerated off-resonant coupling (see Ref. [116] for measurements of this effect).

The phase-locked loop outputs the needed shifts in the AOM frequency to maintain a locked beatnote while taking on the phase noise characteristics of the HP 8640 signal generator [116]. The phase noise of the AOM drive signal is then only limited by the quality of the signal generator. In practice, the signal has more phase noise than the signal generator, but it is significantly reduced compared to no phase-locked loop. The output of the phase-locked loop is blocked by a fast RF switch which in turn is controlled by ARTIQ through a TTL channel. When the switch opens, the RF signal is amplified up to 2 W for driving the 355 nm inner arm AOM. Based on quantum gate performance, the lock has been working well for several years.

355 nm Alignment

Aligning the 355 beams in space and time may take a seemingly infinite amount of time the first try, but somehow it can be completed in a couple of days every subsequent attempt. The procedure to align non-co-propagating 355 nm Raman beams is roughly as follows: triple check optical path lengths, align arms separately by detecting AC stark shifts with Ramsey experiments, tune optical path length and alignment with Rabi oscillations.

Triple checking the optical path length is essential. It will save you much time later on. Measure the total distance of each arm from the beamsplitter to the ion. A string or plastic measuring tape work best. Repeat this measurement three times and average your results. Adjust mirror positions and repeat until the measurements are within 3-5 mm of each other. Our laser model is quoted to have a 12 ps pulse width implying a packet size of 3.6 mm in space. Hence the want for 3-5 mm accuracy.

Now that the optics are in roughly the correct positions, it is time to spatially align each beam to the ion. This is accomplished one beam at a time. Before performing any experiments though, operate the laser in a very low power mode such that you can see the beam on a notecard and it neither catches fire nor hurts your eyes with its brightness. With this very low power beam, map out the boundaries of the trapping area: view the beam

on a notecard after the trap and slowing adjust the final lens translation stage. Record the micrometer readings at relevant dimensions for the trap geometry. For example, in a rod trap, I would record the micrometer readings when the laser hits the very tip of the needles and when the laser hits the top and bottom rf rods when center between the needles. Calculate the micrometer position that will most likely align the beam to the ion and dial in the translation stage to this position. Additionally, the focus can be estimated in a similar way, particularly when directly the beam onto sharp electrode features [106]. Do this for both arms. We added this step to our alignment procedure recently, and it has been invaluable. It is quite helpful for zeroing in the initial alignment, but its true worth is realized when performing a repeated alignment on the same system. We have compared past alignment micrometer positions with new boundary measurements and aligned successfully to the ion. No Ramsey experiments needed.

The beams are now roughly aligned to the ion based on the trap geometry. If there's no Rabi oscillation signal, a more fine-tuned approach is needed. With an excess of σ^+ or σ^- polarization, a single beam can induce a larger and hopefully measurable differential AC Stark shift between the $^2S_{1/2}$ states and $2P_{1/2}$ and $^2P_{3/2}$ states [94, 106]. Adjusting the quarter-wave plate along with pointing the magnetic field along the direct of propagation, the correct polarization and quantization axis can be produced [106]. The AC Stark shift induced by an alignment beam is then measured uses a Ramsey oscillation driven with the $\pi/2$ pulses driven by a microwave horn near the ion. Adjust the final translate stage in your favorite search pattern until a shift is measured. In recent realignments to the ion, we have noticed that very high power UV light is enough to measure a Ramsey signal without needing tilt the magnetic field along the beam propagation direction, but the advantage of this method is relatively small. We typically measure a shift of several kHz if the beam is fairly well aligned, and its focus is near the ion.




Figure 3.9: Beam delivery excluding the 355 nm beams created using the component library [112]. Wavelengths are labeled at their fiber ports. The red shaded mirrors are dichroic optics.

3.3.6 Laser Beam Delivery

The laser beam delivery optics combines the wide range of wavelengths (369.5 nm, 399 nm, and 935 nm) in an overlapping, collimated beam using beamsplitters and Semrock dichroics as shown in Fig. 3.9 and directs them at the ion. The following is a list of the laser powers for each beam measured at their respective fiber ports:

- 369.5 nm doppler cooling: 42.5 uW
- 369.5 nm protection: 68.4 uW
- 369.5 nm protection plus: 348 uW

- 369.5 nm detection: 60 uW
- 369.5 nm optical pumping: 40.1 uW
- 935 nm: 12.8 mW
- 399 nm: 1.39 mW

The beam passes through a $\times 2.5$ telescope and is then focused down with a 150 mm best form lens (LBF254) to a few tens of micron waist at the ion. The 355 nm laser system enters on the opposite side of the vacuum chamber for its own distinct control and alignment. See Sec. 3.3.5 for details about the 355 optical layout.

3.3.7 Microwave Horn

The microwave horn setup is much simpler than the 355 nm system, and it can also drive single-qubit gates. We typically use it for troubleshooting and 355 nm beam alignment, but it can demonstrate interesting results on its own. For example, more than an hour coherence time of a trapped ion qubit was demonstrated using a microwave horn [28].




Figure 3.10: Circuit schematic of the horn drive. Blue lines represent high-frequency cables.

We have two methods for driving the microwave horn (a Fairview horn waveguide antenna) at the desired 12.6 GHz hyperfine transition frequency. The first method mixes two signals to get up to the 12.6 GHz and controls the on-off behavior with an RF switch on one of the mixer inputs as shown in Fig. 3.10. We currently use this method due to the historical

lack of microwave switches at 12.6 GHz. This is no longer the case anymore and directly driving 12.6 GHz is possible. A future project will be to integrate this new high-frequency compatible switch (a ADRF5024-EVALZ-ND or EV1HMC547ALP3), remove the mixer, and generate the signal with a Lab Brick LMS-163 8-16 GHz signal generator. For brevity, I will only describe the currently used method which uses frequency mixing.

The goal of the circuitry is to drive the microwave horn at 12.6 GHz with at least 1 W of power. The horn is located just outside the front view port of the ion trap, under the imaging objective. The signal begins with an HP 8648B at 1989.998835 MHz and 7.0 dBm going through an RF switch (Mini-Circuits ZYSWA-2-50DR+) controlled by ARTIQ. The ~ 2 GHz signal is then mixed (Mini-Circuits ZX05-24 MH-S+) with an always on 10.652826 GHz source output from an HP 8672A. The mixer output is cleaned up by a Fairview 11.7 - 12.8 GHz bandpass filter. A high-frequency cable carries the signal down to a ZX60-153LN-S+ preamp and ZVE-3W-183-S+ amp, and ultimately driving the horn with 29 dBm of power. By controlling the RF switch on the ~ 2 GHz signal before the mixer, the 12.6 GHz drive can be precisely controlled: if the switch is off, the mixer output will be the always on 10.6 GHz signal, but this will be filtered by the bandpass filter and is far off-resonance.

3.4 Imaging System

The imaging system collects scattered photons from the ions which are then used to determine their quantum state. My personal opinion is that the imaging system steals the show. Seeing an individual, isolated atom—a building block of the universe forged by stars—is something more people should experience.

3.4.1 Light Collection

We image the ion with custom assembled optics in what is effectively a microscope as shown in Fig. 3.11. Starting with a 0.28 numerical aperture (NA) triplet (Special Optics 54-17-29-369nm) with $\times 5$ magnification, the light from the ion is focused down through a 100 mm




Figure 3.11: Schematic of the imaging system created using the component library [112]. After magnification by the triplet, best form, and doublet, the light is filtered (represented by green and red elements) then sent to either the camera or the PMT by a flip mirror. The blue line schematically represents scattered photons imaged to the PMT or Camera

diameter pin hole by a best form lens. A $\times 8$ magnification doublet then nearly collimates the light. The optics are held together in a ThorLabs cage system movable by a Newport 562-XYZ translation stage. After the doublet, there are two filters to isolate the 369.5 nm light signal from background photons. The 369.5 nm light enters a dark box which contains a flip mirror and photomultiplier tube (PMT). Cables going into and out of the dark box go through light baffles or sealed ports. If the flip mirror is down, the light is directed to the PMT. We use the PMT when we need high fidelity read out of one or two ions. If the flip mirror is up, the light is directed to an Andor iXon Ultra 897 EMCCD camera. We use the camera for loading, viewing the ion position in the trap, and imaging many ions.

Based on scattering rate measurements, the collection efficiency is 0.16%. Fairly typical for the NA of the imaging system [83]. The scattering rate experiment was performed by measuring the number of photons collected as a function of incident laser power. Subtract off background counts and fit the data to the scattering rate equations for $^{171}\text{Yb}^+$ from the thorough work in Ref. [81].

3.4.2 Photomultiplier Tube Detection

The PMT is a device that can detect individual photons. Incident photons generate free electrons which are amplified, or “multiplied”, to a detectable current. We operate a Hamamatsu H10682-210 PMT which comes with built-in photon counting circuitry. The detection area is 8 mm in diameter with peak sensitivity at 400 nm and 43 % quantum efficiency at 350 nm. We have observed that the PMT introduces nearly zero noise; the maximum dark count rate is specified to be $1 \times 10^{-2} \text{ ms}^{-1}$. The largest background noise is stray light from our detection beam $\sim 0.5 - 1.0$ counts per ms and stray light getting into the dark box $\sim 0.05 - 0.2$ counts per ms.

A single PMT is not used to construct an image of an ion, but it can accurately detect the number of photons emitted by an ion. PMTs typically have better quantum efficiencies than camera technologies making them the common choice for high-fidelity detection in trapped-ion systems. However, detecting multiple ions with one PMT will reduce the fidelity with which the number of bright ions is determined. The overlap of the Poisson distributions of the number of photons detected from N bright ions compared to $N + 1$ bright ions increases quickly as function of N (see Sec. 2.3.2). More overlap means there is less confidence in labeling how many ions were bright and therefore reduces detection fidelity. Overall, the PMT can reach fidelities in the 99% for single ion detection as discussed in Sec. 2.3.2.

Since a PMT outputs a pulse for each photon detected, PMTs can also be used to detection *when* a photon arrives. This time-of-arrival information can increase the detection fidelity by as much as an order of magnitude and reduce detection times by a significant fraction [83]. We typically count the total number of photons in a given time window as this still gives 98% fidelity, but time-of-arrival information has been successfully demonstrated in our lab using ARTIQ software and Sinara hardware. See Sec. 3.5.6 for more information on retrieving time-of-arrival information.

3.4.3 EMCCD Camera

The electron-multiplying charged coupled device (EMCCD) is the long and technical name for what we call the camera. The camera is a Andor iXon Ultra 897 EMCCD with a quantum efficiency of 30% near 369 nm. It detects individual photons on an array of pixels with each pixel physically $16 \text{ micron} \times 16 \text{ micron}$. Since imaging optics magnify by $\times 40$, each pixel effectively images a size of $0.4 \text{ micron} \times 0.4 \text{ micron}$ area at the plane of the ion. The camera can detect individual photons and capture an image at rates of ~ 100 s of microsecond mostly dependent on the size of the region of interest (ROI). The overall photon detection fidelity is worse than a PMT due to lower collection efficiency and increased noise, but an EMCCD camera has spatial information making it the better detector when there are many ions.

We operate the camera in one of two modes. In mode one, we acquire a live stream of images when the ion is irradiated by Doppler cooling light. This mode is used for loading, measuring the position of the ions, and general troubleshooting when speed is not essential. The settings are typically: acquisition mode single, 0.3 s exposure time, 300 EM gain, 10 MHz readout at 16-bit, and $1.7 \mu\text{s}$ vertical pixel shift speed with “normal” vertical clock voltage amplitude. The camera temperature is set to -85° C (no external cooler). These settings give a clear ion image under Doppler cooling light.

In mode two, the camera detects the quantum state of the ion as part of an experiment. The settings are the same as mode one except for the following changes: acquisition mode “kinetic”, triggering “fast external”, 0.5 - 1 ms exposure time, 1000 EM gain, and a custom ROI that excludes as many pixels as possible. These settings allow for fast image acquisition, large ion signal, and minimized background. We have most recently measured an average detection fidelity of 93.7% with these settings for one ion and have seen in the past as high as 98%. We have observed that increasing the readout rate beyond 10 MHz or increasing the vertical shift speed will cause a significant jump in dark counts.

Sometimes we have observed high dark counts or dark ion images that look like there is a dimmed ion as shown in Fig. 3.12. The main cause of this comes from the stochastic start of





Figure 3.12: Comparison of imaging an ion when with the stochastic start of the camera (label “Doppler light leakage”) and when adding a padding delay compensates for the stochastic start (labeled “No light leakage”). The color bar indicates detected photons per pixel. The top left image is of an ion prepared in the bright state. In the upper right plot, the ion is prepared dark, but it is visible due to the camera collecting Doppler light after the detection period. In the bottom right image, the ion is prepared dark, and the padding delay does not allow Doppler light to fall on the camera sensor. The maximum photons per pixel drops from 0.35 to 0.06, and no ion shape is visible.

each image exposure period. The camera is triggered by a TTL pulse to begin the exposure period. However, the camera will not actually begin the exposure until it has finished its current pixel “cleaning” cycle. The cleaning cycles clear out accumulated background signal which is critical for these cameras; clean out cycling is the idle mode of operation. We have observed an upper bound of $\sim 250 \mu\text{s}$ for the camera to finish a cycle. This is problematic

as it makes up a sizable fraction of the typical 500 - 1000 μs detection period. If the camera starts 250 μs late, for example, then consequentially the camera will be collecting photons for 250 μs after the detection laser beam has been off, and worse, it will be exposed for 250 μs into the next repetition of the experiment which is typically Doppler cooling. Under the Doppler cooling beam, ions that were dark will become bright and the resulting image will show an ion when it should have shown nothing. We have fixed this issue adding an extra 300 μs of detection beam time (see bottom half of Fig. 3.12). The camera will still only expose for the desired detection period, but the detection beam will have a 300 μs buffer to account for the camera’s stochastic start.

EMCCD cameras have additional noise sources compared to PMTs. EMCCD technology, in a typical high-gain configuration, follows a Poisson-Gamma-Normal (PGN) noise model due to various physical processes in the camera, which is two more distributions than just the Poissonian statistics of the PMT data [118, 119]. A thorough comparison of photon detector technologies and EMCCD noise sources—along with discussion on the ARTIQ Camera Grabber—can be found in Ref. [119].

3.5 Real-Time Control

Control systems harmoniously orchestrate the wide range of laboratory equipment and instruments with sufficient accuracy and repeatability. Modern atomic laboratories typically need nanosecond resolution, reliable control, and deterministic (defined later on) repeatability to avoid fragility across repetitions of the same experiment and years of use. In other words, an experimentalist wants to program a sequence of events and have the control system execute all those events—on or off, up or down, higher or lower—*exactly* the same way hundreds of thousands (possibly millions) of times across months or even years.

There is a software and hardware side to our control system which work in tandem. The software side is called the Advanced Real-Time Infrastructure for Quantum physics (ARTIQ) software control system. The second part is the Sinara hardware ecosystem [120]. With high-

level programming in ARTIQ and the hardware capabilities of Sinara to physically realize those commands, we can control almost every aspect of lab operations with *nanosecond* precision. ARTIQ was originally spun out of the Ion Storage Group at NIST and is now distributed by M-Labs Limited. A growing number of research groups and partners are supporting this project because it is rather unique: it is open source. Both software and hardware designs are available on their respective GitHub pages [121, 122] and M-Labs advertises its open source licensing.

3.5.1 Advanced Real-Time Infrastructure for Quantum physics (ARTIQ)

ARTIQ is a high-level open-source Python module for real-time control of experimental hardware. At the end of the day, the ARTIQ module (and its designated workflow) is the Python module we use to write and execute every experiment in the lab. We use it a lot to say the least; after only four years of operation, on November 29th, 2023, we ran our 250,000th experiment! As of this writing, ARTIQ experiments have generated over 4 Gigabytes data. ARTIQ accepts experiment commands written in Python, compiles these instructions, and then executes them on the Sinara Kasli crate dedicated FPGA. Though many operations need only millisecond or microsecond timing, critical quantum operations such as single- and two-qubit gates benefit greatly from nanosecond resolution as quantum control techniques which modulate drive amplitude, frequency, or phase can be readily implemented [123, 124, 64, 125].

The backbone of ARTIQ’s real-time control capability is the idea of “the timeline,” the set of all input and output events across all channels [111]. A user may put events on the timeline, shift the timeline forward *or backward*, and record event timestamps. These operations may seem quite abstracted away from a physical laboratory, but they cover all essential parts of an experiment as will be explained later. Further technical details on the timeline can be found under the ARTIQ docs section ARTIQ Real-Time I/O Concepts [111].

To program an experiment, starting from the top of the software stack, we write a new python class which inherits from `artiq.language.environment.EnvExperiment` and overloads

its methods: `build`, `prepare`, `run`, and `analyze`. This makes the anatomy of an ARTIQ experiment as follows (see Fig. 3.13).

build. This is where user inputs are defined, devices are initialized, and initial parameters can be read in. ARTIQ provides some basic input types:

- boolean `artiq.language.environment.BooleanValue()`
- float or integer `artiq.language.environment.NumberValue()`
- string `artiq.language.environment.StringValue()`

There is also `artiq.language.environment.EnumerationValue()` which is effectively a list of strings from which a user can select one option. We have used this, for example, in selecting the measurement basis or pulse shape type. Lastly, `artiq.language.scan.Scannable()` is effectively a range of equally spaced floats or integers. We use this, for example, in defining the time axis of a Rabi oscillation. We have yet to use `artiq.language.environment.PYONValue()`, but it may be a powerful option in some cases.

Devices are initialized in two ways with the only difference being what variable the device is defined as. First, using `.HasEnvironment.setattr_device(key)` which will set the name of the device to `.HasEnvironment.key`. Second, `.HasEnvironment.get_device(key)` which will return the device drive to whatever variable the user chooses. In our lab, we mostly use the second method to allow for descriptive variable names such as `pmt_ttl` instead of `ttl0` where no one would be able to remember what is attached to TTL channel 0. See Fig. 3.13 for example code.

Lastly, datasets from previous experiments are read-only in the `.build()` method. We use this feature to load an experiment with relevant global parameters. This is especially useful when user interface values need to be bounded to avoid damage to people or hardware. For example, we will import the maximum allowable drive for the outer 355 AOM to: 1)

avoid accidentally damaging the AOM and 2) allowing dangerously high amounts of UV light through.

prepare. Computational space to prepare for the experiment. Empty data arrays are initialized with `nan` (not a number) values; these will be replaced during the experiment. See the Fig. 3.13 example. All experimental values are pre-computed in this method such as converting all laser pulse time values to nanoseconds, complicated laser pulse sequences for ground state cooling or pulse shaping, etc. The input values the user selected are saved to a dataset at this point. Lastly, we compute the estimated execution time of the experiment and have the value logged for the user to see in case they may wish to cancel an unexpectedly long experiment.

run. Hardware commands are programmed in this method. This method comes with a kernel flag `@kernel` to indicate that the code will be executed on the Sinara core device (not the host PC). See the Fig. 3.13 example. The core device has a much stricter subset of acceptable Python syntax, data structures, and NumPy features. This allows the “softcore” CPU on the core device to be fast enough to interpret the commands and queue them on the timeline before they are meant to happen. Programming this method feels more like scheduling a calendar than typical linear programming. For example, if nothing is to take place, you must program this using the `delay(time)` function. If a command takes processing time such as counting pulses on a TTL channel, you must program in a delay in the schedule to allow the CPU however many microseconds of computational time it needs to process the TTL input before it goes back to scheduling the next event.

analyze. Computational space for processing results. We are back to “regular” Python code in this method. We typically use this method to lightly process the data such as compute the mean and standard deviation of a repeated measured, fit data, make a plot, and save relevant results. Sometimes the fits fail, however. It would be convenient if this function

could be called individually without rerunning the entire experiment; the ARTIQ browser feature promises to do this. Unfortunately, as of this writing, there is no documentation for the browser. Our solution has been to write “blank” experiments where just the analyze method has been overloaded and much more fitting control is exposed to the user. If an analysis initially fails, the user can run the corresponding blank experiment to retry the fit, or even fit historical data.

ARTIQ experiment scripts are saved in the repository folder. This folder may have subfolders and hidden experiments by putting an underscore “_” at the beginning of the class name. We have only fully tested and polished experiments in the repository. Old versions of experiments are saved outside the repository. Experiments under longer testing phases are saved in a testing subfolder. If they are not resolved promptly, they are moved out of the repository. Filenames are lowercase with underscores for readability.

Experiment python scripts are typically hundreds of lines of code each with a preference for line or inline comments instead of docstrings. In Table 3.1, static code metrics are shown for our current repository folder. The key metrics that reflect our Sisyphean quest for human readable code are Blanks and Comments / SLOC. The high blank line code is due to the spaced styling of the code for visual grouping and hence easier skimming. Additionally, we prefer line comments within an experiment for each step as shown by the large 50 % Comments / SLOC ratio. We have minimal multi-line docstring comments for experiment scripts as evidenced by the lower (Comments + Multi-line)/LOC percentage. They are too far away from the source code to be helpful. When explaining, or teaching about, what each line does in an experiment, we find that line and inline comments are best.

