

UCLA

UCLA Electronic Theses and Dissertations

Title

Tips and Tricks for Building a Quantum Computer with $^{133}\mathrm{Ba}^{+}$

Permalink

<https://escholarship.org/uc/item/4072m646>

ISBN

9798190947496

Author

Wall, Zachary James

Publication Date

2026-09-11

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Tips and Tricks for Building a Quantum Computer with $^{133}\text{Ba}^+$

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Physics

by

Zachary James Wall

ABSTRACT OF THE DISSERTATION

Tips and Tricks for Building a Quantum Computer with $^{133}\text{Ba}^+$

by

Zachary James Wall

Doctor of Philosophy in Physics

University of California, Los Angeles, 2026

Professor Eric Hudson, Co-Chair

Professor Wesley Campbell, Co-Chair

Building a fault-tolerant quantum computer requires more than improving physical hardware in isolation. Hardware-level design choices must be guided by the architecture of the full quantum computer and by the bottlenecks that ultimately limit its performance. The $^{133}\text{Ba}^+$ ion is well suited for this kind of architecture-driven optimization. Its simple nuclear spin structure, with $I = 1/2$, enables fast, high-fidelity state preparation and measurement (SPAM). We leverage this advantage to demonstrate record SPAM fidelity, along with a state-preparation time more than three orders of magnitude faster than previous record-level SPAM demonstrations. This makes $^{133}\text{Ba}^+$ especially promising as a reusable ancilla qubit for quantum error correction.

The long-lived $D_{5/2}$ state, with $\tau = 30$ s, also makes $^{133}\text{Ba}^+$ uniquely positioned to take advantage of the optical-metastable-ground (OMG) qubit protocol, which enables same-species sympathetic cooling. Within this protocol, we demonstrate a magic polarization condition that allows global gates with 532 nm light to be applied in one qubit subspace while leaving other subspaces largely unaffected. We also demonstrate stimulated Raman gates in both the ground and metastable qubit subspaces with errors below the fault-tolerance threshold, supported by a direct measurement of the spontaneous Raman scattering rate. The main practical drawback of $^{133}\text{Ba}^+$ is its radioactivity. Here, we mitigate this challenge by demonstrating ion loading from a ^{133}Ba target below the Nuclear

Regulatory Commission exempt quantity of 10 μCi . This is enabled by a new autoionization loading scheme that substantially reduces loading times for barium ions.

The dissertation of Zachary James Wall is approved.

Leo Zhou

Paul Hamilton

Wesley Campbell, Committee Co-Chair

Eric Hudson, Committee Co-Chair

University of California, Los Angeles

2026

To my family and Jacqlyn for love and support.

TABLE OF CONTENTS

1	Introduction	1
1.1	Are Classical Computers Stuck?	3
1.1.1	A Brief History of Compute	3
1.1.2	Quantum Chemistry Scaling	6
1.2	What is a Quantum Computer?	10
1.2.1	The Quantum Phase Estimation Algorithm	11
1.2.2	Where Does the Speed Up Come From?	13
1.3	How To Build a Quantum Computer	16
1.3.1	Is It Possible to Build a Quantum Computer?	16
1.3.2	Where Are We Now?	19
2	The Ion Trap	22
2.1	Introduction	22
2.2	Solving the Equations of Motion	25
2.3	Determining Parameters of Stable Motion	32
2.4	Numerical Simulation of Real Trap Geometries	36
2.5	Excess Micromotion	41
2.6	Ion Heating Rates	48
3	Atomic Structure and Atom Laser Interactions	54
3.1	Barium Atomic Structure	55
3.1.1	The Spin-Orbit Interaction	58
3.1.2	The Hyperfine Interaction	60

3.2	The Zeeman Shift	62
3.2.1	First Order Magnetic Field Insensitive Transitions	65
3.3	The Light Shift	67
3.4	Doppler Cooling	69
3.5	Dipole Matrix Element Calculations	74
4	The Experiment	82
4.1	First Generation $^{133}\text{Ba}^+$ Trap	82
4.2	Second Generation $^{133}\text{Ba}^+$ Trap	85
4.3	Trap Electronics	88
4.4	Vacuum Chamber	94
4.5	Ion Sources and Loading Procedures	96
4.6	Imaging	102
4.7	Laser Systems	103
4.7.1	Cooling and Readout Lasers	105
4.7.2	Raman Laser System	106
4.7.3	1762 Stable Laser Systems Cavity	107
4.8	Doppler Cooling Cavities	109
4.9	Microwave Drive Chain	111
4.10	Experimental Control Software	112
5	The Limits of Spontaneous Raman Scattering Errors	114
5.1	Raman Gates	114
5.2	Single Qubit Gates	116
5.3	Spontaneous Raman Scattering Error	124