Finally, all data saved to an ARTIQ dataset during an experiment will be automatically saved to an h5 file under the **results** folder. Files are automatically organized by time to the hour. Filenames consist of the unique experiment identification number called the “rid” and the class name. As mentioned previously, ARTIQ experiments have generated over 4 Gigabytes data in its first four years of operations.

a)

```

1  from artiq.experiment import *
2  import numpy as np
3
4  class DetectionCheck(EnvExperiment):
5      """Detection Check"""
6
7      def build(self):
8          # Hardware setup
9          self.setattr_device("core")
10         self.setattr_device("urukul0_cpld")
11         self.doppler_cooling = self.get_device("urukul0_ch3") # Doppler cooling ADM
12         self.detection = self.get_device("urukul0_ch2") # Detection ADM
13         self.setattr_device("ttl0") # PMT channel
14         # GUI inputs
15         self.setattr_argument("detection_time", NumberValue(500, ndecimals=0, step=100))
16         self.setattr_argument("n_reps", NumberValue(100, ndecimals=0, step=10))
17
18     def prepare(self):
19         # Initialize array to later save pmt counts
20         self.pmt_counts = [np.int64(0) for _ in range(self.n_reps)]
21         # Import parameters
22         self.doppler_cooling_time = self.get_dataset("parameters.doppler_cooling_time")
23         self.doppler_atten = self.get_dataset("parameters.doppler_atten")
24
25     @kernel
26     def init_aoms(self):...
27
28     @kernel
29     def doppler_cool(self):...
30
31     1 usage
32     @kernel
33     def exp_start_up(self):...
34
35     @kernel
36     def run(self):

```

b)

```

48
49     1 usage
50     @kernel
51     def exp_start_up(self):...
52
53     @kernel
54     def run(self):
55         self.exp_start_up()
56
57         for i in range(self.n_reps):
58             # Doppler Cool (turn doppler and protection beams on)
59             self.doppler_cool()
60             # State Detection (pulse detection beam and count photons at pmt)
61             self.detection.sw.on()
62             gate_end = self.ttl0.gate_rising(self.detection_time + us)
63             with parallel:
64                 self.detection.sw.off()
65                 self.doppler_cooling.sw.on()
66             self.pmt_counts[i] = self.ttl0.count(gate_end)
67             delay(500 * us)
68
69         # Leave cooling beams on
70         self.doppler_cooling.sw.on()
71         # Save Data
72         self.set_dataset("detection.pmt_counts", self.pmt_counts, broadcast=True)
73
74     def analyze(self):
75         # Convert list to numpy array for simpler computation syntax
76         pmt_counts = np.array(self.get_dataset("detection.pmt_counts"))
77         # Compute mean counts per pulse time
78         mean_count = np.mean(pmt_counts)
79         # Print mean and standard deviation of counts
80         print('Mean number of counts:', mean_count)
81         # Save mean value to broadcasted database
82         self.set_dataset("detection.mean_count", mean_count, broadcast=True)
83

```

Figure 3.13: Minimal working example of a practical ARTIQ experiment. (a) The left half are lines 0 - 61; (b) the right half are lines 48 to the end. This experiment will measure the fluorescence of an ion irradiated by the detection beam and Doppler cooled beforehand by the doppler beam. This is repeated `n_reps` times. The `build()`, `prepare()`, `run()`, and `analyze()` methods are spelled out with the supporting function definitions not show for brevity.

ARTIQ is aware of what devices it can control through the `device_db.py` file. The file is a list of devices with specific dictionary formats. Default settings for hardware can be set here, though we have only modified it for the Phaser as described in Sec. 3.5.5. If you wish to control non-Sinara hardware such as a camera or wavemeter, remote procedure calls RPCs can be defined which allow ARTIQ to act, in limited way, as if the device was part of Sinara [111].

When using ARTIQ there are a few considerations I would like to briefly mention. First, experiments are executed sequentially. Parallel processes are possible within an experiment, but these will typically end once the experiment finishes. This is different than, for example, an LabView VI which may run in parallel with other VIs continuously. For example, a software PID is not a good match for an ARTIQ experiment. There are solutions to most situations (even that PID one I just mentioned), but the point is that it is important to be

aware of ARTIQ’s design principles and their consequences when programming or building experiments. Second, ARTIQ has historically had a slow compiler ($\text{ARTIQ} < 7$); this is changing in future updates $\text{ARTIQ} \geq 7$. It has never been burdensome to us, but it may be something to keep an eye on if you wish to do experiment scheduling such as an auto calibration routine [126].

The last consideration is an anecdotal opinion that I hope is valuable for new ARTIQ users. When ARTIQ first arrived, we were quite frustrated and at a loss for how to do anything. There are many unknown unknowns for first time users—the worst kind of roadblock. We quickly realized that ARTIQ and the associated Sinara hardware are an open source research *project* not a shiny *product*. Do not let this discourage you. Once that idea was settled in our minds, we spent time learning and experimenting. We can now do anything we dream up with ARTIQ. Remember that this is an *open source* effort and as such there is a community of people willing and able to answer questions:

- Read and reread the ARTIQ docs [111]
- Read over relevant GitHub issues [121] or wiki pages [122]
- Look up questions and ask your own on the M-Labs forum [127]
- Talk with students and postdocs from other groups. (This is how we figured out how to use the Phaser correctly.)

It truly is a community effort which is a rather unique situation compared to any other piece of equipment or software in the lab. In fact, we have had the *most* useful and successful support for this equipment than for anything else in the lab.

3.5.2 ARTIQ Dashboard

ARTIQ has a user interface called the ARTIQ Dashboard, although ARTIQ can be run headless. Our philosophy for experiment GUIs has been to maintain an agile and adaptable

GUI that gives the user as much control as possible. This has been a successful strategy working across a wide range of studies and ever-changing interests.

The dashboard primarily consists of five views as shown in Fig. 3.14:

- A list of experiments Fig. 3.14(a)
- Inputs for various experiments Fig. 3.14(b)
- Logging Fig. 3.14(c)
- Experiment scheduling status Fig. 3.14(d)
- Applets Fig. 3.14(e)

These windows consist of our “mission control” from which we submit experiments, monitor their status reports, plot data in real-time, and generally conduct ourselves scientifically.

Experiments are selected from the file explorer (Fig. 3.14(a)). Each experiment represents an individual Python script in the `repository` folder. Names of each experiment are defined from the docstring in the script (see Fig. 3.13). If there is no docstring, the experiment class name will be used. Starting an experiment class name with an underscore “_” will hide an experiment from this view. Subfolders are displayed with a toggle menu. Once an experiment is selected, the GUI will be immediately generated from the build method of that experiment class and displayed in the main window shown in Fig. 3.14(b).

The dashboard’s main window can display many experiment windows. All available input parameters are displayed for each experiment, respectively. Once the inputs are set, click submit. There is, unfortunately, a pesky bug for dual monitor Windows systems where new experiment GUIs pop-up beyond the far edge of view. It requires a few seconds of dragging the GUI back toward the others. This is a known issue with Windows but unresolved at the time of writing.

Once an experiment is submitted, its schedule status will pop up in the schedule window shown in Fig. 3.14(d). All logging related updates such as error reporting, or user




Figure 3.14: Image of our typical ARTIQ dashboard setup. Red lines have been added to highlight the main five views of the dashboard. In panel a) is the list of available experiments. Panel b) is a large window where experiment GUIs are displayed. The user can adjust input parameters and submit experiments here. Panel c) is the logging area. We program experiments to log anything from mundane experimental monitor updates to data analysis information. All of this is displayed in this panel. Panel d) reports the scheduling status: preparing, running, analyzing, deleting. Many experiments can be queued here if need be. They will be executed sequentially or by priority level. Panel e) contains data display panels from just numbers to full figures or images. These display panels are called applets.

programmed reports, are displayed in the log window (Fig. 3.14(c)). We often use this to report back expected experiment run time and fit parameters which we may want to record in the lab journal (a OneNote document) as mentioned previously.

Finally, the last view is really a collection of views called applets (Fig. 3.14(e)). Applets can display a number, plot, image, etc. With some work they can be customized as well, though we have avoided that additional complexity for the time being. Applets are aware of datasets and will update whenever its corresponding dataset updates. This allows plots to update in real-time as data comes in from an experiment. Applet plots can be quickly copied and pasted into an electronic experiment journal such as OneNote.

With all these displays and features, many experiments can be choreographed to make a serious scientific endeavor. All the way from daily calibrations to brand new experiments.

3.5.3 ARTIQ Data Flow

Data flow is how data from different experiments, or outside of ARTIQ, gets where it needs to go. Our philosophy is to use ARTIQ datasets as the centralized source of data as much as possible and only access saved data files for after-the-fact analysis. There are two main situations to consider. First, how to share global parameters across experiments. Second, how to use results from previous experiments to inform future experiments.

We define global values using a blank experiment. There is a laundry list of possible input parameters: laser cooling time, detecting time, AOM frequencies, etc. These are made accessible on the experiment’s GUI. When the blank experiment is executed, the values are saved to a ARTIQ dataset which any experiment can access when its turn comes along. This strategy allows us to be extremely agile while also centralizing variable definitions to avoid bugs.

We currently do not employ a methodology for automatically importing results from previous experiments. This is mostly applicable in the context of auto calibration routines or cascaded experiments. We currently consider these strategies to be an over optimization and require the user to record results and input them in manually for later experiments. However, it is an emerging area of research [126, 128], and my opinion is that ARTIQ would handle such strategies well.

3.5.4 Subartiq-lib

There are many duplicate processes across different experiments: laser cooling, state preparation, detection, initialized hardware, etc. As such, ARTIQ is often most effectively used when there is an additional software layer—of some variety—that sits on top of ARTIQ. NIST has the NIST Framework [129]. Oxford has ndscan [130]. Duke has DAX [131]. At

the time of writing, these are all the modules of which I am aware.

We maintain a philosophy of agility, control, and transparency when it comes to writing experiments, as has been mentioned before. However, this could lead to massive code duplication and over customization which will lead to fatal, troublesome bugs. To resolve this, we have avoided the class structure solutions of other groups. We find such strategies abstract code too much leading to too much overhead for our nimble group.

Instead, we have built a library of transparent functions. We have centralized all duplicate code in the lightest possible option—a python module called `subartiq.lib`. The module is much lighter than the previously mentioned projects. The module consists of a library of functions which only do one thing and cover everything from common hardware commands to common statistical distributions. We find this extremely adaptive and successful for our group. In accordance with our design philosophy, we do not need to push this optimization. Only functions that we expect to be shared by multiple experiments—and which have been thorough verified—need to be in `subartiq.lib`. This also greatly clarifies what functions do for first time users, and it does not deviate from the general ARTIQ workflow explained in the ARTIQ docs lightening our docs burden.

Table 3.1 shows some static code metrics which give evidence that we have successfully held to our design choices with acronyms defined in the Table and caption. The key metrics to observe are the blank lines to LLOC ratio, comments / LOC, and (comments + multi-line) / LOC. First, the blanks to LLOC ratio $\text{Blanks} / \text{LLOC} = 17\%$ shows that the library has healthy spacing between functions for visual grouping and easy skimming. Second, comments / LOC is low at only 9% in contrast the ARTIQ repository is 27%. This is because `subartiq.lib` is a library of concise functions primarily documented using multi-line docstrings. (ARTIQ repository scripts are long class definitions which benefit more from line or inline comments hence the higher comments / LOC ratio.) Last of all, (comments + multi-line) / LOC shows that almost *half* of all the LOC in `subartiq.lib` consist of docstring comments. This the most critical evidence that the library follows our design philosophy.

More specifically, this last metric shows that the library of functions is well documented allowing quick comprehension by the user. Physicists may not be programmers, but we try in our lab.

Metric	subartiq_lib	repository
Lines of Code (LOC)	4687	12796
Logical Lines of Code (LLOC)	1518	6394
Source Lines of Code (SLOC)	1590	6938
Comments	438	3455
Single Comments	450	3272
Multi-line Strings	1850	250
Blanks	797	2336
Comments / LOC	9%	27%
Comments / SLOC	28%	50%
(Comments + Multi-line) / LOC	49%	29%
Maintenance Index	A	A
Cyclomatic Complexity	A (2.09)	A (2.37)

Table 3.1: Static code metrics for subartiq.lib and our ARTIQ repository which holds our experiment scripts. Metrics are defined by Radon python package as the following [132]. Lines of code is the total number of lines of code. Logical lines of code is the number of lines of code that contains exactly one statement. Source lines of code is the number of source lines of code (not necessarily corresponding to the LLOC). Comments is the number of comment lines. Multi-lines strings is the number of lines which represent multi-line strings. Blanks are the number of blank lines (or whitespace-only ones). Maintenance index measures how maintainable (easy to support and change) the source code is. Cyclomatic complexity (also called McCabe number) is the number of linearly independent paths through the code. Additional information can be found in Ref. [132]

The subartiq.lib module makes as thin of a layer as possible on the software stack. It consists of three main folders: experiment, analysis, and comms as shown in Fig. 3.15. The main

scripts within the experiment folder mimic the overloaded ARTIQ methods: `build.py` has functions useful for the `.build()` method; `prepare.py` has functions useful for the `.prepare()` method; and `experiment.py` has functions useful for the `.run()` method. There are also `calibrations.py` and `utilities.py` functions which support `.prepare()` or `.run()` computations. The analysis folder also reads in a fairly straightforward manner: `functions.py` is a list of mathematical functions which are often used in fitting data; `fits.py` has helper functions for fitting data; and `analysis.py` is a list of helpful miscellaneous functions. The comms folder is an experimental feature for quickly extracting a table of values from ARTIQ datasets.




Figure 3.15: Partial file structure of the `subartiq_lib` module. Folder are denoted by black diamonds, files by black circles. The most used functions are found in the scripts on the far right. The file structure is meant to mimic ARTIQ as much as possible for clarity.

When updates to `subartiq_lib` are pushed to GitHub, GitHub Actions run the available tests. Checks are performed for pep8 code styling [133] and pep257 docstring styling [134]. Documentation is automatically generated from the function docstrings using our own document generator script.

3.5.5 Kasli Crate




Figure 3.16: Photo of our Kasli crate as is. The different instruments are outlined in labeled and outlined in different colors.

The Kasli crate holds the open source Sinara ecosystem hardware used to control the experiment. It is also an open-source project and condenses much of traditional ion trap hardware into an impressively small rack mountable form factor of 3U.

Our crate shown in Fig. 3.16 has the following devices: the Kasli core device, Grabber, DIO SMA, two Urukuls, Sampler, Zotino, and Phaser. In the Sinara fashion, each device has a unique name; I will briefly highlight their more colloquial names and how we employ them.

The Kasli, called the core device, is an FPGA (XC7A100T-2FGG484I) carrier board which can control other hardware within the crate or even other crates [122]. It is the hub from which all the real-time control commands come and go. We currently have all twelve Eurocard Extension Modules (EEMs) filled, so we would need another crate if we want new equipment. The Kasli is currently referenced to an external 10 MHz Rb clock (FS725 Standard Research Systems) instead of its 125 MHz internal clock.

The Grabber is a CameraLink port from which a CameraLink cable can directly stream data from an Andor camera to ARTIQ. This strategy avoids slow USB or Ethernet connec-

tions [122]. Pixel data is collected from predefined ROIs and stream along ROI engines at 40 MHz for Andor camera. However, we have found that the number of ROIs is limited, and data acquisition is slow enough to cause significant blocking delay during experiments. Currently, we operate the Andor camera from its own software (SOLIS) where frames are streamed over USB to the experiment PC after the experiment is over.

The DIO SMA board is an eight channel isolated digital input/output device with a 5 ns minimum pulse width [122]. We refer to these as TTL channels since we use them for logic. Grouped in two sets of four, the TTL channels can be either input or output. Currently, the PMT and an RF period trigger are connected to the input TTL channels. On the 4 TTL outputs channels is a switch for the microwave horn output, a switch controlling the 355 nm inner arm AOM RF drive, and the camera external trigger.

The Urukul is a 4-channel DDS-based frequency synthesizer [122]. It offers a frequency range of < 1 to > 400 MHz with subHz resolution, attenuation control from 0 - 31.5 dBm with 0.5 steps, and fast update control. Urukul drives all the 369 Brimrose AOMs (with ZHL-1-2W-S+ amplifiers): optical pumping, protection, detection, and doppler. When driving multiple tones simultaneous, such as during an MS gate, it is necessary to have phase coherence between tones. If one wanted to combine Urukul outputs for such as purpose, deterministic phase control is possible but requires careful adjustments which can be found on the M-Labs forum and Urukul GitHub pages [127, 122].

The Sampler is an 8-channel, 16-bit analog-to-digital converter (ADC) with a 1.5 MSPS update rate across all channels simultaneously [122]. We use this device to measure voltages. Gains can be adjusted so the sampler is sensitive to input ranges from ± 10 mV to ± 10 V. The amount of data acquired is something to keep in mind as the data will need to be stored in memory and eventually saved to a ARTIQ master dataset. A large dataset could lead to significant lag for the host PC.

The Zotino is a 32-channel, 16-bit digital-to-analog converter (DAC) with a 1 MSPS update rate divided between the channels [122]. We use it to output voltages to the trap

electrodes. Some channels act as TTL signals for slow optical shutters such as the imaging flip mirror, the 935 and 399 nm c-mount shutters, and the protection+ ThorLabs shutter.

Last of all is the Phaser. The Phaser is a quad channel 1.25 GS/s RF generator card with dual IQ upconverter and dual 5 MS/s ADC and FPGA [122, 135]. Per the ARTIQ docs, the kasli supports two 25 MS/s data streams for up to 5 numerically controlled oscillators each (5 for each phaser channel) [111]. This means, for one phaser channel, the fastest the channel output will be updated is 40 ns ($1 / 25 \text{ MS/s}$). Additionally, each oscillator within that channel can then only be scheduled to change a parameter—such as frequency, amplitude, or phase—every 8 ns ($1 / (25 * 5)$). Attempting to schedule RTIO events faster than this will lead to unexpected outputs and will, unfortunately, not throw any errors. An additional consequence is that if you want to make changes across multiple oscillators which come into effect simultaneously, the changes must be scheduled within the 40 ns update rate of the channel. (See examples below.) Always test and confirm outputs. The output passes through a ZX75BP-204-S+ bandpass filter, ZFL-500LN+ preamp, and finally a ZHL-03-5WF+ amplifier before driving the 355 nm outer arm AOM. We use the Phaser to drive all quantum operations and ground state cooling. For MS gates, two phase coherent tones are required (or more). Pulse shaping capabilities are also possible though still developing (see the M-Labs forum about MIQRO [127]). A single Phaser channel (it comes with two) has up to five oscillators with full frequency (20 MHz bandwidth), amplitude, and phase control.

For deterministic phase coherence between the oscillators, the phase accumulators must be reset with deterministic timing relative to the coherent operations. This is accomplished by doing two things. First, align each repetition of an experiment to the 320 ns Kasli communication cycle. Otherwise, there would be 40 ns jitter as noted by pathfinder49 in GitHub issue #1659 [121]. In ARTIQ 7, a convenient function `.get_next_frame()` was added to resolve this issue. By moving the time cursor to this future frame using `at_mu(time)` the start of any experimental repetition could have deterministic alignment to the Kasli communication cycle as shown in the code below


```
# Measure the frame timestamp
self.phaser.measure_frame_timestamp()
# Set the "now" of the timeline cursor to a frame far enough into that future
    that it doesn't cause an RTIOUnderflow Error
at_mu(self.phaser.get_next_frame_mu() + delay_pad)
```

where `delay_pad` is some positive integer. It is also a good idea to add some padding to the delay to avoid an `RTIOUnderflow` error that is an integer number of 320 ns. We use `set delay_pad = 96  $\mu$ s (=  $300 \times 320$  ns)`.

Second, clear the phase accumulation in the oscillators before the start of each coherent operation. For a straight MS drive this will do just fine. If phases are to change during the experiment, there may be subtle considerations for *when* the phase accumulation is restarted. For example, it may be necessary to have the phase accumulation reset right when the gate operation starts, or it may be acceptable to have it reset several microseconds beforehand. We have a defined function in `subartiq_lib`, but here is how it might look fully written out if three oscillators were being used

"Aligns relative drive phases for phase coherent gate operations"

```
# First clear the phase accumulators by setting clr=1
self.phaser_ch0_osc0.set_amplitude_phase_mu(amps[0], phases[0], clr=1)
delay_mu(8)
self.phaser_ch0_osc1.set_amplitude_phase_mu(amps[1], phases[1], clr=1)
delay_mu(8)
self.phaser_ch0_osc2.set_amplitude_phase_mu(amps[2], phases[2], clr=1)

# Need to wait 40ns since the first instructuion before new instructions can
    take place
delay_mu(40 - 2*8)
# Save timestamp for when phases start accumulating again
```

```

self.phaser_ch0_osc0.set_amplitude_phase_mu(amps[0], phases[0], clr=0)
delay_mu(8)
self.phaser_ch0_osc1.set_amplitude_phase_mu(amps[1], phases[1], clr=0)
delay_mu(8)
self.phaser_ch0_osc2.set_amplitude_phase_mu(amps[2], phases[2], clr=0)

```

Each board is discoverable by ARTIQ through the `device_db.py` as mentioned previously. This lists all the hardware as well as chip settings. When we first got the Phaser, we calculated and updated some of the trf chip settings. Specifically need to set 8 in binary for the setting `{"tx_div_sel": 0b11}` to get the correct frequency range for the Phaser to drive the 355 nm outer arm AOM near 210 MHz.

3.5.6 Advanced Use Cases

There are some advanced use cases of ARTIQ that make it quite powerful but may also need some illustration and demonstration. That is what this section is for. The main cases I will highlight here are: multiple direct memory access (DMA) sequences; retrieve time-stamped events such as photon time-of-arrival timestamps from a PMT; negative delays; and pulse shaping with the Phaser.

Some event sequences happen so quickly that the Kasli CPU cannot schedule everything in time (throwing an `RTIOUnderFlow` error). Resolved sideband cooling, which may schedule thousands of events each spaced by only a few microseconds or less, is a prime example. ARTIQ offers a solution called the direct memory access DMA extension [111]. This feature stores pre-defined sequences of output RTIO events into the core device's SDRAM (system memory) and can play them back at high speeds [111].

In some cases multiple DMAs may be helpful. Consider an MS gate with pulse shaping. The experiment would need one DMA for ground state cooling, another for ramping up the Phaser output power and another for ramping down the Phaser output power. To record multiple DMAs, first record each sequence using its own separate handle with

`self.core_dma.record(handle1):` in the `.run()` method. Once complete, *then* retrieve all the playback handles one at a time. The sequences can be played back by the normal command when desired. If these lines are called out of order, the DMAs will not work properly, and even worse, they may not throw any errors to inform you of the problem. Test and verify behavior—as always. Below is an example code snippet from our lab to illustrate the order of when to call the DMA methods.

```
@kernel
def run(self):
    # Record sbc pulse schedule to dma engine
    with self.core_dma.record(self.sbc_dma_handle):
        self.phaser_sbc_sched_multi_mode(0, 1, self.rsb_aom_detunings,
            self.sbc_pulse_schedule_mu)
    delay(1 * s)
    # Record ramp up for MS pulse
    with self.core_dma.record(self.ramp_up_dma_handle):
        self.phaser_pulse_ramp(0, [3, 2], self.ramp_up_pa, [self.red_phase,
            self.blue_phase], self.win_t_step_mu, False)
    # Record ramp down for MS pulse
    with self.core_dma.record(self.ramp_down_dma_handle):
        self.phaser_pulse_ramp(0, [3, 2], self.ramp_down_pa, [self.red_phase,
            self.blue_phase], self.win_t_step_mu, True)

    # Get dma playback handles to be used later
    sbc_pulse_handle = self.core_dma.get_handle(self.sbc_dma_handle)
    ramp_up_handle = self.core_dma.get_handle(self.ramp_up_dma_handle)
    ramp_down_handle = self.core_dma.get_handle(self.ramp_down_dma_handle)
```

Recording the timestamp of input events. Accessing timestamps is not our usual mode

of operation, so I will walk through how to do this using the PMT as an example. The PMT outputs a pulsed signal for each detection event. These pulses come to ARTIQ through a TTL channel in input mode. We then program ARTIQ to monitor the input events over a specific duration. The code for counting the number of input events (not timestamps) may look something like this:

```
# Turn on detection beam

self.detection_aom.sw.on()

# Record input signals from pmt
gate_end = self.pmt_ttl.gate_rising(det_time * us)

# Turn off detection beam

self.detection_aom.sw.off()

# Count number of pmt events that were recorded

pmt_counts = self.pmt_ttl.count(gate_end) # Blocking

# Delay to give CPU time to count events

delay(self.delay_counts * us)
```

For detection on a PMT, it can be advantageous to know *when* a photon was detected. There are at least two reasons. First, there is a probability of a bright ion going dark during the detection period through off-resonant coupling to the dark state and vice versa, though the probabilities are asymmetric [83]. The additional information of when a photon arrives at the PMT can help discriminate whether an ion started in the bright or dark state at the beginning of the detection period even if the ion switched states later on. This turns out to be quite powerful and can decrease detection infidelity by an order of magnitude or more as discussed previously in Sec. 3.4. A second reason time-of-arrival is useful is that photon emission can be correlated with the motion of the ion. In particular, we can detect micromotion by comparing photon time-of-arrival and the phase of the trap RF frequency. We have used this in the past to help reduce micromotion (see Sec. 2.1.2).

To retrieve the machine unit precise (nanosecond) timestamp of an input event, loop the

function `TTLInOut.timestamp_mu(...)` and save the output. Once the output is -1 , there are no more events as shown in the code example below.

```
# Turn on detection beam

self.detection_aom.sw.on()

# Record input signals from pmt
gate_end = self.pmt_ttl.gate_rising(det_time * us)

# Turn off detection beam

self.detection_aom.sw.off()

# Extract all pmt event timestamps
for i in range(self.max_events):
    self.timestamps[i] = self.ttl0.timestamp_mu(gate_end)
    if self.timestamps[i] < 0:
        break

# Delay to allow for timestamp extraction
delay(500 * us)
```

Notice how it cannot be known beforehand how many timestamps will be saved, but a predefined array will need to be initialized so there is something for the timestamps to be saved to. The resolution to this is rather simple. Initialize the empty array to be large compared to the expected number of events. In this example, that would mean setting the size of the timestamps array (`self.max_events`) to be much larger than the average number of expected photons. This has worked well for us in the past, it is fairly easy to predict an overestimate of the number of events, and it is computationally cheap to inflate an array with the capacity of a hundred or even few thousand elements.

Although it may be impossible to go backward in time in real life, we can easily do so in ARTIQ using negative time shifts. We only recently discovered that delays, or shifts of the timeline, can be negative thanks to talks with Leon Riesebois and Marissa D’Onofrio from the Duke Quantum Center. This is particularly helpful when programming the 355 nm




Figure 3.17: Timeline schematics illustrating an example case of negative delays. The left column is the frame of reference of the Kasli as programmed by the user (RTIO timeline). The right column is ion's frame of reference (ion timeline). If a latency between previously programmed events and, for example, a TTL pulse, there will be a timing error from the ion's perspective. The TTL pulse will arrive late as shown in the top right. If the events are scheduled with negative delays accounting for the latency as shown in the bottom left, then events at the ion's timeline can occur seamlessly as shown in the bottom right.

AOMs as they need to operate simultaneously or to remove system latencies as shown in Fig. 3.17. We have previously considered using the ARTIQ `parallel:` and `sequential:` block commands, but that does not lend itself to modular control. The code for each arm would need to be inserted separately into the blocks and is quite cumbersome to other events on the timeline. In fact, we could not produce a modular definition of 355 nm AOM operations in this paradigm. Modular control meaning any event can be placed in which ever order we wish without consequence or influence on other events is the desired behavior, and it is possible when using negative delays. For example, below is a slightly modified version of how we pulse both 355 nm AOMs simultaneously. Ignore the function arguments.

```
@kernel
```

```
def raman_pulse(self, ch_ind, osc_inds, phaser_amps,
                phaser_phases, pulse_time_mu, leave_arms_on=False):

    "Ignoring docstring for this example"
```

```

# Turn on outer AOM for pulse_time_mu nanoseconds

self.raman_pulse_outer_arm(ch_ind, osc_inds, phaser_amps,
                           phaser_phases, pulse_time_mu, leave_arms_on)

# raman_pulse_outer_arm takes
# pulse_time_mu + self.phaser_sample_window_mu nanoseconds
# but light is only pulsed for pulse_time_mu.

# Go backward in time to program the inner arm AOM
delay_mu(- pulse_time_mu - self.phaser_sample_window_mu)

# Drive inner arm AOM for pulse_time_mu nanoseconds

self.raman_pulse_inner_arm(pulse_time_mu, False, leave_arms_on)

delay_mu(self.phaser_sample_window_mu)

```

In the above code, the outer arm AOM drive is first scheduled. The time cursor is then moved backward in time using a negative delay to when the moment the outer arm AOM drive was scheduled to start. We then program the inner arm to pulse. The end behavior is a simultaneous drive of both AOMs. In this way, the control of both AOMs is containerized by the subartiq-lib function `raman_pulse()` and acts like a single event. This function can be called totally independent of any other commands—exactly how a function should work. This would not be as neat with a parallel and sequential block coding design when considering additional commands.

The last advanced use case described here is pulse shaping. Pulse shaping, also known as amplitude modulation, is a “quantum control” technique that can be used to limit off-resonant coupling, control phase space trajectories, and make MS gates more robust to some noise sources [123, 124, 64]. The idea is to vary the RF power driving an AOM which then varies the laser beam intensity such that the beam intensity has a shaped profile in time—instead of the typical off-on-off square profile. Given a beam intensity for each arm of the 355 nm I_1 and I_2 , the Rabi frequency Ω of the overall drive is proportional to the square root of the intensity from each arm $\Omega \propto \sqrt{I_1 I_2}$. This implies that if one beam



Figure 3.18: Screenshot of a photodiode signal of the 355 nm inner arm beam at low power. The pulse sequence is labeled on the image and consists of square and Tukey-like pulses. The sequence mimics a Trotter sequence of interweaved single and two-qubit gates. The light pulses are shaped (or amplitude modulated) by the Phaser as outlined in the text. The photodiode averages over the individual light pulses; oscillations and the vertical asymmetry are artifacts.

is shaped, the overall interaction is effectively shaped. We shape using the outer AOM which is controlled by the Phaser. The Phaser was originally designed with pulse shaping hardware [135] but getting those features into a stable version of ARTIQ has been slow. See the MIQRO project for updates. Instead, we have explored “hacking” the Phaser’s fine scale, and programmable, power control to get shaped pulses—despite this not being the intended method for pulse shaping. There are two main considerations for this approach. First, the Phaser’s output power is not in physical units or remotely related to any unit that we can tell. We manually calibrate the relation between the “amplitude” units the Phaser uses and physical RF power units such as milliwatts and dBm. These results are saved under the `subartiq_lib` `calibrations.py` script, so we can convert between units at will. The calibration measurements have shown to be stable over several months. Second, the Phaser can make

changes to its output ever 40 ns (25 MS/s). Therefore, when computing a pulse shape, the changes in the shape need to be defined at time points which are integer multiples of 40 ns.

We have successfully sent shaped RF pulses to the 355 nm outer arm AOM as shown in Fig. 3.18. Initial results showed pulse shaping improved MS gates. Future projects will more thoroughly quantify the benefits of this technique for an MS gate in our lab.

In summary, ARTIQ software and Sinara hardware are powerful and capable. With careful design, effort, and testing, most any control goal is achievable.

3.6 Typical Experiment Workflows

In this chapter, I have detailed out the various tools that comprise the experimental apparatus. Now, the discussion will shift to a bigger picture of operations, troubleshooting, and useful procedures.

A typical daily procedure [97, 106], across almost any experimental study is the following:

- Warm up lasers: 369 nm, 399 nm, 935 nm, and 355 nm (30 - 60 minutes; Fig. 3.7).
- Lock the 369 nm and 935 nm laser frequencies. Bring the 399 nm within an acceptable frequency range (± 10 MHz).
- Turn on the ^{171}Yb oven ($\sim 1 - 3$ min at 0.57 V and 2.59 A). Flash the 355 nm to assist photoionization. Turn off the oven once an ion is on the camera image.
- Check 369 operations such as Doppler cooling, detection statistics, optical pumping rate.
- Lock the 355 nm beat note. Check the 355 nm alignment with Rabi oscillations. Touch up if needed. (Can be stable without touch up for months at a time.)
- Fine tune ion position to minimize micromotion (see Sec. 2.1.2 for methods).
- Measure mode frequencies.

- Optimize ground state cooling.

At this point, the ion and laser systems should be ready for most any experiment or loading more ions. We record the results of these calibration experiments in a table in the electronic lab notebook (OneNote). Observations such as ion position on the camera, the trap RF reflection voltage, 369 nm beam powers near the trap, etc. are also recorded. Monitoring these parameters confirms the reliability of the apparatus and ensures repeatability across years of operation.

3.7 Troubleshooting

Much of the work in experimental physics is spent troubleshooting the apparatus. Troubleshooting and rubbing your chin while looking over data from a diagnostic experiment are a graduate student's most honed skills. What ends up in a published paper is maybe < 0.1% of the data it took to get the system running and optimized. Many others have written on the topic of troubleshooting ion trap experiments [97, 106, 105, 108]; so, I only add my prospective and points that seem to have been missed.

3.7.1 Trapping the First Ion

When trapping an ion for the first time, begin by reading sections of the following theses Rajibul Islam [97], Crystal Senko [106], or Marissa D'Onofrio [105]. Consider careful imaging of the trap geometry using the imaging camera and an incident laser beam. Recording the laser beam's translation stage micrometer values can be very helpful for positioning the beam close to the ion and verifying expected geometries. Lastly, in our experience of trapping for the first time after a vacuum bake, the primary issue *every time* has been a DC or RF issue: incorrect wiring, melted electrodes [55], etc.

3.7.2 General ARTIQ Advice

Whatever you program commands in ARTIQ always *directly* verify the behavior. For example, if you are driving an AOM, plug the signal into a scope first to check the pulse is the width you programmed. Measure the signal power with an RF meter. Measure the timing of the light pulses with photodiodes and an oscilloscope and so forth.

If you are experiencing errors in ARTIQ and just can't figure it out, read the docs [111]. They really can be helpful, so take time to read relevant sections and even explore. It may take many reads of the same sections for things to sink in. If the answer is not there, ask your question on the M-Labs forum [127]. I have always had my posts answered.

3.7.3 Latencies and Odd Early Time Behavior

For a time, Rabi oscillations would measure an initial lag. The first $1.4\ \mu\text{s}$ was flat. A seemingly unrelated problem was that camera images of dark ions looked like bright ions. (See Sec. 3.4.3 for details on the camera.) These bugs were all rooted in one issue: frame of reference. The ion's frame of reference is *the* ultimate frame of reference. All operations need to have their latencies accounted for such that programmed events originating from different sources happen at the ion at the expected times. Laser pulse timings are easily checked with a pair of photodiodes and an oscilloscope. Account for latencies using negative delays as discussed in Sec. 3.5.6.

3.7.4 Motional Coherence

If an ion shows poor motional coherence, for example a quickly decaying blue sideband oscillation, consider adjusting the trap RF frequency. We have found a few kHz change can dramatically improve, or reduce, the motional coherence. We currently suspect this poor motional coherence happens when the RF frequency Ω_{RF} is a harmonic of the secular harmonic frequencies $\Omega_{\text{RF}} = \sum_k^N n_k \omega_k$ where each n_k is an integer and ω_k is the k th mode frequency. This can be an unstable trapping condition [136] (also see Ref. [5] Sec. 4.1.5).

Unfortunately, it is quite easy to satisfy this condition when there are many modes present. Maintaining cold ion temperatures and a high trap frequency Ω_{RF} will reduce this effect [5]. For our experiment, changing the RF frequency to improve the motional coherence is not always repeatable. Sometimes changing the RF frequency and coming back to the original value gives the best motional coherence. There are likely thermodynamic effects confounding our troubleshooting attempt such as causing the trap geometry to expand and contract at different RF frequencies and powers.

3.7.5 Lasers

740 nm Doubling Efficiency As mentioned previously, the doubling cavity in the 369 beam line losses efficiency over time. In the last two years, we have found that by carefully cleaning the front face of the doubling crystal (with a single wipe) the doubling efficiency immediately improves and gives about another year worth of light.

Stray Optical Pumping Light A few tens of nanowatts of leaking optical pumping light have disrupted our experiments in the past. This is easily checked with a long Ramsey experiment (Crystal Senko [106] suggestions another method). The major source of this issue was poor beam quality of the 369.5 nm light at the optical pumping section of the optics table. An iris before the optical pumping section resolved the issue. It was more effective than AOM or fiber realignment. It is worth rechecking the qubit coherence time every few months to confirm light leakage is minimal.

355 nm Stray Light If there is a significant amount of stray light coming down the beam path, it has often been due to circumstances internal to the 355 nm laser housing. Adjusting the 355 nm to a new spot using the 355 nm control program has resolved this in the past.

355 nm Power If the 355 nm power is dropping below 8 W in lightloop mode, follow the instructions for reoptimization in the lab folders. Find the notes on previous times this has

been done. It takes a few hours since some of the adjustments are temperatures, but by closely following the manual it has always improved.

3.7.6 RF Breakdown and DC Shorts

Not all is lost... maybe. In 2023, I fried our rod trap while raising the RF voltage. We wanted higher secular frequencies but hadn't realized we were near an arching/RF breakdown event. We detected the issue from four sources. First, the reflection off the helical resonator was not acting like a resonator—there was no resonant frequency. Second, the power supplies for the DC electrodes were suddenly showing nonzero current values. Third, the pressure in the trap spiked. Fourth, the trap was *glowing* orange. Why this happened is almost impossible to know for sure. However, we speculate that the long term use of the oven may have coated the isolating sections between electrodes with Yb metal making an breakdown event much more likely. The glowing was at the junction of insulating macor and an electrode. It is my opinion that the thermodynamics of the RF voltage heating the electrodes and the macor insulator having poor thermal conductivity are important somehow as well, but this is nearly impossible to test with our current apparatus.

There are some anecdotal stories and even a published article Ref. [137] out from Sandia about RF breakdown, or arching, events between electrodes. Their work is with surface traps which have much closer geometries than our rod trap, but we seem to be subject to some of the same issues. To avoid RF breakdown in the first place, increase the RF power slowly. Counterintuitively, breakdown events may not happen immediately [137]. Second, if an event happens, drop the RF voltage immediately. Monitor the RF reflection as well. We noticed the RF reflection blew up when we had an RF breakdown event, but after lowering the RF voltage, the reflection behaved like a resonator somewhat (with a shifted resonance frequency but poor Q factor). In our case, the DC power supplies started driving current as well implying there was a short, which we assumed was related to the high voltage RF and orange glowing trap. This is pretty much a break-vacuum-and-repair situation. However,

there is a small chance that a high current could burn out the short. We attempted this by increasing the current from the DC power supplies up to as high as 12.0 V and 24.5 A, and after some intense moments, the current from the supplies suddenly dropped to zero suggesting we successfully burned out the short. We have had to operate with less RF voltage in the trap than before, but a working trap is better than no trap.

3.7.7 Miscellaneous

A mix of collected advice for the experimentalist, though if applied out of context may not be that advisable:

- When it comes to machining, you can only have it two of three ways: fast, cheap, or accurate.
- Record the absolute state of the experiment before making adjustments.
- Do not adjust something you know is working to correct for something that's broken. Resolve the issue with the broken thing.
- Variance is as important if not more important than mean: show errorbars.
- If your numerics are off by ~ 2 , check for factors of 2 or π . If your numerics are off by ~ 10 check for factors of 2π

Chapter 4

Optimized Pulse Sideband cooling

The cooling of mechanical oscillators to near their ground motional state is of fundamental importance to fields as varied as atomic clocks [49, 138, 139], quantum computation and simulation [22, 140, 141, 7], quantum sensing and transduction [142, 143, 144], and even gravitational wave detection [145]. Particularly for atom-based platforms, Doppler laser cooling provides a fast and straightforward method for reducing the kinetic energy of the system by orders of magnitude to reach the quantum regime [6, 146]. Even so, recoil effects during photon emission typically prevent Doppler-cooled systems from achieving their absolute motional ground state, requiring the implementation of sub-Doppler cooling methods [147, 148, 102, 113].

For trapped ion experiments, resolved sideband cooling (SBC) is the most popular sub-Doppler cooling technique used to prepare systems near their motional ground state [103, 149, 150, 151]. Its widespread use stems largely from its applicability to most trapped-ion setups, since its effectiveness does not rely on using a specific ion species or trap geometry [5, 152]. In practice, SBC allows trapped ions to be initialized in a nearly-pure ground state of motion, with a typical average harmonic occupation $\bar{n} \lesssim 0.05$ [103]. However, SBC is often the longest time component in an experimental cycle by a significant factor [22, 23]. The situation is only more complicated when many motional modes need to be cooled. Although individual addressing can facilitate some speedups in long ion chains [152], to date no general method is known for determining the optimal SBC protocol.

This chapter presents a framework for calculating the optimal sequence of SBC pulses for near-ground-state cooling. The optimal cooling strategy is applicable to any trapped ion experiment using pulsed SBC and flexible enough to incorporate decoherence effects or heating models if desired. The optimized SBC sequences are benchmarked using a trapped

$^{171}\text{Yb}^+$ ion, finding close experimental agreement with theory predictions as well as significant improvements compared with traditional cooling protocols. The chapter begins with a review of the standard theory of pulsed resolved SBC. Next, the pulsed SBC problem is cast into a matrix formalism that allows for efficient numerical optimization of SBC pulse sequences. Then, optimizations under different experimental constraints are compared. Finally, we discuss implications of the graph framework, potential add-ons, and future directions.

4.1 Resolved Sideband Cooling Theory

When a trapped ion of mass m is confined to a 1D harmonic potential of frequency ω , resolved SBC allows for sub-Doppler cooling of the ion temperature. Prior to the onset of SBC, we assume that the ion has been Doppler cooled using a transition of linewidth Γ to the Doppler cooling limit [86, 153]

$$\bar{n}_i \approx \frac{\Gamma}{2\omega}. \quad (4.1)$$

Following Doppler cooling, the probability of finding the ion in the n^{th} harmonic oscillator level is well-described by the thermal distribution

$$p_{\text{th}}(n) = \frac{\bar{n}^n}{(\bar{n} + 1)^{n+1}} \quad (4.2)$$

which is solely parameterized by the average harmonic state of the ion $\bar{n}$ as derived in Sec. 2.2.2.

SBC protocols may be implemented for both optical and hyperfine qubits; here we begin by focusing on the latter. Typically, far-detuned Raman transitions of wavelength λ and linewidth $\gamma_{\text{rad}} \ll \omega$ are used to manipulate the electronic and motional states of the ion. When the Raman transition frequency is in resonance with the qubit splitting, it drives a “carrier” transition between qubit levels $|\downarrow\rangle$ and $|\uparrow\rangle$ at Rabi frequency Ω , with no change to the motional state. Detuning the Raman frequency by integer multiples of the trap secular frequency ω excites a “sideband” transition, coupling spin flips to a change in motional state

from $|n\rangle$ to $|n'\rangle$, at Rabi rate [8, 5]

$$\Omega_{n,n'} = \Omega e^{-\eta^2/2} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta^{|n-n'|} \mathcal{L}_{n_{<}}^{|n-n'|}(\eta^2), \quad (4.3)$$

where $n_{<}$ ($n_{>}$) is the lesser (greater) of n and n' ,

$$\mathcal{L}_n^{(\alpha)}(X) = \sum_{i=0}^n (-1)^i \binom{n+\alpha}{n-i} \frac{X^i}{i!} \quad (4.4)$$

is the generalized Laguerre polynomial, and

$$\eta \equiv \Delta k x_0 = 2 \sin(\theta/2) \frac{2\pi}{\lambda} \sqrt{\frac{\hbar}{2m\omega}}. \quad (4.5)$$

is the Lamb-Dicke parameter for counter-propagating Raman beams which intersect at an angle θ . In this study, we will refer to an $n - n' = 1$ transition as a first-order red sideband (RSB) transition and an $n - n' = -1$ transition as a first-order blue sideband (BSB) transition.

SBC of hyperfine qubits is typically characterized by a sequence of discrete RSB pulses interleaved with optical pumping (Sec. 2.4.3). A traditional pulsed SBC protocol (which we will call the “classic” protocol) executes as follows [103, 150, 98]. After Doppler cooling to an average harmonic occupation $\bar{n}_i$, and optical pumping to the qubit state $|\downarrow\rangle$, an initial motional level $n_i \gg \bar{n}_i$ is selected as the entry point for SBC. A first-order RSB π -pulse is then applied for $t = \pi/\Omega_{n_i, n_i-1}$ followed by fast optical pumping, to drive the transition $|\downarrow, n_i\rangle \rightarrow |\downarrow, n_i-1\rangle$. Then another iteration is performed using $t = \pi/\Omega_{n_i-1, n_i-2}$, and so on, until the sequence concludes with a final $t = \pi/\Omega_{1,0}$ pulse. In principle, this protocol sweeps the fraction of population for which $n \leq n_i$ into the motional ground state.

By starting at larger n_i and iterating for more pulses, the classic SBC protocol can theoretically reach the SBC limit of $\bar{n}_{\min} \approx (\gamma_{\text{rad}}/2\omega)^2 \ll 1$ [5, 154, 8, 155] where γ_{rad} is the effective decay rate of the optical pumping step. In terms of the motional ground state

$n = 0$, the cooling limit is fundamentally set by the off-resonant coupling during the optical pumping. The minimum probability of reaching $n = 0$ being $p_0 \approx 1 - (\gamma_{\text{rad}}/2\omega)^2$.

In practice, the achievable final $\bar{n}$ may be limited by effects such as imperfect RSB π -pulses, motional heating, and nearly-infinite RSB π -times (Sec. 4.2.1); this is indeed the case for several trapped-ion experiments [156, 104, 54]. Nevertheless, post-SBC temperatures of $\bar{n} \lesssim 0.05$ are routinely achieved with the classic method [103, 150], particularly when the initial state before SBC is in the “low η - $\bar{n}_i$ regime”: $\eta \ll 1$ and $\bar{n}_i \lesssim 10$.

4.1.1 Alternative Cooling Techniques

Although this chapter will focus on the most common ground state cooling techniques, pulsed sideband cooling, we briefly discuss additional cooling strategies for completeness.