5.4	Experimental Verification of the Raman Scattering Rate in $^{133}\text{Ba}^+$	134
6	Magic Angle Raman Gates in Magnetic Clock Qubits	152
6.1	The AC Light Shift of a Clock Qubit	152
6.2	The <i>omg</i> Protocol	159
6.3	Experimental Measurements of Magic Angle Gates	160
6.4	Calculation of the Critical Magnetic Field	173
6.5	Experimental Fitting Details	176
7	Efficient loading of barium ions using autoionizing resonances	178
7.1	Experimental Procedure	181
7.2	Spectroscopy From the $5d6p\ ^3D_1^o$ State	182
7.3	530.8 Transition	184
7.4	Comparison with Previous Photoionization Schemes	184
7.5	Doppler Broadening	189
7.6	Calculation of Saturation Intensity	190
8	State Preparation and Measurement in $^{133}\text{Ba}^+$	193
8.1	What is SPAM?	193
8.2	State Preparation	195
8.3	Single Qubit Gates	203
8.4	Electron Shelving	207
8.5	State Detection	210
9	Future Experiments: The Next Generation	216
9.1	Where is the Bottleneck to Building a Fault-Tolerant Quantum Computer Now?	216

9.1.1	Ion Loading, Shuttling, and Recooling	216
9.1.2	Parallel Operations and Integrated Laser Delivery	220
9.1.3	QEC Syndrome Extraction	224
9.2	SPAM in the Age of Quantum Error Correction	229
9.3	Metastable Same-Species Sympathetic Cooling	233
9.4	Autler–Townes Shielding of Readout Light	243
9.5	Microwave-Dressed Magnetic-Field-Insensitive Qubits in Barium	252
References		264

LIST OF FIGURES

1.1	Algorithmic Scaling	6
2.1	Ion Fourier Spectrum	32
2.2	Mathieu Stability Diagram	35
2.3	Light Collection	36
2.4	Numerical Simulation of Laplace's Equation for a Blade Trap	38
2.5	Parametric Modulation Micromotion	44
2.6	Micromotion Heat Map	47
2.7	Electrical Noise Spectrum	52
3.1	Radial wavefunctions of Ba^+	57
3.2	Ba^+ level structure	59
3.3	Zeeman Shifts in the $D_{5/2}$ of $^{137}\text{Ba}^+$	63
3.4	Substructure with Zeeman Splitting	64
3.5	Clock states in $^{133}\text{Ba}^+$	65
4.1	Ion Cloud	83
4.2	Diagram of RF blades	86
4.3	Trap Wiring	87
4.4	RF drive chain to generate the trap signal.	89
4.5	Trap Inner View	92
4.6	Aluminum pieces on the front and bottom of the can machined for mechanical stabilization.	93
4.7	Bakeout Data	96

4.8	CAD rendering of the vacuum chamber with sliced view.	97
4.9	Views of the six-shooter target assembly and its position relative to the trap. .	98
4.10	View through the chamber window of the two pure radioactive targets.	101
4.11	Zeemax of the light collection system for the .6 NA objective.	103
4.12	Full diagram of light delivery to the blade trap.	104
4.13	Laser optical table diagram and AOM configuration.	105
4.14	Raman laser diagram for driving motionally dependent transition in $^{133}\text{Ba}^+$. .	107
4.15	Diagram of the internal components of the SLS 1762 nm cavity stabilization system.	108
4.16	Full schematic of the current microwave drive chain to generate 9.925 GHz. . .	111
5.1	Energy level diagram of stimulated Raman transition.	119
5.2	Raman Rabi Rate	123
5.3	Energy Level Diagram For Raman Transitions	136
5.4	$^{133}\text{Ba}^+$ level structure	137
5.5	Pulse Sequence	139
5.6	Experimental pulse sequence to measure the total Raman scattering rate. . . .	141
5.7	Stimulated Raman Scattering Rate	142
5.8	SRS Branching Fraction	144
5.9	Two-qubit Gate Error	147
6.1	OMG Architecture	162
6.2	Experimental Sequence for Metastable Gates	164
6.3	Raman Spectroscopy	165
6.4	Measured m -qubit frequency splitting	166
6.5	Microwave Ramsey	167

6.6	Magic Polarization Conditions	169
7.1	Level Diagram for Photoionization	179
7.2	Spectrum of autoionizing resonances originating from the $5d6p\ ^3D_1^o$ state . . .	182
7.3	Autoionization resonances to the $5d_{3/2}15d_{3/2}\ J = 2$ state	183
7.4	Spectrum of the autoionizing resonances near 530.8 nm	185
7.5	Average number of ions loaded per ablation shot as a function of 530.8 nm laser power	186
7.6	Comparison of relevant photoionization cross sections for barium loading . . .	187
7.7	frequency scan of the 413 nm transition	190
8.1	Level diagram of the typical optical pumping scheme in $^{137}\text{Ba}^+$ using σ^+ polarization.	196
8.2	Level diagram of the typical optical pumping scheme in $^{133}\text{Ba}^+$ using mixed polarization.	197
8.3	7 level system simulation to model the effective pumping rate and fidelity of $^{133}\text{Ba}^+$ optical pumping with mixed polarization light at a 5 G magnetic bias field.	200
8.4	Comparison of different intensities to ultimate optical pumping fidelity. . . .	201
8.5	EOM sideband power distribution at $\beta = 1.84$. This parameter quantifies the RF power input strength.	202
8.6	Comparison of different intensities to ultimate optical pumping fidelity, while using an EOM to drive the $ F = 1 \rightarrow F = 1\rangle$ transition.	203
8.7	Diagram of the calibration sequence using a Ramsey experiment.	204
8.8	Single qubit gate benchmarking. A sequence of π pulses with a random phases are applied to extract the single pulse gate error.	205