For optical qubits, continuous SBC is an alternative protocol for achieving near-ground state cooling [157, 100]. In this approach, a RSB is driven continuously on a narrow optical transition while optical pumping is accomplished by spontaneous emission from the excited state. Given the slow decay rate of narrow transitions, spontaneous emission may be enhanced by temporarily coupling the excited state to a dipole-allowed transition. In $^{40}\text{Ca}^+$, for instance, coupling the quadrupole $D_{5/2}$ qubit level to the dipole-allowed $P_{3/2}$ state can lead to cooling rates of $\dot{\bar{n}} = 5 \text{ ms}^{-1}$ when strongly saturating the RSB transition [157]. As we will show in Sec. 4.2, this rate is comparable to the pulsed SBC rate in hyperfine qubits driven by a carrier Rabi frequency $\Omega \approx 2\pi \times 10 \text{ kHz}$. For our experiments, we set $\Omega = 2\pi \times 65 \text{ kHz}$, leading to an initial cooling rate of $\dot{\bar{n}} \approx 30 \text{ ms}^{-1}$.

Continuous SBC has been well-described via detailed theoretical models [149, 153, 100] and validated in experiments [157, 100]. For a given optical pumping rate, the optimum RSB parameters for achieving the lowest final $\bar{n}$ may be estimated from the full set of atomic rate equations [149], or determined experimentally by scanning over different values of RSB power and frequency [158] or optical pumping rates. In contrast, the discreteness of pulsed SBC protocols prevents a similar rate-equation type analysis while greatly expanding the

parameter space of possible cooling sequences. For these reasons, finding a pulsed SBC model that allows for efficient determination of optimal sequences has remained elusive to date; we seek to address this open question in the remainder of this chapter.

4.2 Optimized Pulsed SBC Protocols

For hyperfine qubits, the intuitive ‘classic’ protocol introduced previously is not the most *efficient* pulsed SBC method for reducing an ion’s motional energy. Given a chosen $\bar{n}_i$, which sets the number of pulses, there are no adjustable parameters that may be used to optimize the cooling rate per pulse or per unit time. When starting from small Doppler-cooled $\bar{n}_i$, only a few cycles are needed and the deviation from optimal is small; when $\bar{n}_i$ is large ($\gtrsim 10$), the deviation from optimal widens considerably. If $\bar{n}_i$ is large enough, the classic method will fail to prepare ions in the ground motional state as mentioned previously.

Here, we introduce two globally-optimized pulsed SBC protocols: a single-parameter protocol called the “fixed” method, and a full-parameter protocol called “optimal” method. For a given number of pulses, the optimal method provides the lowest possible $\bar{n}$ after first-order SBC. When $\bar{n}_i$ is large, we show how these protocols can be extended to higher-order SBC to avoid the limitations of first-order cooling. To compute these optimized SBC protocols, we must first numerically simulate the complicated interplay between each π -pulse and its affect on the *entire* harmonic oscillator population $p(n)$. Below, we develop a graph-theoretic description of pulsed SBC to accomplish this task and provide a framework for fast optimization of pulse sequences.

4.2.1 Graph-Theoretic Description of Pulsed Sideband Cooling

We embed SBC into a graph $G = (V, E)$ with a set of vertices V and edges E . The vertices V represent a truncated set of the harmonic states $n = [0, n_{\max}]$ where $n_{\max} \gg \bar{n}_i$ is well satisfied. Each vertex is weighted by the probability corresponding to its harmonic state $V = \{p(0), p(1), \dots, p(n_{\max})\}$, as shown in Fig. 4.1. Each vertex has an undirected edge




Figure 4.1: Graph G representing first-order SBC. The set of vertices V is represented by circles and weighted by the current harmonic probability distribution $p(n)$. The set of edges E is represented by lines: loops weighted by $a_n(t)$ and directed edges weighted by $b_n(t)$.

loop weighted by the probability of not cooling: $a_n(t) = \cos^2(\Omega_{n,n-1}t/2)$ in the case of first-order cooling as shown in Fig. 4.1. The probability of cooling $b_n(t) = \sin^2(\Omega_{n,n-1}t/2)$ weights a directed edge from the n to $n-1$ vertices. For m th-order cooling, the directed edges would connect to their m th leftmost neighbor with the associated Rabi frequency $\Omega_{n,n-m}$. All edges exiting a vertex must sum to 1 to conserve the harmonic state distribution probability.

To model one SBC pulse of time t_0 , all vertex weights take one traversal of their respective edges resulting a new set of vertex weights: $V_n^{(1)} = a_n(t_0)V_n^{(0)} + b_{n+1}(t_0)V_{n+1}^{(0)}$. To model N SBC pulses, the graph is traversed N times. In general, each traversal may have its own associated pulse time $\{t_0, t_1, \dots, t_{N-1}\}$.

We numerically represent the graph and SBC process as a matrix equation. The initial vertex values map to the vector $\vec{p}_{\text{th}} = \{p_{\text{th}}(0), \dots, p_{\text{th}}(n_{\text{max}})\}$, where $p_{\text{th}}(n)$ is the initial thermal distribution following Doppler cooling (Eq. (4.2)). One traversal of the graph maps to the upper triangular matrix

$$W(t) = \begin{pmatrix} 1 & b_1(t) & 0 & \dots \\ 0 & a_1(t) & b_2(t) & \dots \\ 0 & 0 & a_2(t) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (4.6)$$

which is shown graphically in Fig 4.2(a) for $t = 1.016 \times 2\pi/\Omega$. $W(t)$ acting on $\vec{p}_{\text{th}}$ results in an updated probability vector $\vec{p} = \{p(0), \dots, p(n_{\text{max}})\}$




Figure 4.2: The first 30×30 matrix elements of the weight matrix (Eq. (4.6)) are shown graphically for (a) a single pulse and (b) 25 repetitions of the pulse applied in (a).

$$\begin{pmatrix} p(0) \\ p(1) \\ p(2) \\ \vdots \end{pmatrix} = \begin{pmatrix} 1 & b_1(t) & 0 & \dots \\ 0 & a_1(t) & b_2(t) & \dots \\ 0 & 0 & a_2(t) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} p_{\text{th}}(0) \\ p_{\text{th}}(1) \\ p_{\text{th}}(2) \\ \vdots \end{pmatrix}. \quad (4.7)$$

To encode the effects of multiple SBC pulses, all individual pulse matrices $W(t)$ are multiplied together: $W(t_{N-1}) \dots W(t_1)W(t_0)$. In the simplest case, when all pulses are of the same duration t_0 , the SBC interaction is encoded as a matrix power of $W(t_0)$. For example, the final harmonic level occupation after 25 identical pulses can be calculated as $\vec{p} = W^{25}(t_0)\vec{p}_{\text{th}}$, with the low- n matrix elements of W^{25} shown in Fig. 4.2(b).

Fixed protocol

Optimized pulse sequences may be efficiently computed within the graph-theoretic framework introduced above. To begin, we consider a single-parameter optimization that we call the “fixed” protocol. Each of the SBC pulses is chosen to have the same duration $T_{\text{fixed}} =$

$\{t_0, \dots, t_0\}$, similar to SBC schemes implemented in some trapped-ion studies [159, 160, 156]. Here we explicitly seek to minimize the function

$$\bar{n}(t_0) = \sum_{n=0}^{n_{\max}} n \left[W^N(t_0) \vec{p}_{\text{th}} \right]_n \quad (4.8)$$

to find the time t_0 which yields the lowest possible $\bar{n}$ given N identical SBC pulses.

The optimal pulse time for the fixed method can be computed quickly since there is only one parameter to optimize for any number of pulses N . The most costly step in minimizing Eq. (4.8) is the calculation of $[W^N(t_0) \vec{p}_{\text{th}}]_n$ for different t_0 . However, standard numerical packages, such as python's NumPy module [161], can exponentially reduce the number of matrix multiplications needed when computing a power of a matrix through binary decomposition. Assuming $N > 3$, a binary decomposition recursively squares the matrix, exponentially increasing the matrix power: 2, 4, 8, and so on. The implementation is adapted to allow for arbitrary matrix powers, with a computation time scaling with N as $\mathcal{O}(\log_2(N))$ and with system size $n_{\max}$ as $\mathcal{O}(n_{\max}^3)$.

Optimal Protocol

We now consider the optimal protocol, which is a full-parameter optimization where each pulse time is treated as an independent variable. Given a set of experimental parameters and restricting—for now—to first-order RSB pulses, the remaining degrees of freedom are the durations of each SBC pulse. The optimal protocol searches the full available parameter space of N distinct pulse times, yielding the lowest possible $\bar{n}$ for any given value of η , $\bar{n}_i$, Ω , and N .

The optimal protocol, using first-order RSBs, executes as follows. First, the initial harmonic populations $\vec{p}_{\text{th}}$ and Rabi frequencies $\Omega_{n,n-1}$ are calculated over a truncated range of harmonic states $[0, n_{\max}]$ ($n_{\max} \gg \bar{n}_i$), based on the experimental parameters η , $\bar{n}_i$, and Ω .

Next, a gradient descent algorithm is applied to minimize the equation

$$\bar{n}(t_0, t_1, \dots, t_{N-1}) = \sum_{n=0}^{n_{\max}} n [W(t_{N-1}) \dots W(t_1) W(t_0) \vec{p}_{\text{th}}]_n \quad (4.9)$$

to find the pulse schedule $T_{\text{optimal}} = \{t_0, t_1, \dots, t_{N-1}\}$ that gives the lowest average harmonic occupation $\bar{n}(t_0, t_1, \dots, t_{N-1})$ following N SBC pulses.

Since each pulse time in the pulse schedule T_{optimal} is an independent variable, computing the optimal T_{optimal} scales exponentially with the number of pulses. For large $n_{\max}$ or N , this can cause calculations to exceed readily available computational resources. However, we find that careful bounding of the gradient descent minimization can help reduce computation times. For example, using a standard laptop we observe that a 50-pulse SBC optimization takes less than 90 seconds to compute, which is a factor of two faster than for the unbounded case.

The predicted performance of the optimal, fixed, and classic protocols is compared in Fig. 4.3. Simulations are performed using the parameters $\bar{n}_i = 15.36$ and $\eta = 0.18$, which are similar to those of our experimental system. For fewer than ~ 50 SBC pulses, the classic method not only takes the longest absolute time to implement (Fig. 4.3(a)), but also yields the highest final $\bar{n}$ (Fig. 4.3(b)). In comparison, the fixed (solid black) and optimal (dashed gray) methods perform nearly identically, both in overall cooling time and final ion temperature. For larger $\bar{n}_i$, the classic method drifts further away from optimal, while the fixed method retains its near-optimal behavior.

Multiorder Optimization

When outside of the low $\eta\text{-}\bar{n}_i$ regime, the trapping of harmonic population in high- n states can limit first-order RSB cooling [160, 156, 104]. As shown in Fig. 4.3(c), the first-order RSB Rabi frequency approaches zero for specific high- n harmonic levels (approximately $n = 112$ for our chosen parameters). As a consequence, any initial population $n \gtrsim 112$ will be trapped




Figure 4.3: The classic, fixed, optimal, and multiorder protocols are compared for an initial temperature of $\bar{n}_i = 15.36$, and $\eta = 0.18$ (see text for definitions). (a) The total sideband cooling time (excluding optical pumping), and (b) the cooled $\bar{n}$ as a function of the number of SBC pulses. (c) Scaled frequencies for the first-order (solid) and second-order (dash-dot) RSB showing the near-zero frequency of the first-order RSB at $n = 112$. (d) Initial thermal distribution (solid light red) and distributions after 50 pulses of first-order fixed (solid black) and multiorder fixed (dash-dotted purple).

in these high- n states, even while the remaining population $n \lesssim 112$ is swept towards the ground state.

This population trapping effect is visible in Fig. 4.3(d), which shows the harmonic population distribution following 50 first-order SBC pulses. A significant population near $n = 112$ remains uncooled, contributing approximately 0.3 motional quanta to the final value of $\bar{n}$: an order of magnitude higher than the SBC cooling limit and large compared to what is considered near-ground-state cooling. This effect also explains why the three first-order methods in Fig. 4.3(a)-(b) begin to converge at large numbers of pulses: the trapped population contributions to $\bar{n}$ dominate at colder temperatures.

To avoid population trapping at high- n , higher-order RSB pulses can be incorporated into the SBC protocol. We refer to this scheme as “multiorder” cooling. Particularly in experimental regimes where η or $\bar{n}_i$ are large, trapped populations may be so significant that multiorder cooling is *required* to achieve near-ground-state temperatures [160, 156, 104]. This is because the harmonic levels with near-zero RSB Rabi frequencies shift to smaller n as η increases, and because larger fractions of the initial population will be trapped at high- n as $\bar{n}_i$ increases.

Multiorder cooling circumvents population trapping since, for different RSB orders, the Rabi frequencies approach zero at different values of n . This is illustrated in Fig. 4.3(c), where it can be seen that higher RSB orders exhibit their first zeros at higher values of n . This allows for multiorder pulse sequences which first move population from high- to intermediate- n , then employ first-order pulses to reach the ground state.

The graph-theoretic framework we introduced in Sec. 4.2.1 can easily incorporate higher-order pulses. For an m^{th} order pulse of time t , the probability of not cooling is $a_n(t) = \cos^2(\Omega_{n,n-m}t/2)$ and is mapped to the diagonal of the weight matrix $W(t)$. Likewise, the probability that the m^{th} order pulse takes $|n\rangle \rightarrow |n-m\rangle$ is $b_n(t) = \sin^2(\Omega_{n,n-m}t/2)$ and is mapped to the m^{th} upper diagonal of $W(t)$. Both the fixed and optimal protocols may then be calculated for multiorder cooling once the $W(t)$ matrices are constructed.

We simulate and optimize a multiorder fixed protocol with N_3 third-order pulses, N_2 second-order pulses, and N_1 first-order pulses fixing the total number of SBC pulses $N = N_1 + N_2 + N_3$ and allowing the pulse time to vary per order

$$\bar{n}(t_1, t_2, t_3) = \sum_{n=0}^{n_{\max}} n \left[W^{N_1}(t_1) W^{N_2}(t_2) W^{N_3}(t_3) \vec{p}_{\text{th}} \right]_n. \quad (4.10)$$

N_1 , N_2 , and N_3 were selected by brute force optimization of a block sequence (detailed in the next paragraph). Figure 4.3(b) shows multiorder cooling (dash-dotted purple) working significantly faster than the optimal first-order method, cooling from $\bar{n}_i = 15.36$ to a final $\bar{n} = 0.06$ after only 50 pulses. In addition, the multiorder protocol avoids the high- n population trapping present in the first-order sequences. This can be seen in Fig. 4.3(d), where population is much more efficiently transferred from high- n to low- n when multiorder pulses are used.

Multiorder cooling introduces further optimization and experimental challenges. For an N pulse SBC protocol that includes k_m pulses of order m , there are a factorial number of permutations ($N! / \prod k_m!$) in which the pulse orders may be sequenced, and an exponential number of $\{k_m\}$ choices which satisfy $\sum k_m = N$. For small numbers of pulses ($N \lesssim 20$), we used a brute force computation to conclude that a “block” sequence is best: all k_m pulses of the same order m stay together in a “block,” and higher-order m blocks are applied before lower orders. Under this restriction, the number of possible sequences becomes polynomial in the number of applied orders m , scaling as $\mathcal{O}(N^{m-1})$.

In practice, applying pulses with arbitrarily high orders is not experimentally feasible. Transition linewidths narrow for higher orders, making resonant excitation difficult. In addition, transition rates decrease, making pulse times impractically long (Eq. 4.3). In our experimental demonstration, we reliably address RSB transitions up to 3rd order. If higher RSB orders are needed, but not possible to apply, alternation between lower orders may still remove trapped population [156] at the cost of longer pulse sequences.

4.3 Summary and outlook

Sideband cooling has been a popular and powerful technique for the near ground-state preparation of trapped ions. Yet, historical approaches to SBC can be made more efficient. In this chapter, we have shown how to calculate the optimal pulsed SBC protocol for any experimental setup characterized by a cooling laser geometry and wavelength, an ion wavepacket width (which depends upon the ion mass and trap frequency), and an initial ion temperature (which depends on the trap frequency and atomic linewidth).

Our efficient numeric simulations and optimizations were enabled by expressing pulsed SBC within a graph-theoretic framework. This approach is powerful for optimizing SBC pulse sequences, and is particularly important in regimes with high Doppler-limited initial temperatures $\bar{n}_i$, or extended ion wavepackets (which correspond to a large η). We observe that repeated SBC pulses with a single optimized time perform nearly-identically to fully-optimized pulse sequences, while traditional protocols were the least efficient per pulse and per unit time.

In future work, we anticipate that the graph representation of pulsed SBC may be expanded to include noise models for ion heating (Ch. 6), decoherence, off-resonant couplings, and effects of rf-driven micromotion. Such additions could be smoothly incorporated into the matrix formalism and would allow for further SBC optimization in the face of realistic experimental imperfections. Extending to multiple ions and multiple modes is another natural direction that fits nicely within the matrix representation of pulsed SBC.

Chapter 5

Trapped-ion Thermometry

Estimating ion temperatures and heating rates are essential characterizations in ion trap experiments [162, 5] since they inform the efficacy of cooling protocols and potential sources of noise for quantum operations. Yet, standard methods for measuring the average occupation number $\bar{n}$ near the ground state implicitly assume the motion is well-described by a thermal distribution of harmonic oscillator levels [103, 150]. When this assumption is violated, as is the case for Fock states, coherent states, or states following significant SBC [62, 104, 98], more sophisticated thermometry methods must be employed to accurately characterize ion motional temperatures.

In this chapter, we develop an improved thermometry technique to more accurately measure $\bar{n}$ —particularly following SBC. Our method requires only the experimental hardware already needed for implementing pulsed SBC. We benchmark our new thermometry technique using a trapped $^{171}\text{Yb}^+$ ion finding close experimental agreement with theory predictions as well as significant improvements compared with traditional cooling and thermometry protocols. The chapter is structured as follows. First, we review current thermometry techniques. Next, introduce the new thermometry method. Finally, compare the techniques at the Doppler cooling limit and after SBC. We summarize with concluding remarks.

5.1 Introduction

In the quantum regime, full ion thermometry requires knowledge of the probabilities $p(n)$ for occupying each harmonic level n , so that the average occupation $\bar{n} = \sum np(n)$ may be calculated. Given the impracticality of measuring dozens or hundreds of probabilities $p(n)$ to high accuracy, thermometry techniques must make assumptions about the underlying distribution $p(n)$. The most common one is to assume that $p(n)$ is thermal, meaning the

distribution follows Gibbs statistics (Sec. 2.4.3), in which case $\bar{n}$ may be extracted by taking the ratio of first-order RSB and BSB transition probabilities [103]. However, Sec. 4.2 and Fig. 4.3(d) demonstrate that sideband-cooled ions can have dramatically non-thermal distributions $p(n)$, depending on the cooling protocol, the number of RSB orders, and the number of cooling pulses. Thus, common ion thermometry methods may give widely inaccurate results following extensive sideband cooling, motivating development of a new approach.

In this section, we begin by outlining two common ion thermometry methods, their underlying assumptions, and the reasons they fail to correctly measure ion temperatures following significant sideband cooling. We then introduce a new technique for ion thermometry which has been specifically tailored to reveal ion temperatures after sideband cooling and depends only on the time-averaged value of RSB transitions.

5.2 Existing Methods

5.2.1 The Ratio Method

Nearly all experiments measuring trapped-ion temperatures deep in the quantum regime follow the approach used in Ref. [103], which we call the “ratio” method. The ion is first initialized in the state $|\downarrow\rangle$, and the first-order red and blue sidebands are then driven with the same power for the same time. If the ion motional distribution is thermal, then the ratio of RSB to BSB transition probabilities can be related to the average harmonic level occupation $\bar{n}$ (Sec. 2.4.3):

$$r \equiv \frac{P_{\uparrow}^{\text{RSB}}(t, \Delta)}{P_{\uparrow}^{\text{BSB}}(t, \Delta)} = \frac{\bar{n}}{\bar{n} + 1}. \quad (5.1)$$

This ratio r may be experimentally determined by fitting absorption lineshapes to frequency scans over the red and blue sidebands (as in [103]), or by driving red and blue sidebands on resonance and taking the ratio of the resulting time series.

The ratio method is powerful due to its direct dependence on $\bar{n}$ and experimental ease.

However, the ratio method relies on the assumption of a thermal harmonic distribution which is inherently mismatched to the motional distribution of ions following significant sideband cooling (see Fig. 4.3(d)). As we will show in Sec. 5.5, this assumption can lead to an order-of-magnitude underestimate of the final $\bar{n}$ after only moderate sideband cooling.

5.2.2 The SVD Method

When the underlying motional distribution is known to be non-thermal, alternative thermometry methods may provide a better estimate of $\bar{n}$. One popular method performs a frequency-domain analysis of a BSB Rabi oscillation, using singular value decomposition (SVD) to extract the harmonic level probabilities $p(n)$ [62]. In this method, a BSB oscillation is described as a matrix of transition probabilities $b_n(t_i) = \sin^2(\Omega_{n,n-1}t_i/2)$ acting on the level probability vector $\vec{p}$ to yield the measured fluorescence at each time step t_i . SVD is then used to pseudo-invert the transition probability matrix and isolate the vector of $p(n)$ values (see Sec. 2.4.3 for more detail). This technique has been successfully implemented to measure $\bar{n}$ for both thermal states as well as coherent states [62].

Although SVD is a flexible method for measuring $\bar{n}$ in non-thermal distributions, there are several drawbacks. First, data acquisition can take a long time since long-oscillation time series are necessary to accurately determine as many harmonic state probabilities as possible. This is further compounded by the need to perform many thousands of repetitions to keep quantum projection noise low and avoid potential overfitting during the SVD. Additionally, the output probabilities from SVD have no physical boundary constraints such as $0 \leq p(n) \leq 1$ or $\sum p(n) = 1$. This has been found to produce large errors when applied to distributions with many non-negligible probabilities at high harmonic level n [62], as is the case for the distributions shown in Fig. 4.3(d).

5.3 Modeling Post-SBC Distributions

The primary reason that the ratio and SVD methods fail to accurately estimate $\bar{n}$ following SBC is that they are not well-matched to the motional state distributions shown in Fig. 4.3(d). After SBC, the largest contributions to $\bar{n}$ are often driven by the residual population remaining at large n , which is neglected when using a simple thermal approximation or when focusing on only the low- n populations. Thus, improved modeling of the probability distribution $p(n)$ following SBC is a prerequisite for higher-accuracy estimation of ion temperatures.

To date, the most detailed modelling of post-SBC motional distributions was outlined in [104]. Using simulated multi-order SBC pulses, it was found that the harmonic level populations were well-approximated by a double thermal distribution:

$$p_{\text{double}}(n) = \alpha p_{\text{th}}(n|\bar{n}_l) + (1 - \alpha) p_{\text{th}}(n|\bar{n}_h) \quad (5.2)$$

where $\bar{n}_l$ captures the distribution for low n states, $\bar{n}_h$ captures the distribution for high n states, and the total average occupation is $\bar{n} = \alpha\bar{n}_l + (1 - \alpha)\bar{n}_h$. Our numeric simulations of multiorder SBC in Fig. 4.3(d) likewise demonstrate that the final state populations are well-described by this double-thermal model. In [104], $\bar{n}$ was experimentally determined by first fitting the simulated distribution to extract $\bar{n}_h$, then fitting the experimental data to Eq. (5.2) with $\bar{n}_h$ as a fixed parameter.

Here, we seek to generalize Eq. (5.2) and develop a measurement protocol that avoids dependence on numeric simulations. To begin, we propose direct measurement to find the harmonic level populations $p_{\text{meas}}(n)$ up to $n = k$, where $k > \bar{n}_l$. Using this, we compute the remaining population fraction in all levels $n > k$:

$$p_{\text{rem}}(n > k) = 1 - \sum_{n=0}^k p_{\text{meas}}(n). \quad (5.3)$$

Next, we propose direct measurement of the initial thermal state $\bar{n}_i$ before SBC, which we identify as $\bar{n}_h$ in Eq. (5.2). Once again the quantity $p(n > k)$ is calculated, this time for the initial thermal distribution

$$p_{\text{th}}(n > k) = \sum_{n=k+1}^{\infty} \frac{\bar{n}_i^n}{(\bar{n}_i + 1)^{n+1}} \quad (5.4)$$

The ratio of Eqs. (5.3) and (5.4) estimates the fraction of states remaining in an approximate thermal distribution of average occupation $\bar{n}_i$. The final $\bar{n}$ is then estimated as

$$\bar{n} \approx \sum_{n=0}^k n p_{\text{meas}}(n) + \frac{p_{\text{rem}}(n > k)}{p_{\text{th}}(n > k)} \sum_{n=k+1}^{\infty} n \frac{\bar{n}_i^n}{(\bar{n}_i + 1)^{n+1}}. \quad (5.5)$$

The advantage of Eq. (5.5) is that it leverages the most information available from measurement with no direct dependence on simulation. The only remaining element needed is a robust method to measure the individual probabilities of the low-lying harmonic levels, $p(n \leq k)$. In the following section, we introduce a simple technique that reveals these desired motional state populations.

5.4 Time-averaged Thermometry

We propose a “time-average” measurement protocol which, when combined with Eq. (5.5), provides a high-accuracy estimate of $\bar{n}$ following SBC. This approach is constructed to measure the individual probabilities of the first few harmonic levels. Suppose a trapped ion is initialized in the state $|\downarrow\rangle$. Then, the expected probability of finding the ion in the $|\uparrow\rangle$ state when driven with an m^{th} order RSB is given by:

$$P_{\uparrow, m}^{\text{RSB}}(t) = \sum_{n=0}^{\infty} \frac{1}{2} [1 - e^{-\gamma_n t} \cos(\Omega_{n+m, n} t)] p(n + m). \quad (5.6)$$

where no assumptions have been made about the probability distribution $p(n)$, and decoherence effects at rate $\gamma_n = \gamma_0(n+1)^\alpha$ [62] have been included for generality. We have measured

$\alpha = 0.47$ by preparing different Fock states, driving a blue sideband, and then fitting the dynamics to an exponentially decaying oscillation.

The running time average of Eq. (5.6) is

$$\begin{aligned}\bar{P}_{\uparrow,m}^{\text{RSB}}(t) &= \frac{1}{t} \int_0^t P_{\uparrow,m}^{\text{RSB}}(t') dt' \\ &= \frac{1}{2} \sum_{n=0}^{\infty} p(n+m) \left[1 - \frac{\gamma_n}{(\Omega_{n+m,n}^2 + \gamma_n^2) t} \right. \\ &\quad \left. + \frac{e^{-\gamma_n t} (\gamma_n \cos(\Omega_{n+m,n} t') - \Omega_{n+m,n} \sin(\Omega_{n+m,n} t))}{(\Omega_{n+m,n}^2 + \gamma_n^2) t} \right].\end{aligned}\quad (5.7)$$

We observe that for long times ($t \gg \gamma_n/(\Omega_{n+m,n}^2 + \gamma_n^2)$, $t \gg \Omega_{n+m,n}/(\Omega_{n+m,n}^2 + \gamma_n^2)$), the time average converges to a partial sum of motional state probabilities

$$\bar{P}_{\uparrow,m}^{\text{RSB}}(t) \approx \frac{1}{2} \sum_{n=0}^{\infty} p(n+m). \quad (5.8)$$

To extract the individual harmonic probabilities, consider driving with a first-order RSB:

$$\begin{aligned}\bar{P}_{\uparrow,1}^{\text{RSB}}(t) &\approx \frac{1}{2} \sum_{n=0}^{\infty} p(n+1) \\ &\approx \frac{1}{2} [1 - p(0)]\end{aligned}\quad (5.9)$$

from which $p(0)$ can be directly estimated

$$p(0) \approx 1 - 2\bar{P}_{\uparrow,1}^{\text{RSB}}(t). \quad (5.10)$$

Higher harmonic state probabilities may then be estimated by driving with sequentially higher-order RSBs and applying the recursion relation

$$p(m-1) \approx 2(\bar{P}_{\uparrow,m-1} - \bar{P}_{\uparrow,m}). \quad (5.11)$$

This time-average approach provides an efficient and robust method for extracting motional state populations. Compared with existing methods, relatively few points are needed to determine the time average of the RSB oscillation. Although these points should be taken at long times (relative to the RSB Rabi frequency), we note that Eq. (5.8) does not depend on the decoherence rate γ , and indeed converges *faster* when decoherence is included. Rather, we anticipate that the largest errors in time-average measurements will arise from real-time changes in $p(n)$ driven by motional heating. Such trap heating effects have been studied (Ref. [162] and Ch. 6) and can be incorporated into the motional state analysis if needed.

Some noise sources can induce a bias in the estimate by adjusting the 1/2 reference point to a slightly lower value. Static or fluctuating AC Stark shifts, asymmetric detection errors, or detuned drives may all move the half contrast value of 1/2 in Eq. (5.9), resulting in an under estimation of $p(n)$. These effects are not limiting however as they can be systematically measured by independent measurements and removed.

5.5 Experimental Thermometry

In this section, we experimentally demonstrate the effectiveness of our time-averaged thermometry method. We begin by measuring the temperature of a trapped ion following Doppler cooling and comparing the time-average method to several existing techniques. We then repeat our measurements and comparisons using an optimized sideband cooling sequence from Sec. 4.2, finding that the time-average method most closely agrees with theory predictions.

Thermometry experiments are performed on a single $^{171}\text{Yb}^+$ ion confined in a linear Paul trap with axial frequency $\omega_z = 2\pi \times 0.670 \pm 0.008$ MHz. In our setup, the Lamb-Dicke parameter $\eta = 0.18 \pm 0.01$, the Rabi carrier frequency $\Omega = 2\pi \times 64.9 \pm 0.5$ kHz, and the optical pumping time is 5 μs . Doppler cooling is performed with 369.5 nm light along the $^2S_{1/2}|F=0\rangle \rightarrow ^2P_{1/2}|F=1\rangle$ and $^2S_{1/2}|F=1\rangle \rightarrow ^2P_{1/2}|F=0\rangle$ transitions (linewidth

$\Gamma = 2\pi \times 19.6$ MHz), while red and blue sideband transitions are performed with far-detuned Raman beams at 355 nm. After each experiment, the qubit state is determined by irradiating the ion with 369.5 nm light resonant with the $^2S_{1/2}|F = 1\rangle \rightarrow ^2P_{1/2}|F = 0\rangle$ transition and capturing the spin-dependent fluorescence on a photomultiplier tube.

5.5.1 Thermal Distribution

When an ion is cooled to its Doppler-limited temperature, the motional state is well-characterized by a thermal distribution (Eq. (4.2)). Given our axial trap frequency, this temperature corresponds to an average harmonic occupation $\bar{n}_{\text{Dop}} = 14.6 \pm 0.2$ (Eq. (4.1)). We take this value as the theoretical prediction, against which we compare several different methods for trapped-ion thermometry.

We begin by using the ratio method to estimate the Doppler-cooled ion temperature. Figs. 5.1(a)-(b) show frequency scans over the red and blue sidebands, respectively, with error bars smaller than the size of the markers. Sinc squared functions are fit to the data with excellent agreement and shown as solid lines. Taking the ratio of the RSB and BSB transition strengths (Eq. (5.1)) yields $\bar{n}_{\text{ratio}} = 14.3 \pm 1.5$, in good agreement with the Doppler-limited prediction.

Two additional estimates of the Doppler-limited temperature may be extracted by driving a first-order BSB oscillation. In the first method, the data is fit to a thermally-weighted Rabi oscillation $P_{\uparrow}^{\text{BSB}}(t) = \sum_{n=0}^{800} p_{\text{th}}(n) \sin^2(\Omega_{n,n+1}t/2)$, shown as the solid light blue curve in Fig. 5.1(c). This single-parameter fit finds an estimated $\bar{n}_{\text{thermal fit}} = 14.9 \pm 0.7$. Using the same BSB data set, we also employ the SVD method to estimate $\bar{n}_{\text{SVD}} = 16.4 \pm 2.1$. In Fig. 5.1(c), the dashed black curve is calculated by weighting a BSB oscillation function $P_{\uparrow}^{\text{BSB}}(t) = \sum_{n=0}^{n_{\text{SVD}}} p_{\text{SVD}}(n) \sin^2(\Omega_{n,n+1}t/2)$ with the SVD-computed probabilities $p_{\text{SVD}}(n)$.

Finally, the first (dark green) and second (dark purple) RSBs are driven over a long period of time, with their respective running time averages (light green and light purple) shown in Figs. 5.1(d)-(e). We take an excess of data points in our demonstration to confirm




Figure 5.1: Thermometry comparisons of thermally-distributed ion motional states. (a) and (b) are red and blue sideband frequency scans used to determine $\bar{n}$ from the ratio method. (c) shows BSB Rabi oscillation data (blue points) fit by both a thermally-weighted Rabi oscillation function (solid blue) and a SVD analysis (dashed black). (d) and (e) are long Rabi oscillations of the first- and second-order RSBs, respectively, with their running time-average values shown as solid lines. (f) compares these different thermometry methods against the calculated Doppler cooling limit of $\bar{n}_{\text{Dop}} = 14.6$.

the accuracy of this new technique, though we note that only ~ 20 data points at long times are needed to find the same $\bar{n}$ to within 5%. From the first-order RSB time average in Fig. 5.1(d), we estimate $p(0)$ using Eq. (5.10). Using the second-order RSB time average in Fig. 5.1(e) and the value for $p(0)$, $p(1)$ may be obtained from Eq. (5.11). Finally, fitting $p(0)$ and $p(1)$ to a thermal distribution yields $\bar{n}_{\text{time avg}} = 14.6 \pm 1.2$.

All extracted values of $\bar{n}$ are compared to the Doppler-limited prediction in Fig. 5.1(f). We conclude that all approaches studied here are viable methods for extracting the average harmonic occupation $\bar{n}$ when applied to thermal distributions. In the following section, we will re-apply these measurement techniques to sideband cooled ions, whose motional distributions are predicted to be significantly non-thermal.

5.5.2 Sideband Cooled Distribution

In this set of experiments, the ion is initially cooled to the Doppler limit of $\bar{n} = 14.6$, then further cooled using 25 first-order fixed SBC pulses (see Sec. 4.2.1). As shown in Fig. 4.3(b), this small number of pulses cannot reach the ground state using any SBC protocol when starting from such a large initial $\bar{n}$. Nevertheless, we will show that 25 SBC pulses is already sufficient to induce large discrepancies between different thermometry techniques.

The inherent nonthermal distribution of the sideband cooled ion is predicted to cause a significant bias in the ratio method's estimation of $\bar{n}$. Fig. 5.2 illustrates this point for the given experimental parameters. In Fig. 5.2(a), a simulated distribution after 25 first-order fixed SBC pulses (solid) is compared to a thermal distribution with the same $\bar{n}$ (dashed). The wide discrepancy indicates that a thermal state is a poor approximation for the post-SBC distribution.

To quantify the potential error in assuming a thermal distribution, Fig. 5.2(b) compares the $\bar{n}$ of the simulated distribution (solid) to the predicted result from the ratio method (dashed). The ratio method drastically underestimates $\bar{n}$ after just a few pulses, with almost a full order of magnitude difference by 25 pulses. We caution that when ratio-method




Figure 5.2: (a) Simulated motional state distribution after 25 first-order fixed pulses (solid), and a thermal distribution with the same $\bar{n}$ (dashed). (b) For any number of SBC pulses, the estimated $\bar{n}$ from the ratio method (dashed) is predicted to significantly underestimate the true $\bar{n}$ as calculated from the simulated distribution (solid).

thermometry is applied after significant SBC, it may result in misleadingly low estimates of ion temperatures and motional heating rates.

Following SBC, we show the Rabi oscillations of first-order red and blue sidebands in Fig. 5.3(a). The data points are connected (not fitted) to guide the eye, and errors at each point are the size of the marker. Under the assumptions of the ratio method, the ratio of the RSB to BSB at any point in the time provides a valid estimate of $\bar{n}$. We have calculated this ratio for all points in Fig. 5.3(a), and have plotted the corresponding $\bar{n}$ in Fig. 5.3(b).

For thermal distributions, as assumed by the ratio method, $\bar{n}$ should be constant at all times. In Fig. 5.3(b), the substantial differences in extracted $\bar{n}$ with time provide experimental evidence that the underlying state distribution is nonthermal. To estimate $\bar{n}$ in Fig. 5.3(b), we average over the varying $\bar{n}$ to find $\bar{n}_{\text{ratio}} = 0.58 \pm 0.56$. This value is a drastic *underestimate* of the predicted value $\bar{n}_{\text{sim}} = 3.57 \pm 0.58$, by almost a full order of magnitude. Furthermore, the simulated $\bar{n}$ does not account for ion heating or noise effects, which if included would make the discrepancy even larger.

Next, we applied a SVD analysis to the first-order BSB in Fig. 5.3(a). Since the tail of the SBC distribution is predicted to be long, we chose the length of the level probability vector $\vec{p}$ to maximize the number of physically constrained probabilities, $0 \leq p(n) \leq 1$. Nevertheless, the BSB time-series data remained poorly fit for any length of $\vec{p}$, and the most accurate SVD result ($\bar{n}_{\text{SVD}} = 8.0 \pm 1.3$) still significantly disagrees with the simulated average harmonic occupation.

Lastly, we apply our time-average measurement technique to a sideband cooled ion. We begin by driving the first (dark green), second (dark purple), and third (dark orange) RSBs for a long time period, as shown in Figs. 5.3(c)-(e). Following the time average procedure outlined in Sec. 5.4, $p(0)$, $p(1)$, $p(2)$, and $p(n > 2)$ are estimated from the measured time averages. Substituting these probabilities into Eq. (5.5) results in a measured $\bar{n}_{\text{time avg}} = 4.1 \pm 0.7$.

The estimated level distributions from the simulation, time average method, and SVD




Figure 5.3: Thermometry comparisons of a sideband cooled ion. (a) the measured first-order RSB and BSB time series. Points are connected to guide the eye. (b) $\bar{n}$ estimation at each time point using the ratio method (excluding the first few time steps). (c)-(e) long-time Rabi oscillations for the first, second, and third RSBs, respectively, with their running time averages drawn as solid lines. (f) population distributions as estimated by numeric simulation (black), time-averaged method (blue), and SVD (gray). (g) $\bar{n}$ measurements from the ratio method, SVD, and the time-average method are compared to a numeric simulation of SBC. Only the time-average method closely estimates $\bar{n}_{\text{sim}}$.

method are compared in Fig. 5.3(f). The numerically simulated distribution (black) follows a monotonic decrease in population for increasing n . The time average method (blue) finds similar monotonic behavior, with a relative excess of population in the $n = 1$ and $n = 2$ levels which we attribute to ion motional heating out of the $n = 0$ state [162]. In contrast, the distribution estimated by the SVD method (gray) is non-monotonic and exhibits a steep drop-off in population between $n = 7$ and $n = 8$, suggesting unphysical behavior which cannot be explained by standard heating models [162]. Of all the considered thermometry techniques, the time average method best matches the simulated level distributions, and it is the only method that does not significantly disagree with the simulated prediction $\bar{n}_{\text{sim}}$ (Fig. 5.3(g)).

5.6 Conclusion

We have introduced a new thermometry technique which more closely models the state distribution after SBC, and experimentally validated its performance. In contrast, we observe that the most common measurement technique can severely underestimate ion temperatures if extensive SBC is performed. We have argued that careful understanding of the expected state distributions is a necessary precondition for accurate thermometry.

Finally, the time-average technique can open new possibilities for improved thermometry. With this method, for instance, it should be possible to probe the time-dependent population dynamics of trapped-ion motional states and observe how the harmonic level distribution changes in response to external noise sources which will be covered in the next chapter (Ch. 6). Such experiments would provide an additional set of characterizations which may help elucidate mechanisms responsible for anomalous ion heating.

Chapter 6

Measurement-induced Heating

We experimentally study the heating of trapped atomic ions during measurement of their internal qubit states. During measurement, ions are projected into one of two basis states and discriminated by their state-dependent fluorescence. We observe that ions in the fluorescing state rapidly scatter photons and heat at a rate of $\dot{n} \sim 2 \times 10^4$ quanta/s, which is ~ 30 times faster than the anomalous ion heating rate. We introduce a quantum trajectory-based framework that accurately reproduces the experimental results and provides a unified description of ion heating for both continuous and discrete sources.

6.1 Introduction

A trapped ion heats when it gains motional energy from its surrounding environment. Ion motion is the medium by which quantum information is transferred between qubits [45], and motional heating is detrimental to high-fidelity quantum operations. For instance, heating is known to disrupt the ions' phase space trajectories during the application of entangling operations and necessarily leads to quantum gate errors [163, 164, 64]. Moreover, as ions heat, they become even more susceptible to errors arising from non-closure of phase space trajectories, noise in the driving fields, or motional frequency drifts [165, 64, 166]. Effects such as anomalous ion heating are so pernicious that they set constraints on ion trap designs [167, 162, 168] and entangling gate timings [19, 18], and they motivate specialized preparation of trap electrode surfaces [169, 170] and operation of traps in cryogenic environments [171, 172, 173].

Less explored are the heating effects from measurements in the middle of a quantum circuit, which are fundamental to multiple areas of quantum information processing. For instance, many quantum error-correcting protocols rely upon 'mid-circuit' measurements and

feedforward to correct errors and provide fault-tolerant operations [174, 30, 31]. In addition, mid-circuit measurements are central to measurement-based quantum computing schemes [175, 176] and entanglement phase transitions in quantum many-body systems [177, 178, 179, 180, 181]. They may also provide a more efficient way of encoding quantum algorithms on NISQ-era hardware [182, 183, 23]. Yet it is currently unknown whether heating effects from mid-circuit measurements, which arise from fundamental photon-atom interactions, are important compared to anomalous heating. Consequently, it is unclear whether dedicated recooling strategies are truly required to avoid significant quantum gate errors following a mid-circuit measurement.

Here, we experimentally study the measurement-induced heating of a trapped ion. We first establish the anomalous heating rate under ambient conditions, which is then compared to the observed heating rate during measurement of the qubit state. We find a measurement-induced heating rate that is ~ 30 times larger than the ambient heating in our trap. We develop a generalized theoretical framework to describe heating from continuous noise sources (leading to anomalous heating) and discrete noise sources (such as photon absorption and emission, responsible for heating during measurement). This framework supports our experimental observation that measurement-induced heating cannot be avoided by specific choices of detection laser parameters. We conclude that dedicated recooling strategies will be required for high-fidelity quantum operations following mid-circuit ion measurements.

6.2 Ambient Heating

Before characterizing the effects of measurement-induced heating, we first measure our baseline ambient heating conditions and introduce a predictive theoretical framework to describe heating in general. Motional state heating $\dot{\bar{n}}$ is defined as the rate at which the average phonon occupation $\bar{n} = \langle \hat{a}^\dagger \hat{a} \rangle$ increases per unit time [162]. Under ambient experimental conditions, $\bar{n}$ will increase due to the interaction of ions with various sources of external noise [5, 162, 184], which generically may be modeled as a continuous, time-dependent fluctuating

field [185].

Experiments are performed with a single $^{171}\text{Yb}^+$ ion confined in a four-rod linear Paul trap detailed in Ref. [54] with radial secular frequency $\omega = 2\pi \times 1.09$ MHz. Doppler cooling of the ions is accomplished by irradiating the $369.5 \text{ nm } ^2\text{S}_{1/2} |F=0\rangle \rightarrow ^2\text{P}_{1/2} |F=1\rangle$ and $^2\text{S}_{1/2} |F=1\rangle \rightarrow ^2\text{P}_{1/2} |F=0\rangle$ transitions, with the dark qubit state $|0\rangle \equiv ^2\text{S}_{1/2} |F=0\rangle$ prepared via optical pumping. Far-detuned Raman beams at 355 nm drive carrier transitions between the hyperfine qubit states $|0\rangle$ and $|1\rangle \equiv ^2\text{S}_{1/2} |F=1, m_F=0\rangle$ as well as red and blue sideband transitions [5]. The ion is cooled to near the motional ground state through pulsed resolved sideband cooling using second and first-order red sideband pulses [98]. The internal qubit states are detected by resonantly irradiating the $^2\text{S}_{1/2} |F=1\rangle \rightarrow ^2\text{P}_{1/2} |F=0\rangle$ transition of the ion at 369.5 nm and collecting the state-dependent fluorescence on a photomultiplier tube (PMT) for 1 ms. The combined state preparation and measurement errors are estimated to be $< 0.3\%$.

Our measurement of the ambient motional heating rate does not assume that the motional states follow a thermal distribution [104, 98, 100] and is executed as follows. The ion is first cooled to near the motional ground state. Then, the ion is left in the dark for a precise delay time gaining motional quanta due to ambient heating sources. A blue sideband is then driven from 0 - 300 μs (covering five periods) with 60 time points and 500 repetitions per point. The median value for the low-energy motional state probabilities is then computed using the Singular Value Decomposition (SVD) method [62, 186, 98] with Monte Carlo error propagation to estimate asymmetric 1σ confidence intervals.

As the delay time for ambient heating is varied, the motional state dynamics under ambient heating conditions are shown in Fig. 6.1(a). We plot the occupation probability for the lowest two motional energy levels $|n=0\rangle$ and $|n=1\rangle$ which have the largest dynamic range. Heating dynamics are observed as probability decays out of $|n=0\rangle$ and into $|n=1\rangle$, leading to a temporary increase of the $|n=1\rangle$ probability at early times. Afterwards, population in the states $|n=0\rangle$ and $|n=1\rangle$ monotonically decays towards thermal equilibrium with the

environment.