8.9	Dissipative preparation of the $ 1\rangle$ state after the application of a microwave π pulse.	206
8.10	The temperature dependence of the cabinet shelving method is compared for two different starting thermal distributions.	208
8.11	Simulation of the rate of population shelving from the $F = 1$ state versus the rate population accidentally falls into the $F = 0$ causing an error.	209
8.12	Off-resonant shelving rate	210
8.13	Photon count distribution	211
8.14	Detection time optimization	213
8.15	Experimental SPAM measurement	214
9.1	Proposed metastable same-species sympathetic-cooling cycle and the present optical-reset simulation. The 532 nm Raman red-sideband pulse removes one phonon from a shared motional mode while transferring the coolant within the $D_{5/2}$ manifold. A 614 nm quench, spontaneous emission, recycling of the D manifolds, and a selective 1762 nm pulse return the coolant to its initial state. The neighboring data ion remains in the $S_{1/2}$ clock manifold throughout the cycle. The right panel shows the simulated final ground-state distribution for equal σ^+ , π , and σ^- components in the quench field.	236
9.2	Concept for spectroscopic shielding of a spectator ion during fluorescence readout. A measured ancilla emits approximately 10^3 photons during a $100\ \mu\text{s}$ readout, giving a high-end reabsorption estimate of approximately 31% for a spectator $5\ \mu\text{m}$ away. A strong 650 nm control field splits the spectator's 493 nm absorption resonance and reduces its overlap with the fluorescence spectrum. For a $^{133}\text{Ba}^+$ ground-state clock qubit at 5 G, the calculated differential Stark shift can be cancelled near a circular-polarization parameter of $ \mathcal{A} \simeq 0.29$. The spectra in the center panel are schematic and are not plotted from the numerical optical-Bloch calculation.	246

9.3	Microwave dressing of magnetic-field-sensitive clock states in barium. The left panel shows the $^{137}\text{Ba}^+$ ground-state topology used in the exploratory calculation. An off-resonant microwave field couples the qubit states to neighboring Zeeman levels through several detunings. The center panel illustrates the desired result. The undressed qubit has a first-order-insensitive point near 2.35 G but retains quadratic curvature. Proper dressing can, in principle, create a broader plateau and make the operating frequency insensitive to both magnetic-field and microwave-amplitude fluctuations. The plotted frequency curves are schematic and are not fitted to a validated barium solution. The right panel distinguishes the results already supported by the exploratory notes from the preliminary candidate range and the still-unresolved double-magic condition.	256
-----	--	-----

LIST OF TABLES

3.1	First order field insensitive transitions considered in $^{133}\text{Ba}^+$	67
4.1	Resonator parameters and estimated blade voltage.	91
4.2	Selected components of the trap and vacuum system.	95
4.3	Measured isotopic distribution of barium in the sample.	97
4.4	Ablation Targets	100
5.1	Shorthand notation for the five manifolds considered here.	138
5.2	Laser Intensity vs Scattering Rate	143
9.1	Simulated final population after quenching and recycling an initially prepared $ D_{5/2}, F = 3, m_F = 3\rangle$ coolant ion with equal σ^+ , π , and σ^- polarization com- ponents.	238

ACKNOWLEDGMENTS

When I first arrived at UCLA, I knew absolutely nothing about experimental AMO physics. The only electives I had skipped as an undergraduate were optics and electronics, which turned out to be the two most useful classes for my PhD. Fortunately, I quickly found that everyone in the UCLA AMO community was kind, generous, and willing to teach. First, I want to thank Justin Christensen for allowing me to learn on the experiment he created. I now understand what it means to let a new student learn to drive for the first time in the sports car you built. The lessons I learned from being allowed to fail were innumerable, and Justin's patience was endless. I have tried to carry this practice forward with the next generation of barium trappers in the lab, Steven Diaz, Jiaxiang Wang and Elijah Mossman. They arrived in the lab at exactly the right time, when I was beginning to doubt my ability to single-handedly achieve our research objectives. Steven has revitalized the lab with a healthy dose of optimism, and Jiaxiang has brought fresh ideas to help drive the science forward. I have also been fortunate enough to get to know the other founder of the UCLA barium lab, Dave Hucul. Dave has made me believe in the dream of a $^{133}\text{Ba}^+$ quantum network, and I hope to see it realized one day. Dave and Justin crawled so that I could walk, and hopefully Steven, Elijah, and Jiaxiang can run.