The motional state dynamics may be modeled as an ion coupled to a reservoir, or ‘bath’, outlined and experimentally observed in Ref. [186]. The dynamics of the n th diagonal element of the motional state density matrix $\rho_{n,n}(t)$ follow the solution to the master equation for a harmonic oscillator weakly coupled to a high temperature reservoir—which is the case for ambient experimental conditions—and given by [186]

$$\begin{aligned} \rho_{n,n}(t) = & \frac{1}{1 + \dot{n}t} \sum_{j=0}^n \left(\frac{\dot{n}t}{1 + \dot{n}t} \right)^j \left(\frac{1}{1 + \dot{n}t} \right)^{2n-2j} \\ & \times \sum_{l=0}^{\infty} \left(\frac{\dot{n}t}{1 + \dot{n}t} \right)^l \binom{n+l-j}{n-j} \\ & \times \binom{n}{j} \rho_{n+l-j, n+l-j}(0) \end{aligned} \quad (6.1)$$

where $\dot{n}$ is the linear heating rate and the only free parameter. Equation (6.1) is fit to the measured motional state probabilities, yielding an ambient heating rate of $\dot{n} = 770 \pm 20 \text{ s}^{-1}$. In Fig. 6.1(b) we show the estimated $\bar{n}$ from a cumulative fit to the bath model. To estimate $\bar{n}$, we describe the initial motional state by a double thermal distribution (Appendix B.3), estimated from the initial measured $|n=0\rangle$ and $|n=1\rangle$ values and extended to include the first 100 motional states [104, 98] (Appendix B.3). We then compute multiple fits of the motional heating rate $\dot{n}$ to the data using the bath model. A cumulative fit is used for each point, such that $\dot{n}_k$ determined at the k th time point only includes data up to time t_k . Each heating rate $\dot{n}_k$ is then used to propagate the initial motional state following the bath model dynamics, resulting in an estimate of $\bar{n}$ at time t_k . Both the data and the bath model exhibit linear heating, as expected for trapped ion systems subject to continuous fluctuations of electric fields.




Figure 6.1: Motional state dynamics under ambient heating conditions. (a) Motional state probabilities are measured using the SVD method. Both a bath model [186] (dashed red) and our Quantum Trajectory Theory (QTT) model (solid black) show good agreement with the experimental data. (b) The estimated $\bar{n}$ under ambient conditions shows linear heating and is well described by both models.

6.2.1 QTT Model

We introduce a framework based on semi-classical quantum trajectory theory (QTT) [187, 188] which can accurately predict motional state dynamics, $\bar{n}$, and $\dot{\bar{n}}$ for both continuous and discrete sources of ion heating. We first describe the approach, then establish its validity for continuous noise sources under ambient heating conditions by comparing it to our

experimental data and to the bath model.

Ambient heating sources are modeled as an effective fluctuating electric field $E(t)$ which captures a wide range of physical noise sources [185]. To compute the quantum trajectory along the radial direction of interest, the classical center-of-mass phase space coordinate

$$\alpha(t) = \sqrt{\frac{m\omega}{2\hbar}}\hat{x}(t) + \frac{i}{\sqrt{2m\omega\hbar}}\hat{p}(t) \quad (6.2)$$

with harmonic frequency ω and particle mass m is recorded as $E(t)$ shifts the ion in phase space $\alpha \rightarrow \alpha + \alpha_k$. We consider a time interval $\Delta t = t_{k+1} - t_k$ that is long compared to the correlation time of the electric field fluctuations and satisfies $|\alpha_k| \ll 1$. For a fluctuating electric field, a particle of charge e is shifted by [189]

$$\alpha_k = \frac{ie}{\sqrt{2m\omega\hbar}} \int_{t_k}^{t_{k+1}} E(t) e^{i\omega t} dt. \quad (6.3)$$

We then compute the motional state dynamics from the displacement $\hat{D}(\alpha)$ of the initial motional state $|n\rangle$ to Fock state $|m\rangle$. For $m \geq n$,

$$\begin{aligned} p_m(n) &= |\langle m | \hat{D}(\alpha) | n \rangle|^2 \\ &= \frac{n!}{m!} |\alpha|^{2(m-n)} e^{-|\alpha|^2} \left[\mathcal{L}_n^{(m-n)}(|\alpha|^2) \right]^2 \end{aligned} \quad (6.4)$$

where $\mathcal{L}_n^{(k)}(x)$ is the generalized Laguerre polynomial [190]. Combining this with the definition of α_k in Eq. (6.3) (see Appendix A.1), we recover the commonly defined heating rate in the literature [5, 162], $\dot{n} = e^2 S_E(\omega)/(4m\omega\hbar)$, where $S_E(\omega)$ is the spectral density of electric field fluctuations at the trap frequency. Quantum trajectories randomly sample electric fields, using a $S_E(\omega)$ independently determined by the bath model, and are averaged together to provide the expected motional heating. On average, the microscopic phase space kicks α_k are connected to the macroscopic observable $\bar{n}$ since $\langle |\alpha_k|^2 \rangle = \dot{n} \Delta t$.

Figure 6.1 shows that the QTT model, averaged over 1000 trajectories, closely agrees

with the experimental data and the bath model in predicting motional state dynamics and $\bar{n}$. We emphasize that our QTT model contains no adjustable parameters and depends only on the physical quantities in Eq. (6.3). In the next section, we will show how QTT may be readily adapted to discrete heating sources, such as photon kicks, by adjusting the definition of α_k accordingly.

6.3 Measurement-Induced Heating

We now consider the heating of a trapped ion irradiated with a resonant detection beam, which is the standard configuration for qubit state readout. Two experiments are performed: one where an ion is prepared in the dark qubit state ($|0\rangle$), and a second where an ion is prepared in the bright qubit state ($|1\rangle$). In each case, the ion is first cooled to near its motional ground state by pulsed sideband cooling [98]. If a bright qubit state is desired, we drive a carrier π -pulse from $|0\rangle$ to $|1\rangle$ with measured fidelity 99.3 ± 0.1 %. The detection beam, which has saturation parameter $s = 1.27 \pm 0.02$ and scatters photons at a rate $\Gamma = 2\pi \times 1.07$ MHz [81], is then turned on for a variable heating time. An optical pumping pulse then resets the qubit state to $|0\rangle$ while preserving the newly excited phonon state $|n\rangle$. Finally, the phonon state probabilities are measured using the same SVD method outlined in the previous section.

We expect that ions in the dark state will scatter no photons during measurement and exhibit the same heating rate as the ambient case. This expectation is confirmed in Fig. 6.2(a), which shows the probability of $|n=0\rangle$ for dark ions during measurement. Data were taken out to 8 ms and fit to a global bath model, yielding an estimated motional heating rate of 780 ± 40 s⁻¹. This rate is indistinguishable from the ambient case presented in Fig. 6.1(b) (770 ± 20 s⁻¹).

In stark contrast, ions in the bright qubit state will undergo stochastic momentum kicks of order $\sim \hbar k$ due to photon absorption and emission. This process cannot be generically described by the bath model of Eq. (6.1). Although photon scattering does not increase the




Figure 6.2: Motional state dynamics during measurement-induced heating. (a) Motional ground state $|n=0\rangle$ probability for a dark ion. The bath model (dashed red) under the ambient conditions of Fig. 6.1(a) is included for comparison. (b) The motional ground state probability for bright-state ions decays much more rapidly than for dark-state ions. (c) Early-time population dynamics of a bright ion undergoing measurement-induced heating. The QTT model accurately describes the data with no adjustable parameters. (d) Estimates of $\bar{n}$ based on the panel (c) data, comparing the QTT model to a thermal model. Shaded bands indicate 1σ model uncertainties.

average velocity of the ion ($\langle v \rangle = 0$ in a trap), the stochastic timing of photon recoil events relative to the harmonic motion of the ion is expected to increase velocity fluctuations $\langle v^2 \rangle$ and therefore increase motional energy [53]. This mechanism is observed in Fig. 6.2(b) where the $|n = 0\rangle$ probability of a prepared bright ion decays rapidly, indicating fast heating out of the motional ground state.

We model this process using our QTT framework, where the phase-space kick α_k now encodes the effects of atomic scattering processes. We write (Appendix A.2):

$$\alpha_k = \frac{ie^{i\omega t}\hbar k}{\sqrt{2m\omega\hbar}} \left[\sqrt{f_x} + \sqrt{f_{sx}^{(k)}} + \frac{8\sqrt{f_x}\Delta\omega n}{\gamma^2(1+s') + 4\Delta^2} \right] \quad (6.5)$$

where $\omega = 2\pi \times 1.09$ MHz is the trap frequency along the radial x -axis, $k = 2\pi/(369.5 \text{ nm})$ is the wave vector, and n is the ion motional state. In Eq. (6.5), absorption is described by the geometric factor $f_x = 1/4$ which accounts for the projection of the incident laser beam along the x -axis. Emission is treated with geometric factor $f_{sx}^{(k)}$, randomly chosen to recreate an isotropic emission pattern [53, 188]. Finally, the Doppler effect is included as the final term in Eq. (6.5) and depends on parameters such as the natural linewidth γ , modified saturation parameter s' , and laser detuning Δ [81] (Appendix A.2).

We report close agreement between the QTT model and the experimental data as shown in Fig. 6.2. The model accurately predicts the dynamics of the motional ground state (Fig. 6.2(a)-(b)) and all measured low-lying motional states (Fig. 6.2(c)). This agreement at short times is particularly noteworthy, since most heating models assume thermal state distributions and would fail to capture the initial non-thermal behavior following sideband cooling [104, 98, 100]. As before, the QTT model in Fig. 6.2 contains no adjustable parameters and depends only on the atomic physics and laser properties described in Eq. (6.5).

Figure 6.2(d) shows the rapid increase in $\bar{n}$ during measurement-induced heating. The QTT prediction (dark blue line) shows a linear heating rate of $\dot{\bar{n}} = 22900 \pm 240 \text{ s}^{-1}$, which is nearly 30 times faster than the heating rate for ions in the dark state or under ambient heating

conditions. For comparison, we also fit the measured low-energy states $|n = 0\rangle$, $|n = 1\rangle$, and $|n = 2\rangle$ to a presumed thermal distribution at each time point (dashed black line). The thermal fit significantly underestimates $\bar{n}$ at early times, consistent with previous numeric simulations and experimental observations [104, 98, 100] and only matches the experimental data after several hundred scattering events have taken place.

6.4 Long-time Behavior

Since the typical qubit measurement time for trapped ions is significantly longer than the 100 μs timescale shown in Fig. 6.2(c)-(d), we investigate whether this rapid linear heating persists at longer times. As seen in Eq. (6.5), the choice of laser detuning Δ plays a key role in determining the typical magnitude of phase-space kicks during photon scattering. Therefore, we study measurement-induced heating in the long-time limit for three different choices of Δ , spanning the red-detuned, near-resonant, and blue-detuned regimes.

We estimate $\bar{n}$ for heating times up to 2 ms using a measured carrier Rabi oscillation. Following measurement-induced heating of the trapped ion, a carrier oscillation is driven from 0 - 40 μs (covering 6 periods) with 60 time points and 300 repetitions each. The carrier oscillation is fit with $\bar{n}$ as a free parameter. Far-detuned 355 nm Raman beams couple to the x and y principal axes in this experiment, so both modes affect the carrier Rabi oscillation. A carrier Rabi oscillation with x and y COM mode couplings is given by

$$P_{\text{bright}}(t) = \sum_{n_x, n_y=0}^{\infty} p_{n_x} p_{n_y} \sin(\Omega_{n_x, n_y} t/2)^2 \quad (6.6)$$

where the Rabi frequency for the x and y mode couplings is given by

$$\Omega_{n_x, n_y} = \Omega_0 e^{-\eta_x^2/2} e^{-\eta_y^2/2} \mathcal{L}_{n_x}(\eta_x^2) \mathcal{L}_{n_y}(\eta_y^2) \quad (6.7)$$

with $\eta_x = 0.104$ and $\eta_y = 0.112$ measured independently [5]. The thermal contribution from

the y -axis can be removed by approximating its contribution with the ratio of the respective secular frequencies $\bar{n}_y = (\omega_x/\omega_y)\bar{n}_x = 1.48\bar{n}_x$. The motional distribution is assumed to be thermal. The final fitting parameters are then the Rabi frequency and average motional state along the x -axis: $\{\Omega_0, \bar{n}_x\}$. This fitting procedure assumes a thermal motional state distribution, which is in agreement with our data after $\sim 50 \mu\text{s}$ (Fig. 6.2(d)).

Figure 6.3 shows the sensitivity of ion motional heating to laser detuning Δ in the long-time limit. We perform the experiment with three different detection beam detunings: $\Delta = 2\pi \times \{-11, -1, 9\}$ MHz, corresponding to red-detuned, near-resonant, and blue-detuned, respectively. In Fig. 6.3, the x -axis is scaled by the photon scattering rate of each detuning, $\Gamma = 2\pi \times \{0.56, 0.94, 0.67\}$ MHz, with our typical state detection requiring $\Gamma t \approx 5 \times 10^3$ scattering events. For the red-detuned case, where Δ is chosen near the optimal Doppler detuning, the data equilibrates to the Doppler cooling limit $\bar{n} \approx 12.7$. All other choices for Δ leave the ion with higher motional energy in the long-time limit.

The expected $\bar{n}(t)$ under these conditions may be predicted by extending the QTT model with Eq. (6.5) to long times, or by adapting the equations governing laser cooling [53] to describe measurement-induced heating. We find (Appendix A.3)

$$\bar{n}(t) = \left[\left(\bar{n}(0) + \frac{R}{D\hbar\omega} \right) e^{\Gamma_0 Dt} - \frac{R}{D\hbar\omega} \right] \quad (6.8)$$

where we define $\Gamma_0 \equiv \gamma(s/18)/(1 + s' + 4\Delta^2/\gamma^2)$ as the effective scattering rate, $R \equiv (f_x + f_{sx})\hbar^2 k^2/(2m)$ as the effective recoil energy, and $D \equiv 8\Delta\hbar f_x k^2/(m\gamma^2(1 + s') + 4m\Delta^2)$ as a Doppler-effect term that takes into account the laser detuning Δ . In the expressions above, we assume that spontaneous emission is isotropic ($f_{sx} = 1/3$) and that the initial Doppler shift $\vec{k} \cdot \vec{v}$ is small compared to the natural linewidth γ . The parameters in Eq. (6.8) are independently measured, and the predicted curves for each detuning are plotted as solid lines in Fig. 6.3. Shaded bands illustrate our ± 2 MHz experimental uncertainty in measuring Δ .

From this model and the data in Fig. 6.3, we find that rapid measurement-induced




Figure 6.3: Measurement-induced heating depends sensitively on the detuning Δ of the detection beam. Darker solid lines are predicted values according to Eq. (6.8) with shaded areas showing ± 2 MHz uncertainty bands in detuning. Points and lines of the same color have the same detuning.

heating cannot be avoided by judicious choice of detection beam parameters. At the start of a measurement, the short-time behavior of Eq. (6.8) may be approximated as $\bar{n}(t) \approx \bar{n}(0) + \Gamma_0 R t / (\hbar \omega)$. This corresponds to linear initial heating at rate $\dot{\bar{n}} = \Gamma_0 R / (\hbar \omega) \sim 2 \times 10^4 \text{ s}^{-1}$, in agreement with measurements in Fig. 6.2(d). The only tunable beam parameter is the photon scattering rate Γ_0 , which cannot be reduced without sacrificing the qubit state detection fidelity or increasing the detection time (leading to the same final $\bar{n}$). We also note that this measurement-induced heating rate in $^{171}\text{Yb}^+$ is among the *smallest* for common trapped-ion species, since the effective recoil energy R , and hence $\dot{\bar{n}}$, are suppressed by the large ion mass.

6.5 Discussion and Outlook

Mid-circuit measurements will be a critical element to future quantum technologies. This work has quantified the rapid motional heating which takes place during the measurement of a trapped ion, which dominates over anomalous heating and is an unavoidable consequence of photon scattering. We have also provided a unified framework (QTT) to accurately model motional heating effects for both continuous and discrete noise sources. Our experimental data and theoretical modeling indicate that, if left unmitigated, measurement-induced heating will present a substantial roadblock for performing high-fidelity operations following a mid-circuit measurement.

The effects of measurement-induced heating become more complex when detecting the state of specific ions in a larger array. For a collection of co-trapped ions, measurement-induced heating will add motional energy to all vibrational modes in which a bright ion participates, weighted by their mode participation amplitudes. These bright ions would act like local hot reservoirs, with the motional energy spreading quickly relative to typical detection time scales.

We conclude that for mid-circuit measurements to be viable, dedicated mid-circuit recoiling strategies must be implemented to combat measurement-induced heating. Brute-force sympathetic recoiling after measurement has been successful, at the cost of co-trapping multiple ion species and spending more time recoiling than executing quantum gate operations [22, 23]. More sophisticated recoiling schemes may be needed, such as rapid exchange cooling [191], phonon rapid adiabatic passage cooling [192], or perhaps dedicated mid-circuit measurement zones within a QCCD trap [193, 22]. Future work may also consider simultaneous ion measurement and recoiling by utilizing shelving, state teleportation [194], or the *omg* qubit architecture [44].

Chapter 7

Conclusion and Outlook

As a whole, this work has presented a complete view of thermal ion motion: heating, cooling, and thermometry. I anticipate a combined view of the presented works to be fruitful, though currently unexplored. The new time-average thermometry technique has been shown to have greater accuracy than typical techniques. The assumption that motional state probabilities follow a thermal distribution has been questioned and shown may not be the case after ground state cooling to significant effect. A double thermal distribution assumption has been shown to better match simulations and measurements.

A new model of sideband cooling has demonstrated (Ch. 4) efficient modeling of new sideband cooling protocols which are faster than traditional methods. The graph-model approach is adaptable any pulsed cooling scheme. In future work, we anticipate that the graph representation of pulsed SBC may be expanded to include noise models for ion heating (Ch. 6), decoherence, off-resonant couplings, and affects of rf-driven micromotion. Extending to multiple ions and multiple modes is another natural direction that fits nicely within the matrix representation of pulsed SBC. With added degrees of freedom, more advanced computational methods may be needed as the matrices grow. The matrices are sparse however and do not grow exponentially. Speed improvements which take advantage of this could extend this model to 10s of ions. Despite these improvements, ground state cooling can take up to two orders of magnitude longer than the actual quantum operations executed during an experiment [22, 23]. In the near-term, we do not anticipate changing ground state cooling protocols, but later on, we may start ground state cooling with Electromagnetically Induced Transparency (EIT) cooling [101, 99] which can cool to near the ground state in ~ 1 ms and over a wide bandwidth.

A new thermometry technique, the time-average method, has been introduced (Ch. 5)

which more closely models the state distribution after sideband cooling. Its performance has been experimentally validated. In contrast, we observe that the most common measurement technique can severely underestimate ion temperatures after sideband cooling. We have observed that careful understanding of the expected motional state distribution is a necessary precondition for accurate thermometry. Finally, the time-average technique can open new possibilities for probing motional states. With this method, for instance, it should be possible to observe the time-dependent population *dynamics* of trapped-ion motional states. It would be particularly insightful to observe how motional state probabilities change in response to external noise sources to help elucidate mechanisms responsible for anomalous ion heating. We pursue this course in Ch. 6.

Finally, in Ch. 6 we have provided a unified framework to accurately model motional heating effects from both continuous and discrete noise sources. This is a timely result as no unified model previously existed and mid-circuit measurements, which rely on photon scattering, will be a critical element to future quantum technologies. This work has quantified the rapid and unavoidable measurement-induced heating rate to be $\dot{n} \sim 2 \times 10^4$ quanta/s. The experimental data and theoretical modeling indicate that, if left unmitigated, measurement-induced heating will present a roadblock for subsequent high-fidelity operations following a mid-circuit measurement.

More sophisticated recooling schemes are currently being pursued, such as rapid exchange cooling [191], phonon rapid adiabatic passage cooling [192], or perhaps dedicated mid-circuit measurement zones within a QCCD trap [193]. Future work may also consider simultaneous ion measurement and recooling by utilizing shelving, state teleportation [194], or the *omg* qubit architecture [44].

Motional ground state cooling is critical to trapped ion performance, but it is difficult to get to the ground state and stay there. Understanding cooling, heating, and thermometry at a new unprecedented level has been interesting in its own right. I also hope trapped ion systems used as academic scientific instruments or industrial quantum technologies will

benefit from this work.

Appendix A

Unified heating formalism

The theoretical elements of ion heating are outlined for continuous and discrete sources and for longer times. Numerical details can be found in Appendix B.

A.1 Derivation of Continuous Heating Rate from QTT Model

To compute the semi-classical quantum trajectories, the classical center-of-mass phase space coordinate

$$\alpha(t) = \sqrt{\frac{m\omega}{2\hbar}}\hat{x}(t) + \frac{i}{\sqrt{2m\omega\hbar}}\hat{p}(t) \quad (\text{A.1})$$

is recorded as external sources stochastically shift $\alpha \rightarrow \alpha + \alpha_k$ over a small time interval $\Delta t = t_{k+1} - t_k$ following Monte Carlo rules [187, 188]. Ambient heating sources are modeled as an effective fluctuating electric field $E(t)$, which captures a wide range of physical noise sources [185]. The resulting shift in phase space at moment t_k is given by [189]

$$\alpha_k = \frac{ie}{\sqrt{2m\omega\hbar}} \int_{t_k}^{t_{k+1}} E(t') e^{-i\omega t'} dt' . \quad (\text{A.2})$$

The jump in phase space has connections to the common definition of heating due to fluctuating electric fields. Consider

$$|\alpha_k|^2 = \frac{e^2}{2m\omega\hbar} \left| \int_{t_k}^{t_{k+1}} E(t') dt' \right|^2 \quad (\text{A.3})$$

which is related to the single-sided power spectral density and recovers the commonly defined heating rate defined in the literature [162]

$$|\alpha_k|^2 = \frac{e^2}{2m\omega\hbar} S_E(\omega) \Delta t \quad (\text{A.4})$$

where $|\alpha_k|^2 = \dot{n}\Delta t$.

A.2 Derivation of Discrete Heating Rate from QTT Model

Calculation of the change in phase space due to discrete photon scattering events. Near-resonance incident photons apply an average force on $^{171}\text{Yb}^+$ in the x -direction, quoted from Ref. [53] and using the scattering rate from Ref. [81],

$$\langle F_x \rangle = \Gamma_0 \hbar k \sqrt{f_x} \left(1 + \frac{8\Delta n\omega}{\gamma^2(1+s') + 4\Delta^2} \right) \quad (\text{A.5})$$

with variable definitions matching those given the text and

$$\Gamma_0 = \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}}. \quad (\text{A.6})$$

The momentum kick per absorption event Δp_{abs} includes Doppler related frequency shifts

$$\Delta p_{\text{abs}} = \hbar k \sqrt{f_x} \left(1 + \frac{8\Delta n\omega}{\gamma^2(1+s') + 4\Delta^2} \right) \quad (\text{A.7})$$

Momentum kicks due to emission are independent of velocity, delivering a momentum change

$$\Delta p_{\text{em}} = \hbar k \sin(\theta) \cos(\phi) \quad (\text{A.8})$$

with angles θ and ϕ randomly chosen from an isotropic emission pattern [53].