The bulk of my time in the lab was spent alongside Sam Vizvary, who took up the first-generation barium trap, and Michael Bareian, who started the new integrated photonics effort. They were both valuable resources and good friends. The three of us also helped distribute the attention of our two bosses. If one of us had not gotten much work done, we could usually count on another to talk more during the meeting. I also want to thank Matt Boguslawski, who oversaw all three of us at the beginning of my time at UCLA. I will never forget being drenched in sweat while making radioactive targets at 4 a.m. with our team.

Next, I want to thank my two PIs, Eric Hudson and Wesley Campbell. I have been very lucky to have two perspectives on every difficult problem I have faced, especially when

those perspectives come from two of the smartest people I have ever met, and especially when neither of them minds telling the other that they are wrong. Over the years, Eric has molded me into a better physicist. He has motivated me with both sticks and carrots to learn lessons in the lab and in life. Eric is truly a big-picture thinker, and from him I have learned to dream big and to look for interesting science in new places, no matter how unfamiliar they may be to me. From Wes, I have learned to be faster, more technical, and to always put things in dB. Part of being a good scientist is being a good scrutinizer, finding the weak points or inconsistencies in a dataset or method. I have learned to always put error bars on a graph and actually consider what they mean, and above all, to question any graph I see without them. What I have learned from both Eric and Wes is how to collaborate and the value of a good partnership. Science is hard, and I have learned that being open and sharing ideas is much more rewarding than worrying about where credit will be assigned.

During my seven years in the lab at UCLA, I have made so many friends, had so many deep physics conversations, and shared so much super glue that it sometimes felt like I was the only person who ever bought it. On my first day, I was handed off to be mentored by Ricky Elwell, and I have stuck by his side ever since. There is a deep bond between those who are forced to do their PhDs with radioactive materials. Grant Mitts, Clayton Ho, and Hao Wu were huge helps as we concurrently built up ion traps, ARTIQ software packages, and optical tables. Anthony Ransford and Conrad Roman were like my uncles trapping ytterbium in the neighboring lab, helping me when I needed it, while also reminding me that barium was not the best ion for quantum computing (they both trap it now, though, so who's laughing?). Sami Khamis and Randy Putnam, who worked in the adjacent lab, were probably the most helpful because they convinced me to leave the lab and have some semblance of a life outside of physics. We nearly won the intramural basketball league, which, considering we are the physics department, is fairly impressive. Kristian Barajas was the refreshing AMO theorist we all needed, never afraid to ask the simple question that none of us actually understood. Alex Androutsos is hopefully the future generation

of the project and has picked up everything I have taught him very quickly, except for soccer, where he has not improved much. Everyone in UCLA AMO has contributed to my experience here, and I hope the culture we fostered (and the parties) continue for many more years.

To my parents, thank you for caring, for asking, and for genuinely trying to understand my research. Thank you for making me unroll my poster after conferences and present it to you, even though it was probably like listening to a foreign language. Thank you for editing every essay, poem, and now thesis since I was in kindergarten. I am so glad that APS accidentally sends my copy of *Physics Today* to your house, and that you attempt to relay the contents of any article that mentions quantum to me over the phone. You know that I have always looked for adventure, even from a young age, but one of the biggest challenges of my PhD was being so far away from home.

Luckily, I found a home away from home. Jacquelyn, despite the highs and lows of this work that you have witnessed firsthand, I have never once regretted coming to UCLA or doing my PhD, because this is the place where I met you. This journey has changed my life in more ways than one, and I am all the better for it.

VITA

2015–2019 B.S. in Physics, Chemistry and Mathematics, Rhodes College

2019–2026 Graduate Student Researcher in Physics, UCLA

PUBLICATIONS

M. J. Boguslawski*, **Z. J. Wall***, S. R. Vizvary, I. D. Moore, M. Bareian, D. T. C. Allcock, D. J. Wineland, E. R. Hudson, and W. C. Campbell, “Raman Scattering Errors in Stimulated-Raman-Induced Logic Gates in $^{133}\text{Ba}^+$,” *Physical Review Letters* **131**, 063001 (2023).

S. R. Vizvary, **Z. J. Wall**, M. J. Boguslawski, M. Bareian, A. Derevianko, W. C. Campbell, and E. R. Hudson, “Eliminating Qubit-Type Cross-Talk in the OMG Protocol,” *Physical Review Letters* **132**, 263201 (2024).

Z. J. Wall, J. D. Piel, S. R. Vizvary, M. Bareian, S. Diaz, E. Mossman, A. Ransford, C. H. Greene, E. R. Hudson, and W. C. Campbell, “Barium Autoionization for Efficient Ion Trap Loading,” arXiv:2512.12191 (2025).

CHAPTER 1

Introduction

“Information is physical.” – Rolf Landauer

What if, instead of combining individual chemicals in a vial, one could simulate chemistry on a computer?

The marriage of machines and computation in the 20th century has guided science toward many new discoveries. Before modern computers, scientists used raw intuition to develop an incredible foundation of analytical theories to model nature. But understanding the fundamental equations that govern a system does not necessarily mean one can predict its behavior. To solve substantive problems with broad applications in everyday life, one must often search a massive parameter space. Even if the optimal parameters are known, testing the system and calculating the relevant observables may still be extremely labor-intensive. Thus, many of the incredible discoveries in chemistry and physics have emerged from an army of academic and industrial scientists testing, and often serendipitously stumbling upon, the unique concoctions that push human knowledge and ability forward.

By the time I came online in 1996, kindergartners were already using iMacs in school. When I arrived at college, computing technology had advanced to the point that first-year undergraduates could build custom molecules and proteins in software, almost like assembling Tetris blocks. Under the guidance of Professor Mauricio Cafiero at Rhodes College, I learned about modeling software such as *Gaussian*, which applies density functional theory to approximate the quantum mechanical structure of systems containing thousands of atoms (Evans et al., 2018). The goal of this research is to predict candidate

molecules that may bind to a particular enzyme or cellular target and to calculate their associated binding energies. If this process could be made both accurate and efficient, it would dramatically accelerate drug design; for example, by enabling the development of therapeutics that regulate dopamine levels in patients with Parkinson’s disease.

My research at Rhodes led me to the lab of another computational chemist, Professor Graeme Henkelman at the University of Texas at Austin. In his group, students were designing catalysts for the efficient production of hydrogen fuel as a renewable form of energy storage. The aim was to replace platinum, a scarce and prohibitively expensive catalyst used in existing systems. Their simulations involved nanoparticles containing tens of thousands of atoms and were executed in parallel on one of the largest supercomputers in the world. Individual atoms within these complex structures were rearranged or substituted across thousands of candidate configurations to identify those with the best performance. Yet even with such a powerful computing cluster, accuracy had to be sacrificed. In fact, many of these structures were modeled primarily using the purely classical Lennard-Jones potential. It was there, while trying to improve the efficiency of these computations, that I first heard about the prospect of quantum computing.

The challenge with predicting quantum particles is that they are not described by deterministic variables, they live in the world of *probability amplitudes*. When another particle is added to the system, it is not simply described by one more variable, instead an entirely new column and row must be added to the tensor object that describes the entire system. This is because the particles are *entangled*, and the degree to which they are entangled with every other particle in the system is quantified in the new column and row added. The result is an exponential scaling, both in the variables required to be stored, and the operations required to evolve and measure the system. Crucially, it is the effect of this enigmatic entanglement that we hope to calculate in real chemical systems. Thus, the idea naturally comes into view: what if one could use quantum particles to represent and store these probability amplitudes?

The catch is that, in order to make a register of quantum particles and control their

entanglement, they must be isolated from the environment. The environment is, in effect, constantly performing *measurements*, or extracting information. Thus in order to build a quantum computer, and make a useful measurement, one must prevent the environment from measuring the system first. This is a daunting task since in some sense everything in the universe is entangled. A muon passing through a vacuum chamber is, to some degree, entangled with the particles inside. Moreover, if one incorrectly encodes the quantum register because of imperfect control of their experimental levers, this can be equivalent to a loss of information to the environment. This erosion of the coherent phase relationships within the quantum register is known as *decoherence*.

The work presented in this dissertation focuses on developing tools for achieving an unprecedented degree of control over a quantum system while suppressing decoherence beyond previous demonstrations. These advances are realized experimentally in a trapped $^{133}\text{Ba}^+$ ion system. Broadly, the dissertation is organized into three phases. The first phase centers on dramatically improving the loading efficiency of the first-generation $^{133}\text{Ba}^+$ trap and on demonstrations of laser-based Raman gates. The second phase involves the design and construction of a new state-of-the-art $^{133}\text{Ba}^+$ trap, together with the development of a novel photoionization loading method. The final phase culminates in record-breaking demonstrations of state preparation and measurement in the second-generation trap. The final section of this dissertation outlines planned future experiments and presents a possible roadmap for realizing a useful quantum computer in an academic setting.

1.1 Are Classical Computers Stuck?

1.1.1 A Brief History of Compute

There is no reason to believe that classical computing technology has already reached its peak. On the contrary, we may be living through a second renaissance in digital computing. New paradigms are advancing rapidly, including optical computing, analog computing, graphics processing units, and artificial intelligence, to name only a few. Each of these has

already transformed, or has the potential to transform, important computational problems that affect everyday life. Why, then, do we need a quantum computer?