The center-of-mass phase space coordinator is defined as

$$\alpha(t) = \frac{1}{\sqrt{2}} \left[\sqrt{\frac{m\omega}{\hbar}} \hat{x}(t) + \frac{i}{\sqrt{m\omega\hbar}} \hat{p}(t) \right]. \quad (\text{A.9})$$

The change in phase space α_k due to a photon scattering event (absorption and emission)

then is calculated as

$$\begin{aligned}\alpha_k &= \frac{ie^{i\omega t}}{\sqrt{2m\omega\hbar}}\Delta p \\ \alpha_k &= \frac{ie^{i\omega t}\hbar k}{\sqrt{2m\omega\hbar}} \left[\sin(\theta)\cos(\phi) \right. \\ &\quad \left. + \sqrt{f_x} \left(1 + \frac{8\Delta n\omega}{\gamma^2(1+s') + 4\Delta^2} \right) \right]\end{aligned}\tag{A.10}$$

where the change in α due to absorption events has already been averaged over while the change in α due to emission will be averaged over in the quantum trajectory Monte Carlo. A single scattering event is on the order of 10 ns while the trap period is 1 μ s. Therefore, we have assumed the same phase for an absorption and emission cycle.

This microscopic calculation of α_k estimates the exact same heating dynamics $\dot{n}$ as the more common semi-classical Doppler shifted energy equations [53] as will be shown below. The average motional energy $E = \hbar\omega|\alpha_k|^2$ is given by

$$E = \frac{\hbar^2 k^2}{2m} \left[f_x \left(1 + \frac{8\Delta n\omega}{\gamma^2(1+s') + 4\Delta^2} \right)^2 + f_{sx} \right]\tag{A.11}$$

where the emission geometry factor $f_{sx} = \langle \sin(\theta)^2 \cos(\phi^2) \rangle = 1/4$. Applying the binomial approximation, the average motional energy E of Eq. (A.11) recovers the result from the semi-classical Doppler derivation.

$$E = \left[\frac{\hbar^2 k^2}{2m} (f_x + f_{sx}) + \frac{8\Delta\hbar f_x k^2/m}{\gamma^2(1+s') + 4\Delta^2} n\hbar\omega \right]\tag{A.12}$$

This approach is in complete agreement with the semi-classical analysis. The energy change per scattering event is given by $\hbar\omega|\alpha_k|^2$, and the total rate of energy change $dE/dt = \Gamma_0\hbar\omega|\alpha_k|^2$. Comparing this expression with the semi-classical calculation for dE/dt Eq. (A.18), we arrive at the same equation but having derived the result from a microscopic phase space displacement starting point.

A.3 Semi-classical Laser Heating

We treat the regime where near-resonant laser irradiating a harmonically trapped atom can cause heating of motional states. The treatment initially follows Ref. [53] but then deviates to consider the case of heating due to the laser frequency being near-resonance or blue-detuned ($\Delta > 0$).

Consider a resonant beam of intensity I and wavevector $\vec{k}$ near 369.5 nm. When this light hits a trapped ion moving with velocity $\vec{v}$ and a photon is absorbed, a photon is later emitted with wavevector $\vec{k}_s$.

To find the energy rate of change along a direction of interest (i.e. x -axis), we multiply the change in energy by the scattering rate $\Gamma(\omega, \vec{v})$ and average over the absorption and emission scattering directions:

$$\frac{dE_x}{dt} = \langle \Gamma(\omega, \vec{v}) \left[\frac{\hbar^2 k^2}{2m} (f_x + f_{sx}) + \hbar k_x v_x \right] \rangle \quad (\text{A.13})$$

where m is the mass of the trapped ion, v_x is the velocity in along the x -axis, k is the magnitude of $\vec{k}$, $f_x = \hat{k}_x^2$ is the incident photon geometric factor along the x -direction, and $f_{sx} = \hat{k}_{sx}^2$ is the emitted photon geometric factor along the x -direction.

The scattering rate $\Gamma(\omega, \vec{v})$ from Refs. [81, 82] is given by:

$$\Gamma(\omega, \vec{v}) = \frac{\gamma(s/18)}{1 + s' + \frac{4(\Delta + \vec{k} \cdot \vec{v})^2}{\gamma^2}} \quad (\text{A.14})$$

where $\gamma = 2\pi \times 19.6$ MHz is the natural linewidth, $s = I/I_{sat}$ is the saturation parameter, Δ is the detuning from resonance, and

$$s' = \frac{1}{216} \left(\frac{s\gamma}{\delta_B} \right)^2 + \frac{8}{3} \left(\frac{\delta_B}{\gamma} \right)^2 \quad (\text{A.15})$$

with Zeeman splitting $\delta_B = 2\pi \times 5.288$ MHz.

The Doppler shift $\vec{k} \cdot \vec{v}$ is much smaller than the natural linewidth such that the approximation of Eq. (A.14) is given by

$$\Gamma(\omega, \vec{v}) \approx \Gamma_0 \left[1 + \frac{8\Delta(\vec{k} \cdot \vec{v})}{\gamma^2(1 + s') + 4\Delta^2} \right] \quad (\text{A.16})$$

where

$$\Gamma_0 = \frac{\gamma(s/18)}{1 + s' + \frac{4\Delta^2}{\gamma^2}}. \quad (\text{A.17})$$

The change in energy (Eq. (A.13)), assuming $\langle v_i \rangle = 0$ and $\langle v_i v_j \rangle = 0$ for any two directions $i \neq j$, becomes

$$\frac{dE_x}{dt} = \Gamma_0 \left[\frac{\hbar^2 k^2}{2m} (f_x + f_{sx}) + \frac{8\Delta \hbar f_x k^2 E_x / m}{\gamma^2(1 + s') + 4\Delta^2} \right] \quad (\text{A.18})$$

where the average classical harmonic oscillator energy $E_x = m\langle v_x^2 \rangle$ has been swapped in (with no factor of 1/2).

Equation A.18 is a first order differential equation which can be solved analytically. For convenience, we define:

$$R \equiv \frac{\hbar^2 k^2}{2m} (f_x + f_{sx}) \quad (\text{A.19})$$

$$D \equiv \frac{8\Delta \hbar f_x k^2 / m}{\gamma^2(1 + s') + 4\Delta^2} \quad (\text{A.20})$$

then our equation may be written

$$\frac{dE_x}{dt} = \Gamma_0 (R + DE_x) \quad (\text{A.21})$$

which has the solution

$$E_x(t) = \left(E_x(0) + \frac{R}{D} \right) e^{\Gamma_0 D t} - \frac{R}{D}. \quad (\text{A.22})$$

Replacing the energy with the average motional state $\bar{n} = E_x(t)/(\hbar\omega)$, we arrive at equation

in the main text Eq. (6.8).

We examine Eq. (A.22) in a few limits. First, in the long-time limit and $\Delta < 0$ ($D < 0$), the final value of the energy converges to:

$$E_x(\infty) = -\frac{R}{D} = -\frac{\hbar\gamma}{8} \left(1 + \frac{f_{sx}}{f_x}\right) \left[\frac{\gamma(1+s')}{2\Delta} + \frac{2\Delta}{\gamma}\right] \quad (\text{A.23})$$

This agrees with results in the literature [53, 6], which are typically derived by setting $dE_x/dt = 0$ in Eq. (2.21) and define detuning Δ with opposite sign.

Another limit is the special case of zero detuning ($D = 0$). In the limit of $D \rightarrow 0$, the motional energy is readily given by:

$$E_x(t) = E_x(0) + \Gamma_0 R t \quad (\text{A.24})$$

With no damping force from cooling ($\Delta = 0$), the ion experiences linear heating.

Appendix B

Semi-classical Quantum Trajectory Numerics For Ion Heating

Quantum trajectory theory (QTT) simulates the stochastic evolution of a quantum state which may undergo wavefunction “collapse” by using Monte Carlo methods [187]. Many evolutions, or “trajectories,” are computed and ultimately averaged together to find ensemble averages which are equivalent to similar master equation calculations [187]. In a real sense, QTT mimics what would happen in a physical experiment.

This appendix briefly covers the relevant numerics and algorithm for computing the semi-classical quantum trajectories of trapped-ion motional states during heating as discussed in Ch. 6.

B.1 Fluctuating Fields

As stated in the text, fluctuating electric fields capture a wide range of possible noise sources as outlined in Ref. [185]: Mechanical vibrations of the trap, patch charges on the electrode surfaces, spurious man-made or natural electrical noise which contribute differentially, thermal noise on the electrodes, etc.

The Monte Carlo steps for computing quantum trajectories under these heating conditions are computed as follows:

- Evolve time by a small step Δt such that $|\alpha_k| \ll 1$
- Compute change in phase space α_k from Eq. (6.3)
- Record new phase space coordinate $\alpha(t)$
- Repeat previous steps thousands of times to generate statistics

In practice, we approximate the integral in Eq. (6.3), which is a Fourier transform of the

electric field, as a Gaussian distribution about a fixed value [189]. Under this assumption the mean of the Gaussian is then related to the heating rate of the ion by $\sqrt{2\hbar m\omega\dot{n}\Delta t}/e$.

Following the Monte Carlo steps and computing 10,000 trajectories, the final phase space values of all the trajectories are shown in Fig. B.1. The two histograms show the projected distributions on their respective axes. The distributions are nearly ideal Gaussian distributions suggesting the heating dynamics follow a random walk in phase space.




Figure B.1: Final α values of 10000 QTT trajectories under a continuous heating source. The histograms, top and bottom, reveal the normally distributed real and imaginary components of α and thus its random walk like behavior.

B.2 Discrete Photon-recoil

The Monte Carlo steps for computing quantum trajectories under measurement-induced heating are computed as follows. The ion begins in its internal ground state $|0\rangle$ and a motional Fock state $|n\rangle$. The initial motional Fock state is chosen randomly based on the measured initial state. Next, the probability of scattering a photon over a small time interval

Δt with atomic scattering rate R_{scat} is computed:

$$p_{\text{event}} = \Delta t R_{\text{scat}} . \quad (\text{B.1})$$

The small time interval Δt is chosen such that $p_{\text{event}} \ll 1$, a common requirement for Monte Carlo methods [187]. The probability of scattering p_{event} is compared to a random number from a uniform distribution $r_k \sim U[0, 1]$. If $p_{\text{event}} < r_k$, then the ion's phase space trajectory α does not change, but α is still recorded. If $p_{\text{event}} > r_k$, then the ion has scattered a photon (i.e. absorbed and emitted a photon) and its phase space trajectory is adjusted $\alpha \rightarrow \alpha + \alpha_k$ as described in Ch. 6. To compute the kick α_k , the spontaneous emission geometric factor for this random kick $f_{sx}^{(k)}$ is randomly computed such that, on average, it follows an isotropic emission pattern. The new phase space coordinate is recorded. This is then repeated thousands of times to get statistics. The probabilistic stepping of the ion's motion through phase space leads to a stochastic random walk trajectory just like the shape of Fig. B.1.

As an aside, to satisfy $p_{\text{event}} \ll 1$, a small Δt is often needed. In Ch. 6, we used a $\Delta t \approx 2$ nanoseconds. This can be cumbersome as you may wish to plot data with dynamics many orders of magnitude slower—in our case, microsecond time scales. However, computing the state probabilities from Monte Carlo data that is thousands of times finer grain than you need is cumbersome to say the least. I have found it advantageous to resample the time axis of the Monte Carlo data such that it matches the scale useful for a given plot. This can cut computation times down by orders of magnitude when later computing motional state probabilities and average motional state values.

B.2.1 Emission Pattern

Computing the emission pattern is in general nontrivial, but in the case of $^{171}\text{Yb}^+$, it ironically ends up being isotropic. There are two parts to the calculation. First, establish the

emission pattern based on the atomic transitions. Second, transform from the emission basis to the principal motional axes basis.

During resonant scattering in $^{171}\text{Yb}^+$, there are three equally likely spontaneous emission channels out of the $^2\text{P}_{1/2}$ state with probability of emitting into the $\hat{k}_s$ direction [53]

$$P_s(\hat{k}_s) = \frac{3}{8\pi} \sin^2 \theta_s \quad (\text{B.2})$$

for the $\Delta m_F = 0$ transition and

$$P_s(\hat{k}_s) = \frac{3}{16\pi} (1 + \cos^2 \theta_s) \quad (\text{B.3})$$

for the $\Delta m_F = \pm$ transitions where θ_s is the angle between the quantization axis set by a magnetic field $\vec{B}$ and the emission direction $\hat{k}_s$: $\vec{B} \cdot \hat{k}_s = \cos(\theta_s)$ [53].

Adding the three possible emission patterns together with equal probability leads to the isotropic case of $P_s(\hat{k}_s) = 1/4\pi$:

$$\begin{aligned} P_s(\hat{k}_s) &= \frac{1}{3} \left[\frac{3}{8\pi} \sin^2 \theta_s \right] + \frac{2}{3} \left[\frac{3}{16\pi} (1 + \cos^2 \theta_s) \right] \\ &= \frac{1}{8\pi} [\sin^2 \theta_s + 1 + \cos^2 \theta_s] \\ &= \frac{1}{4\pi} . \end{aligned} \quad (\text{B.4})$$

To compute the emission geometry factor along the x -axis f_{sx} , we now need to rewrite the emission vector in the basis of the principal motional axes. An isotropic emission vector, in the emission basis ($\{\hat{x}_s, \hat{y}_s, \hat{z}_s\}$) may be written as $\hat{k}_s = x_s \hat{x}_s + y_s \hat{y}_s + z_s \hat{z}_s$ where x_s , y_s , and z_s are random values but satisfy normalization conditions $|\hat{k}_s|^2 = 1$. For our trap, the two radial principal axes point up at 45° and 135° degrees from the optics table. Therefore, the emission basis and principal motion axes ($\{\hat{x}, \hat{y}, \hat{z}\}$) are related by $\hat{x}_s = (\hat{x} - \hat{y})/\sqrt{2}$, $\hat{y}_s = -\hat{z}$, and $\hat{z}_s = (\hat{x} + \hat{y})/\sqrt{2}$. The axial direction is irrelevant as the laser drives used in the study

of Ch. 6 only couple to the radial modes. The emission vector in the motional basis is then

$$\begin{aligned}\hat{k}_s &= x_s(\hat{x} - \hat{y})/\sqrt{2} - y_s\hat{z} + z_s(\hat{x} + \hat{y})/\sqrt{2} \\ \hat{k}_s &= \frac{x_s + z_s}{\sqrt{2}}\hat{x} - y_s\hat{z} + \frac{z_s - x_s}{\sqrt{2}}\hat{y} .\end{aligned}\tag{B.5}$$

In spherical coordinates $x_s = \sin \theta_s \cos \phi_s$, $y_s = \sin \theta_s \sin \phi_s$, and $z_s = \cos \theta_s$. Looking along just the x -axis as we do for the study in Ch. 6, $\hat{k}_{sx} = (\sin \theta_s \cos \phi_s + \cos \theta_s)/\sqrt{2}$.

Including this in the integral over the probability of emission as done in Ref. [53] Eq. (10b), we recover the geometry factor $f_{sx} = 1/3$, which is the expected value for isotropic emission. For a different spontaneous emission pattern, this may result in a more complicated result such as a dipole pattern with θ_s dependence.

In the QTT Monte Carlo, we randomly compute the emission direction for each scattering event. The directions are sampled from a uniform distribution $\{\hat{x}_s, \hat{y}_s, \hat{z}_s\} \sim U[-1, 1]$. The values are normalized such that $|\hat{k}_s|^2 = 1$. Note that randomly sampling θ_s and ϕ_s values would also work if you prefer spherical coordinates. Note that to get an isotropic emission pattern the probability distributions of θ_s and ϕ_s would not uniform.

B.3 Determination of Double Thermal Distribution

The double thermal distribution has been suggested as a more accurate description of an ion's motional state after ground state cooling [104, 98] and our work in Ch. 6 supports this as well. The distribution can be analytically estimated from measured data without a numerical fit as follows.

The double thermal distribution (defined in Eq. 5.2) is just that, two thermal distributions with a low average $\bar{n}_l$ and a high average $\bar{n}_h$:

$$\begin{aligned}p_{\text{double}}(n) &= \alpha p_{\text{th}}(n|\bar{n}_l) + (1 - \alpha) \\ p_{\text{double}}(n) &= \alpha \frac{\bar{n}_l^n}{(\bar{n}_l + 1)^{n+1}} + (1 - \alpha) \frac{\bar{n}_h^n}{(\bar{n}_h + 1)^{n+1}}\end{aligned}\tag{B.6}$$

with the average motional state being a weighted sum of the two $\bar{n}_{\text{double}} = \alpha\bar{n}_l + (1 - \alpha)\bar{n}_h$ ($0 \leq \alpha \leq 1$).

As shown in Ref. [104], only the first one or two motional states $|n = 0\rangle$ and $|n = 1\rangle$ constitute the $\bar{n}_l$ distribution as $\bar{n}_l \ll \bar{n}_h$. Typically, $\bar{n}_l \ll 1$ and $\bar{n}_h \gtrsim 10$. The accuracy of the analysis in Ch. 6 also suggests that only using the first two motional states to estimate the double thermal distribution is an accurate choice. (This is experimentally serendipitous as the low energy states are typically easier to measure accurately.) Consider, the measured state probabilities are $p_{\text{double}}(0)$ and $p_{\text{double}}(1)$. These would, therefore, entirely capture the contribution to Eq. (B.6) from $\bar{n}_l$. For example, consider the $|n = 0\rangle$ state probability which only depends on $\bar{n}_l$ to good approximation

$$\begin{aligned} p_{\text{double}}(0) &= \frac{\alpha}{\bar{n}_l + 1} + \frac{1 - \alpha}{\bar{n}_h + 1} \\ p_{\text{double}}(0) &\approx \frac{\alpha}{\bar{n}_l + 1} . \end{aligned} \tag{B.7}$$

The only calculation needed, given measured values $p_{\text{double}}(0)$ and $p_{\text{double}}(1)$, is to estimate the remaining population $p(n \geq 2)$ which is dominated by $\bar{n}_h$ as the $\bar{n}_l$ contributions are exponentially suppressed for higher states. We can compute $\bar{n}_h$ from a sum of thermal distribution probabilities

$$\begin{aligned} p(n \geq 2) &= (1 - \alpha) \sum_{n=2}^{\infty} \frac{\bar{n}_h^n}{(\bar{n}_h + 1)^{n+1}} \\ &= (1 - \alpha) \left(1 - \frac{1}{\bar{n}_h + 1} - \frac{\bar{n}_h}{(\bar{n}_h + 1)^2} \right) \\ &= (1 - \alpha) \frac{\bar{n}_h^2}{(\bar{n}_h + 1)^2} . \end{aligned} \tag{B.8}$$

In general, $p(n \geq m) = \bar{n}^m / (\bar{n} + 1)^m$. The value $p(n \geq 2)$ is known from measurement since $p(n \geq 2) = 1 - p_{\text{double}}(0) - p_{\text{double}}(1)$. Doing some algebra, the average motional state $\bar{n}_h$ is

then estimated as

$$\bar{n}_h = (1 - \alpha) \frac{p(n \geq m)^{1/m}}{1 - p(n \geq m)^{1/m}} . \quad (\text{B.9})$$

where $m = 2$ in this example calculation. The individual motional state probabilities for the higher energy states $|n \geq 2\rangle$ are then computed using this estimate of $\bar{n}_h$ (Eq. (B.9)) and ignoring $\bar{n}_l$ terms to good approximation: $p_{\text{double}}(n \geq 2) \approx (1 - \alpha)\bar{n}_h^n/(\bar{n}_h + 1)^{(n+1)}$. The α and $1 - \alpha$ terms can be absorbed into $\bar{n}_l$ and $\bar{n}_h$ as convenience dictates.

Bibliography

- [1] W. Paul, Reviews of modern physics **62**, 531 (1990).
- [2] NobelPrize.org, The Nobel Prize in Physics 1989, <https://www.nobelprize.org/prizes/physics/1989/summary/>, accessed: 2024-03-26.
- [3] Trapped the ion, <https://www.washington.edu/innovation/trapping-the-ion/>, accessed: 2024-03-26.
- [4] H. Dehmelt, Am. J. Phys **58**, 1 (1990).
- [5] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, Journal of Research of the National Institute of Standards and Technology **103**, 259 (1998).
- [6] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, Reviews of Modern Physics **75**, 281 (2003).
- [7] C. Monroe *et al.*, Reviews of Modern Physics **93**, 025001 (2021).
- [8] D. J. Wineland and W. M. Itano, Physical Review A **20**, 1521 (1979).
- [9] NobelPrize.org, The Nobel Prize in Physics 2012, <https://www.nobelprize.org/prizes/physics/2012/summary/>, accessed: 2024-03-26.
- [10] D. J. Wineland, R. E. Drullinger, and F. L. Walls, Physical Review Letters **40**, 1639 (1978).
- [11] R. P. Feynman, Int. J. Theor. Phys. **21**, 467– (1982).
- [12] D. Deutsch, Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences **400**, 97 (1985).
- [13] S. Lloyd, Science **273**, 1073 (1996).

- [14] P. W. Shor, Proceedings 35th annual symposium on foundations of computer science 124 (1994).
- [15] J. I. Cirac and P. Zoller, Physical review letters **74**, 4091 (1995).
- [16] Q. Turchette, C. Wood, B. King, C. Myatt, D. Leibfried, W. Itano, C. Monroe, and D. Wineland, Physical Review Letters **81**, 3631 (1998).
- [17] C. A. Sackett *et al.*, Nature **404**, 256 (2000).
- [18] J. P. Gaebler *et al.*, Physical Review Letters **117**, 060505 (2016).
- [19] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, Phys. Rev. Lett. **117**, 060504 (2016).
- [20] C. R. Clark *et al.*, arXiv preprint arXiv:2105.05828 (2021).
- [21] R. Srinivas *et al.*, Nature **597**, 209 (2021).
- [22] J. Pino *et al.*, Nature **592**, 209 (2021).
- [23] S. Moses *et al.*, arXiv preprint arXiv:2305.03828 (2023).
- [24] J. E. Christensen, D. Hucul, W. C. Campbell, and E. R. Hudson, npj Quantum Information **6**, 35 (2020).
- [25] T. Harty, D. Allcock, C. J. Ballance, L. Guidoni, H. Janacek, N. Linke, D. Stacey, and D. Lucas, Physical review letters **113**, 220501 (2014).
- [26] A. Ransford, C. Roman, T. Dellaert, P. McMillin, and W. C. Campbell, Physical Review A **104**, L060402 (2021).
- [27] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J.-N. Zhang, L.-M. Duan, D. Yum, and K. Kim, Nature Photonics **11**, 646 (2017).