To answer that question, it is useful to look back to a time before the universal digital machine most of us imagine when they picture the word computer. The first known computer was recovered from a shipwreck in 1905 and dubbed the Antikythera mechanism. The device dates to roughly 200 BCE and is of Greek origin. Although some details of its use remain uncertain, it is widely understood to have been a hand-driven system of gears used to predict astronomical cycles, including eclipses.

Knowledge of devices of this kind was lost to history. Much later, one of the first truly useful modern computers was developed by William Thomson in the nineteenth century. The machine consisted of pulleys and a device known as a Scotch yoke, which converted linear motion into rotational motion. Since, the purpose of this device was to predict the tides, it was sometimes called a tidal clock. It illustrates the deep connection between clocks and computers: both are physical systems engineered to encode and process regular patterns of motion. In fact, many of the current hardware candidates for quantum computers, are also used as atomic clocks.

By the time of Thomson's invention, Pierre-Simon Laplace had already derived the differential equations that govern tidal motion. These equations incorporated the motion of the Moon and the Sun, Earth's rotation, and the fluid dynamics underlying the tides. Rather than solving these equations directly, Thomson's machine evaluated a harmonic approximation of the relevant frequencies, calibrated using amplitudes and phases inferred from real tidal data . Once its pulleys were calibrated, the machine had to be hand-cranked for hours to produce tidal plots that would previously have required months of human calculation. Thomson's tide machine is illustrative of what early instantiations of computers looked like. They were devices designed for the efficient calculation of a single algorithm.

Mechanical analog computers found extensive use in the early twentieth century, reaching a peak in devices such as the Norden bombsight. Developed for the U.S. military during

World War II, the bombsight used gyroscopic stabilization and mechanical computation to predict, in real time, the trajectory of a bomb released from an aircraft. Yet despite its sophistication, it fell well short of its legend, achieving a practical accuracy of only about 1,000 feet under combat conditions. This limitation reflected a broader problem with analog computation, because intrinsic physical errors accumulated as these devices were scaled to more complex calculations. Even modern electrical analog computers remain fundamentally limited in precision, with typical errors often no better than about 0.1% (Yang et al., 2025). Moreover, analog computers cannot generally be made strictly fault tolerant through error correction, because continuous errors do not admit discrete correction in the way digital errors do. This same logic was once used to dismiss the possibility of quantum computing, until the discovery of quantum error correction.

Modern computers emerged when information became digitized and was operated on in discrete steps, as in the British Bombe machine developed by Alan Turing at Bletchley Park. Digital computation relies on discrete states. For example, a voltage in a transistor may be interpreted as logical true if it exceeds a threshold, and logical false if it falls below that threshold. This discreteness makes digital computers far less error-prone than analog devices, since small miscalibrations do not matter unless they push the signal across the threshold. It also allows for error corrections or the removal of entropy that builds up in the system. If, after many operations, an input voltage has drifted slightly but remains within the correct threshold, the information can still be read out and copied reliably to a new transistor. Finally, digital machines are highly adaptable, since the same hardware can be reprogrammed through software to perform many different tasks. Quantum computers, by contrast, are not expected to replace classical computers, but rather to work alongside them on problems for which quantum methods offer a distinct advantage. Quantum computers, like the tide predicting machines, are unlikely to resemble today's general-purpose universal digital machines. Only a small number of known algorithms are currently believed to offer a meaningful advantage on a quantum computer.




Figure 1.1: Scaling of different classical algorithms

1.1.2 Quantum Chemistry Scaling

The natural question, then, is what can classical computers not do. More precisely, even with another decade of improvements in microchip engineering and data-center-scale computation, which problems will remain fundamentally out of reach? There are classes of problems that will still fill scientists with dread, even if attacked with the largest supercomputers imaginable. Their defining feature is that the number of operations required to solve them grows exponentially with system size.

Algorithms are typically classified by their asymptotic scaling, or order of growth. Some tasks are linear, $O(N)$; for example, finding the maximum value in a list requires only a single pass through the data. Others are sublinear, such as binary search, which scales as $O(\log N)$ by repeatedly halving the search space. Many of the most practically important problems are polynomial. Matrix multiplication, for example, is often written as $O(N^3)$, and modern GPUs are designed to perform such operations extremely efficiently. But beyond these lie the true leviathans: exponential problems. A classic example is the

traveling salesman problem: given N cities, what is the shortest route that visits each city exactly once and returns to the starting point? Although simple to state, this problem becomes computationally intractable for classical machines at surprisingly small N . This contrast in scaling is illustrated in Fig. 1.1.

Let’s return to our previous example of quantum chemical calculations and analyze the scaling of these problems. The computational challenge of quantum chemistry can be understood by comparing a simple atom such as helium with a strongly correlated molecule such as the iron dimer. In the Born–Oppenheimer approximation, the electronic Hamiltonian for helium is

$$\hat{H} = \sum_{i=1}^2 \left(-\frac{1}{2} \nabla_i^2 - \frac{2}{r_i} \right) + \frac{1}{r_{12}}, \quad (1.1)$$

where the terms represent the electronic kinetic energy, the attraction of each electron to the nucleus, and the electron–electron repulsion. In Hartree–Fock theory, the many-electron wavefunction is approximated by a single Slater determinant, reducing the problem to a self-consistent mean-field description. In a basis of M spatial orbitals, each molecular orbital is expanded as

$$\phi_i(\mathbf{r}) = \sum_{\mu=1}^M C_{\mu i} \chi_{\mu}(\mathbf{r}), \quad (1.2)$$

and the resulting Roothaan–Hall equations are

$$FC = SC\varepsilon. \quad (1.3)$$

The dominant computational cost arises from constructing the Fock matrix,

$$F_{\mu\nu} = h_{\mu\nu} + \sum_{\lambda\sigma} P_{\lambda\sigma} \left[(\mu\nu|\lambda\sigma) - \frac{1}{2}(\mu\lambda|\nu\sigma) \right], \quad (1.4)$$

which depends on the four-index two-electron integrals $(\mu\nu|\lambda\sigma)$. As a result, conventional Hartree–Fock scales formally as $\mathcal{O}(M^4)$ per self-consistent-field iteration, with an additional

$\mathcal{O}(M^3)$ diagonalization step. For helium, the basis size is so small that the entire calculation is trivial on modern hardware.

The formal structure of the Hamiltonian is unchanged for a much more difficult system such as the iron dimer, Fe_2 , for which the electronic Hamiltonian is

$$\hat{H} = \sum_{i=1}^{52} \left(-\frac{1}{2} \nabla_i^2 - \frac{26}{r_{iA}} - \frac{26}{r_{iB}} \right) + \sum_{i<j}^{52} \frac{1}{r_{ij}} + \frac{26^2}{R_{AB}}. \quad (1.5)$$

Here one must describe 52 interacting electrons in the field of two transition-metal nuclei. Hartree–Fock remains computationally feasible because it still optimizes only a single determinant, and therefore retains polynomial scaling in the basis size. However, for Fe_2 this approximation becomes physically inadequate. The partially filled $3d$ and $4s$ shells generate many near-degenerate electronic configurations, so no single Slater determinant provides a satisfactory description of the ground state. This is the hallmark of strong static correlation.

In an orbital occupation basis, this failure of the single-determinant picture may be understood in terms of entanglement. A Hartree–Fock state is specified by a single occupation pattern, apart from the antisymmetry required by fermionic exchange. By contrast, a strongly correlated state must be represented as a superposition of many different occupation patterns. In this sense, electronic correlation is the appearance of nontrivial entanglement between orbital degrees of freedom: the wavefunction can no longer be represented by a single mean-field configuration.

The formally exact remedy within a chosen one-electron basis is full configuration interaction (FCI), in which the wavefunction is expanded in all possible Slater determinants,

$$|\Psi_{\text{FCI}}\rangle = \sum_I c_I |D_I\rangle. \quad (1.6)$$

If the basis contains K spin-orbitals and the system contains N electrons, then the

dimension of the determinant space is

$$\dim \mathcal{H}_{\text{FCI}} = \binom{K}{N}. \quad (1.7)$$

This combinatorial growth is the origin of the exponential difficulty of exact quantum chemistry. For helium, the FCI space remains tiny in any modest basis, so exact diagonalization is straightforward. For Fe_2 , however, the situation is completely different. With $N = 52$ electrons and a chemically realistic basis containing on the order of $K \sim 100$ spin-orbitals, the determinant space has dimension

$$\binom{100}{52} \sim 10^{29}, \quad (1.8)$$

which is far beyond any conceivable classical memory or compute budget. Even storing the wavefunction coefficients would be impossible, let alone constructing and diagonalizing the Hamiltonian in that basis.

The contrast in computational resources is therefore stark. Hartree–Fock requires storing and optimizing one-particle orbitals and constructing the Fock matrix, giving polynomial cost in the basis size and making calculations feasible for quite large systems. FCI, by contrast, requires access to the full correlated many-electron Hilbert space, whose dimension grows combinatorially with the number of electrons and orbitals. For simple atoms such as helium, this distinction is of little practical consequence. For strongly correlated transition-metal systems such as Fe_2 , it becomes decisive. Thus, the central challenge of quantum chemistry is not merely the evaluation of a complicated Hamiltonian, but the representation of the correlated many-electron wavefunction. Mean-field methods remain tractable because they avoid the full correlated structure of the problem, whereas exact methods such as FCI become intractable precisely because they attempt to represent it exactly.

1.2 What is a Quantum Computer?

A quantum computer is a device whose fundamental computational elements obey the laws of quantum mechanics. Building such a machine requires the ability to initialize, preserve, manipulate, and read out quantum states over timescales sufficient to carry out a computation. The basic unit of quantum information, the quantum bit or qubit, must be coherently controllable, so that arbitrary single-qubit rotations can be performed on the Bloch sphere and entangling two-qubit operations, often represented by a CNOT gate, can be implemented. While a quantum computer can in principle reproduce any classical logical operation, its significance lies in algorithms for which quantum devices are known or expected to provide an advantage over the best known classical approaches, typically through more favorable scaling with system size.