- [28] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, *Nature Communications* **12**, 1 (2021).
- [29] L. Egan *et al.*, *Nature* **598**, 281 (2021).
- [30] C. Ryan-Anderson *et al.*, *Physical Review X* **11**, 041058 (2021).
- [31] L. Postler, F. Butt, I. Pogorelov, C. D. Marciniak, S. Heußen, R. Blatt, P. Schindler, M. Rispler, M. Müller, and T. Monz, *arXiv preprint arXiv:2312.09745* (2023).
- [32] R. J. Niffenegger *et al.*, *Nature* **586**, 538 (2020).
- [33] M. C. Reville, *arXiv preprint arXiv:2009.02398* (2020).
- [34] B. Hampel, D. H. Slichter, D. Leibfried, R. P. Mirin, S. W. Nam, and V. B. Verma, *Applied physics letters* **122**, 174001 (2023).
- [35] C. Mordini, A. R. Vasquez, Y. Motohashi, M. Müller, M. Malinowski, C. Zhang, K. K. Mehta, D. Kienzler, and J. P. Home, *arXiv preprint arXiv:2401.18056* (2024).
- [36] R. D. Delaney *et al.*, *arXiv preprint arXiv:2403.00756* (2024).
- [37] J. Sterk *et al.*, *arXiv preprint arXiv:2403.00208* (2024).
- [38] D. Matsukevich, P. Maunz, D. L. Moehring, S. Olmschenk, and C. Monroe, *Physical Review Letters* **100**, 150404 (2008).
- [39] I. V. Inlek, C. Crocker, M. Lichtman, K. Sosnova, and C. Monroe, *Physical review letters* **118**, 250502 (2017).
- [40] V. Krutyanskiy *et al.*, *Physical Review Letters* **130**, 050803 (2023).
- [41] R. Sutherland, R. Srinivas, S. C. Burd, D. Leibfried, A. C. Wilson, D. J. Wineland, D. Allcock, D. Slichter, and S. Libby, *New Journal of Physics* **21**, 033033 (2019).

- [42] D. Hucul, J. E. Christensen, E. R. Hudson, and W. C. Campbell, Physical review letters **119**, 100501 (2017).
- [43] S. Olmschenk, Physical Review A **105**, 042617 (2022).
- [44] D. Allcock, W. Campbell, J. Chiaverini, I. Chuang, E. Hudson, I. Moore, A. Ransford, C. Roman, J. Sage, and D. Wineland, Applied Physics Letters **119**, 214002 (2021).
- [45] K. Mølmer and A. Sørensen, Physical Review Letters **82**, 1835 (1999).
- [46] A. Sørensen and K. Mølmer, Physical review letters **82**, 1971 (1999).
- [47] A. Sørensen and K. Mølmer, Physical Review A **62**, 022311 (2000).
- [48] P. D. Schwindt *et al.*, 2015 Joint Conference of the IEEE International Frequency Control Symposium & the European Frequency and Time Forum 752 (2015).
- [49] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, Reviews of Modern Physics **87**, 637 (2015).
- [50] T. Bothwell, C. J. Kennedy, A. Aepli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, Nature **602**, 420 (2022).
- [51] J. H. Gross, *Mass spectrometry: a textbook* (Springer Science & Business Media, ADDRESS, 2006).
- [52] P. Richerme, Physical Review A **94**, 032320 (2016).
- [53] W. M. Itano and D. Wineland, Physical Review A **25**, 35 (1982).
- [54] M. D’Onofrio, Y. Xie, A. J. Rasmusson, E. Wolanski, J. Cui, and P. Richerme, Physical Review Letters **127**, 020503 (2021).
- [55] Y. Xie, J. Cui, M. D’Onofrio, A. J. Rasmusson, S. W. Howell, and P. Richerme, Quantum Science and Technology **6**, 044009 (2021).

- [56] D. Berkeland, J. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland, *Journal of applied physics* **83**, 5025 (1998).
- [57] D. Leibfried, *Physical Review A* **60**, R3335 (1999).
- [58] J. Gaebler, C. Baldwin, S. Moses, J. Dreiling, C. Figgatt, M. Foss-Feig, D. Hayes, and J. Pino, *Physical Review A* **104**, 062440 (2021).
- [59] N. K. Lysne, J. F. Niedermeyer, A. C. Wilson, D. H. Slichter, and D. Leibfried, *arXiv preprint arXiv:2402.05857* (2024).
- [60] J. Keller, H. L. Partner, T. Burgermeister, and T. Mehlstäubler, *Journal of Applied Physics* **118**, 104501 (2015).
- [61] C. J. Goham and J. W. Britton, *AIP Advances* **12**, 115315 (2022).
- [62] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland, *Physical Review Letters* **76**, 1796 (1996).
- [63] C. Cohen-Tannoudji, B. Diu, and F. Laloe, *Quantum Mechanics* **1**, 898 (1986).
- [64] C. H. Valahu, I. Apostolatos, S. Weidt, and W. K. Hensinger, *Journal of Physics B: Atomic, Molecular and Optical Physics* **55**, 204003 (2022).
- [65] A. Kyprianidis, A. Rasmusson, and P. Richerme, *New Journal of Physics* (2024).
- [66] D. F. James, Technical report, Los Alamos National Laboratory (unpublished).
- [67] C. Flühmann, T. L. Nguyen, M. Marinelli, V. Negnevitsky, K. Mehta, and J. Home, *Nature* **566**, 513 (2019).
- [68] B. De Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, *Nature Physics* **18**, 296 (2022).
- [69] I. Rojkov, P. M. Röggl, M. Wagener, M. Fontboté-Schmidt, S. Welte, J. Home, and F. Reiter, *arXiv preprint arXiv:2305.05262* (2023).

- [70] S. Burd, R. Srinivas, J. Bollinger, A. Wilson, D. Wineland, D. Leibfried, D. Slichter, and D. Allcock, *Science* **364**, 1163 (2019).
- [71] M. Affolter, W. Ge, B. Bullock, S. C. Burd, K. A. Gilmore, J. F. Lilieholm, A. L. Carter, and J. J. Bollinger, *Physical Review A* **107**, 032425 (2023).
- [72] S. C. Burd, R. Srinivas, H. M. Knaack, W. Ge, A. C. Wilson, D. J. Wineland, D. Leibfried, J. J. Bollinger, D. Allcock, and D. Slichter, *Nature Physics* **17**, 898 (2021).
- [73] S. Burd *et al.*, arXiv preprint arXiv:2304.05529 (2023).
- [74] B. M. White, P. J. Low, Y. de Sereville, M. L. Day, N. Greenberg, R. Rademacher, and C. Senko, *Physical Review A* **105**, 033102 (2022).
- [75] N. Greenberg, B. M. White, P. J. Low, and C. Senko, arXiv preprint arXiv:2307.07627 (2023).
- [76] X. Shi, S. Todaro, G. Mintzer, C. Bruzewicz, J. Chiaverini, and I. Chuang, *Applied Physics Letters* **122**, 264002 (2023).
- [77] S. Olmschenk and P. Becker, *Applied Physics B* **123**, 1 (2017).
- [78] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe, *Physical Review A* **76**, 052314 (2007).
- [79] P. J. Low, B. White, and C. Senko, arXiv preprint arXiv:2306.03340 (2023).
- [80] P. Maunz, D. Moehring, S. Olmschenk, K. C. Younge, D. Matsukevich, and C. Monroe, *Nature Physics* **3**, 538 (2007).
- [81] S. Ejtemaee, R. Thomas, and P. Haljan, *Physical Review A* **82**, 063419 (2010).
- [82] D. Berkeland and M. Boshier, *Physical Review A* **65**, 033413 (2002).

- [83] A. Myerson, D. Szwer, S. Webster, D. Allcock, M. Curtis, G. Imreh, J. Sherman, D. Stacey, A. Steane, and D. Lucas, Physical Review Letters **100**, 200502 (2008).
- [84] A. Lee, Ph.D. thesis, University of Maryland, College Park, 2016.
- [85] NobelPrize.org, 1997 Nobel Prize in Physics Press release, <https://www.nobelprize.org/prizes/physics/1997/press-release>, accessed: 2024-03-26.
- [86] S. Stenholm, Reviews of Modern Physics **58**, 699 (1986).
- [87] A. Rasmusson, K. J. Ferris, N. T. Bronn, and O. Lanes, arXiv preprint arXiv:2209.02795 (2022).
- [88] D. A. Steck, (2007).
- [89] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Nature **536**, 63 (2016).
- [90] Y. Wang, S. Crain, C. Fang, B. Zhang, S. Huang, Q. Liang, P. H. Leung, K. R. Brown, and J. Kim, Physical Review Letters **125**, 150505 (2020).
- [91] A. Binai-Motlagh, M. L. Day, N. Videnov, N. Greenberg, C. Senko, and R. Islam, Quantum Science and Technology **8**, 045012 (2023).
- [92] S. Lim, S. Baek, J. Withlow, M. D’Onofrio, T. Chen, S. Phiri, S. Crain, K. R. Brown, J. Kim, and J. Kim, arXiv preprint arXiv:2402.13560 (2024).
- [93] K. R. Brown, A. W. Harrow, and I. L. Chuang, Physical Review A **70**, 052318 (2004).
- [94] W. C. Campbell, J. Mizrahi, Q. Quraishi, C. Senko, D. Hayes, D. Hucul, D. N. Matsukevich, P. Maunz, and C. Monroe, Physical Review Letters **105**, 090502 (2010).
- [95] I. Moore, W. Campbell, E. Hudson, M. Boguslawski, D. Wineland, and D. Allcock, Physical Review A **107**, 032413 (2023).

- [96] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, ADDRESS, 2010).
- [97] K. R. Islam, Ph.D. thesis, University of Maryland, College Park, 2012.
- [98] A. J. Rasmusson, M. D’Onofrio, Y. Xie, J. Cui, and P. Richerme, Physical Review A **104**, 043108 (2021).
- [99] L. Feng, W. Tan, A. De, A. Menon, A. Chu, G. Pagano, and C. Monroe, Physical Review Letters **125**, 053001 (2020).
- [100] E. C. Reed, L. Qi, and K. R. Brown, arXiv preprint arXiv:2403.04891 (2024).
- [101] G. Morigi, J. Eschner, and C. H. Keitel, Physical review letters **85**, 4458 (2000).
- [102] C. F. Roos, D. Leibfried, A. Mundt, F. Schmidt-Kaler, J. Eschner, and R. Blatt, Physical Review Letters **85**, 5547 (2000).
- [103] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, Physical Review Letters **62**, 403 (1989).
- [104] J.-S. Chen, S. M. Brewer, C. W. Chou, D. J. Wineland, D. R. Leibbrandt, and D. B. Hume, Physical Review Letters **118**, 053002 (2017).
- [105] M. D. D’Onofrio, Ph.D. thesis, Indiana University, 2022.
- [106] C. Senko, Ph.D. thesis, University of Maryland, College Park, 2014.
- [107] O. G  hne and G. T  th, Physics Reports **474**, 1 (2009).
- [108] Y. Xie, Ph.D. thesis, Indiana University, 2022.
- [109] J. D. Sivers, L. R. Simkins, S. Weidt, and W. K. Hensinger, Applied Physics B **107**, 921 (2012).

- [110] K. G. Johnson, J. D. Wong-Campos, A. Restelli, K. A. Landsman, B. Neyenhuis, J. Mizrahi, and C. Monroe, Review of Scientific Instruments **87**, 053110 (2016).
- [111] ARTIQ documentation, M-Labs Ltd, <https://m-labs.hk/artiq/manual/index.html>.
- [112] A. Franzen, ComponentLibrary, 2006, this work is licensed under the Creative Commons Attribution-NonCommercial 3.0 Unported License. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc/3.0/>.
- [113] S. Ejtemaee and P. C. Haljan, Physical Review Letters **119**, 043001 (2017).
- [114] C. Baldwin, B. Bjork, M. Foss-Feig, J. Gaebler, D. Hayes, M. Kokish, C. Langer, J. Sedlacek, D. Stack, and G. Vittorini, Physical Review A **103**, 012603 (2021).
- [115] S.-A. Guo *et al.*, arXiv preprint arXiv:2311.17163 (2023).
- [116] R. Islam *et al.*, Optics letters **39**, 3238 (2014).
- [117] E. Mount *et al.*, Quantum Information Processing **15**, 5281 (2016).
- [118] M. Hirsch, R. J. Wareham, M. L. Martin-Fernandez, M. P. Hobson, and D. J. Rolfe, PloS one **8**, e53671 (2013).
- [119] N. Schwegler, Ph.D. thesis, ETH Zurich, 2018.
- [120] S. Bourdeauducq, whitequark, R. Jördens, D. Nadlinger, Y. Sionneau, and F. Kermarrec, ARTIQ, 2022.
- [121] ARTIQ: A leading-edge control system for quantum information experiments, Github repository, <https://github.com/m-labs/artiq>.
- [122] Sinara: Wiki home, Github repository, <https://github.com/sinara-hw/meta/wiki>.
- [123] C. F. Roos, New Journal of Physics **10**, 013002 (2008).

- [124] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt, *Nature Physics* **4**, 463 (2008).
- [125] M. Weber, M. Gely, R. Hanley, T. Harty, A. Leu, C. Löschnauer, D. Nadlinger, and D. Lucas, arXiv preprint arXiv:2402.12955 (2024).
- [126] A. S. Dalvi, J. Whitlow, M. D’Onofrio, L. Rieseboos, T. Chen, S. Phiri, K. R. Brown, and J. M. Baker, arXiv preprint arXiv:2308.10787 (2023).
- [127] M-labs forum, M-Labs Ltd, <https://forum.m-labs.hk/>.
- [128] L. Rieseboos, B. Bondurant, and K. R. Brown, *IEEE Micro* **41**, 57 (2021).
- [129] P. Kent and D. Slichter, The NIST Scan Framework for ARTIQ, https://pages.nist.gov/artiq_scan_framework/.
- [130] D. Nadlinger *et al.*, ndscan – N-dimensional scans for ARTIQ, <https://github.com/OxfordIonTrapGroup/ndscan>.
- [131] L. Rieseboos *et al.*, Duke ARTIQ Extensions (DAX), <https://gitlab.com/duke-artiq/dax>, 2019.
- [132] Radon, Github repository, <https://github.com/rubik/radon>.
- [133] G. van Rossum *et al.*, Style Guide for Python Code (pep8), <https://www.python.org/dev/peps/pep-0008/>, 2001.
- [134] G. van Rossum *et al.*, Style Guide for Python Code (pep257), <https://peps.python.org/pep-0257/>, 2001.
- [135] N. Krackow, Short Time Fourier Transform Pulse Generator for Trapped Ion Quantum Gates, 2020.
- [136] R. Alheit, T. Gudjons, S. Kleineidam, and G. Werth, *Rapid Communications in Mass Spectrometry* **10**, 583 (1996).

- [137] J. M. Wilson, J. N. Tilles, R. A. Haltli, E. Ou, M. G. Blain, S. M. Clark, and M. C. Revelle, *Journal of Applied Physics* **131**, 134401 (2022).
- [138] N. Huntemann, C. Sanner, B. Lipphardt, C. Tamm, and E. Peik, *Physical Review Letters* **116**, 063001 (2016).
- [139] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C.-w. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibbrandt, *Physical Review Letters* **123**, 033201 (2019).
- [140] K. Wright *et al.*, *Nature Communications* **10**, 5464 (2019).
- [141] R. Blatt and C. F. Roos, *Nature Physics* **8**, 277 (2012).
- [142] M. J. Biercuk, H. Uys, J. W. Britton, A. P. VanDevender, and J. J. Bollinger, *Nature Nanotechnology* **5**, 646 (2010).
- [143] C. L. Degen, F. Reinhard, and P. Cappellaro, *Reviews of Modern Physics* **89**, 035002 (2017).
- [144] J. D. Teufel, T. Donner, D. Li, J. W. Harlow, M. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds, *Nature* **475**, 359 (2011).
- [145] C. Whittle *et al.*, *Science* **372**, 1333 (2021).
- [146] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (Springer-Verlag, New York, NY, 1999).
- [147] J. Dalibard and C. Cohen-Tannoudji, *J. Opt. Soc. Am. B* **6**, 2023 (1989).
- [148] M. Kasevich and S. Chu, *Physical Review Letters* **69**, 1741 (1992).
- [149] I. Marzoli, J. I. Cirac, R. Blatt, and P. Zoller, *Physical Review A* **49**, 2771 (1994).
- [150] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, *Physical Review Letters* **75**, 4011 (1995).

- [151] B. E. King, C. S. Wood, C. J. Myatt, Q. A. Turchette, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland, Physical Review Letters **81**, 1525 (1998).
- [152] J.-S. Chen, K. Wright, N. C. Pisenti, D. Murphy, K. M. Beck, K. Landsman, J. M. Amini, and Y. Nam, Phys. Rev. A **102**, 043110 (2020).
- [153] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt, J. Opt. Soc. Am. B **20**, 1003 (2003).
- [154] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt, Physical Review Letters **41**, 233 (1978).
- [155] D. J. Wineland and W. M. Itano, Physics Today **40**, 34 (1987).
- [156] H. Che, K. Deng, Z. T. Xu, W. H. Yuan, J. Zhang, and Z. H. Lu, Phys. Rev. A **96**, 013417 (2017).
- [157] C. Roos, T. Zeiger, H. Rohde, H. Nägerl, J. Eschner, D. Leibfried, F. Schmidt-Kaler, and R. Blatt, Physical Review Letters **83**, 4713 (1999).
- [158] C. Hempel, Ph.D. thesis, Universität Innsbruck, Ph.D. Dissertation, 2014.
- [159] L. Deslauriers, P. C. Haljan, P. J. Lee, K. A. Brickman, B. B. Blinov, M. J. Madsen, and C. Monroe, Physical Review A **70**, 043408 (2004).
- [160] Y. Wan, F. Gebert, F. Wolf, and P. O. Schmidt, Physical Review A **91**, 043425 (2015).
- [161] C. R. Harris *et al.*, Nature **585**, 357 (2020).
- [162] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, Reviews of Modern Physics **87**, 1419 (2015).
- [163] F. Haddadfarshi and F. Mintert, New Journal of Physics **18**, 123007 (2016).

- [164] A. R. Milne, C. L. Edmunds, C. Hempel, F. Roy, S. Mavadia, and M. J. Biercuk, *Physical Review Applied* **13**, 024022 (2020).
- [165] C. D. Bentley, H. Ball, M. J. Biercuk, A. R. Carvalho, M. R. Hush, and H. J. Slatyer, *Advanced Quantum Technologies* **3**, 2000044 (2020).
- [166] M. Cetina, L. Egan, C. Noel, M. Goldman, D. Biswas, A. Risinger, D. Zhu, and C. Monroe, *PRX Quantum* **3**, 010334 (2022).
- [167] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe, *Phys. Rev. Lett.* **97**, 103007 (2006).
- [168] D. An, C. Matthiesen, E. Urban, and H. Häffner, *Physical Review A* **100**, 063405 (2019).
- [169] D. Allcock, L. Guidoni, T. Harty, C. Ballance, M. Blain, A. Steane, and D. Lucas, *New Journal of Physics* **13**, 123023 (2011).
- [170] D. A. Hite *et al.*, *Physical Review Letters* **109**, 103001 (2012).
- [171] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang, *Physical Review Letters* **100**, 013001 (2008).
- [172] G. Pagano *et al.*, *Quantum Science and Technology* **4**, 014004 (2018).
- [173] M. Weber, C. Löschnauer, J. Wolf, M. Gely, R. Hanley, J. Goodwin, C. Ballance, T. Harty, and D. Lucas, *Quantum Science and Technology* **9**, 015007 (2023).
- [174] B. M. Terhal, *Reviews of Modern Physics* **87**, 307 (2015).
- [175] R. Raussendorf, D. E. Browne, and H. J. Briegel, *Physical Review A* **68**, 022312 (2003).
- [176] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest, *Nature Physics* **5**, 19 (2009).

- [177] M. Iqbal *et al.*, arXiv preprint arXiv:2302.01917 (2023).
- [178] M. Iqbal *et al.*, Nature **626**, 505 (2024).
- [179] B. Skinner, J. Ruhman, and A. Nahum, Physical Review X **9**, 031009 (2019).
- [180] M. Ippoliti, M. J. Gullans, S. Gopalakrishnan, D. A. Huse, and V. Khemani, Physical Review X **11**, 011030 (2021).
- [181] M. Foss-Feig *et al.*, arXiv preprint arXiv:2302.03029 (2023).
- [182] E. Chertkov *et al.*, Nature Physics **18**, 1074 (2022).
- [183] M. DeCross, E. Chertkov, M. Kohagen, and M. Foss-Feig, Physical Review X **13**, 041057 (2023).
- [184] D. A. Hite, K. S. McKay, S. Kotler, D. Leibfried, D. J. Wineland, and D. P. Pappas, MRS Advances **2**, 2189 (2017).
- [185] S. K. Lamoreaux, Physical Review A **56**, 4970 (1997).
- [186] Q. Turchette, C. Myatt, B. King, C. Sackett, D. Kielpinski, W. Itano, C. Monroe, and D. Wineland, Physical Review A **62**, 053807 (2000).
- [187] H. J. Carmichael, *Statistical methods in quantum optics 2: Non-classical fields* (Springer Science & Business Media, ADDRESS, 2007).
- [188] L. Horvath, M. Collett, H. Carmichael, and R. Fisher, in *Quantum Stochastic Heating of a Trapped Ion*, Optica Publishing Group (Conference on Coherence and Quantum Optics, ADDRESS, 2007), p. CTuE2.
- [189] D. F. James, Physical Review Letters **81**, 317 (1998).
- [190] F. De Oliveira, M. S. Kim, P. L. Knight, and V. Buek, Physical Review A **41**, 2645 (1990).

- [191] S. D. Fallek, V. S. Sandhu, R. A. McGill, J. M. Gray, H. N. Tinkey, C. R. Clark, and K. R. Brown, *Nature Communications* **15**, 1089 (2024).
- [192] M. Fabrikant, P. Lauria, I. Madjarov, W. Burton, and R. Sutherland, arXiv preprint arXiv:2403.02315 (2024).
- [193] D. Kielpinski, C. Monroe, and D. J. Wineland, *Nature* **417**, 709 (2002).
- [194] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, *Science* **309**, 749 (2005).