To elucidate this idea, we can examine one of the simplest quantum algorithms, Deutsch's algorithm. It shows how a quantum computer can extract a global property of a function using fewer queries than a classical computer. The task is to determine whether a one-bit function $f(x)$, with $x = 0$ or 1 , is *constant*, meaning $f(0) = f(1)$, or *balanced*, meaning $f(0) \neq f(1)$. Classically, both $f(0)$ and $f(1)$ must be evaluated to determine this with certainty. In Deutsch's algorithm, the input qubit is prepared in a superposition of both possible inputs, the function is applied once through a quantum oracle, and quantum interference is used to convert the relative phase information into a measurable outcome, ie.

$$\frac{1}{\sqrt{2}} \left((-1)^{f(0)} |0\rangle + (-1)^{f(1)} |1\rangle \right) \xrightarrow{H} \begin{cases} \pm |0\rangle, & f(0) = f(1) \quad (\text{constant}) \\ \pm |1\rangle, & f(0) \neq f(1) \quad (\text{balanced}) \end{cases} \quad (1.9)$$

After a final Hadamard gate, measurement of the input qubit deterministically reveals whether the function is constant or balanced (Deutsch, 1985). In this way, the function need only be queried once, giving the quantum algorithm an advantage over the classical one.

1.2.1 The Quantum Phase Estimation Algorithm

Let's return to our quantum chemical calculation to see how one might use a quantum algorithm. In the quantum algorithm for molecular energy estimation, the electronic Hamiltonian of the molecule is first expressed in a finite one-electron basis and then mapped onto a register of qubits. In the second-quantized representation, the many-electron wavefunction is written in the occupation-number basis, so that each basis state is specified by which orbitals are occupied and which are empty. In the simplest direct mapping, each qubit stores the occupation of a single orbital. Thus, $|0\rangle$ denotes an unoccupied orbital and $|1\rangle$ denotes an occupied orbital. More commonly in modern quantum chemistry language, one uses spin-orbitals, so that a spatial orbital contributes two qubits, one for spin-up occupation and one for spin-down occupation. A Hartree–Fock reference state is then represented by a computational basis bitstring such as,

$$|110010\dots\rangle, \tag{1.10}$$

while a correlated molecular wavefunction is represented by a superposition of many such occupation strings. In this way, the qubit register directly encodes the electronic Fock space of the molecule.

Once the molecular Hamiltonian has been mapped to a qubit Hamiltonian, quantum phase estimation can be used to determine a molecular energy. If a trial state $|\psi_{\text{trial}}\rangle$ has nonzero overlap with an eigenstate $|E_k\rangle$ of the molecular Hamiltonian H , then under time evolution it acquires a phase proportional to the corresponding eigenenergy:

$$e^{-iHt} |E_k\rangle = e^{-iE_k t} |E_k\rangle. \tag{1.11}$$

Expanding the trial state as

$$|\psi_{\text{trial}}\rangle = \sum_k c_k |E_k\rangle, \tag{1.12}$$

one finds that phase estimation returns the energy E_k with probability $|c_k|^2$. The role of the

ancilla register is to extract this phase, converting the problem of molecular spectroscopy into the problem of measuring an eigenphase of the unitary operator e^{-iHt} . In the recursive version of phase estimation used by (Aspuru-Guzik et al., 2005), this phase can be determined iteratively while using only a small readout register, thereby reducing the total number of qubits required for the energy measurement.

In practice, the molecular electronic Hamiltonian is first mapped onto qubits as a sum of weighted Pauli strings,

$$H = \sum_j c_j P_j, \quad (1.13)$$

where each P_j is a tensor product of Pauli operators acting on the qubit register. The full time-evolution operator e^{-iHt} is then approximated by a Trotter expansion, in which the evolution under the full Hamiltonian is decomposed into a product of short evolutions under the individual Pauli terms,

$$e^{-iH\Delta t} \approx \prod_j e^{-ic_j P_j \Delta t}. \quad (1.14)$$

Each factor $e^{-ic_j P_j \Delta t}$ is implemented as a circuit by first rotating the relevant qubits into a basis in which the Pauli string becomes a product of Z operators, then using a ladder of CNOT gates to accumulate the parity onto one qubit, applying a single R_z rotation, and finally undoing the CNOTs and basis changes. In this way, one Trotter step consists of a long sequence of elementary one- and two-qubit gates that approximates the time evolution generated by the molecular Hamiltonian (Aspuru-Guzik et al., 2005).

A simple estimate of the cost may be obtained from the second-quantized molecular Hamiltonian,

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \frac{1}{2} \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s. \quad (1.15)$$

For M spin-orbitals, the one-electron part contains $O(M^2)$ terms, while the two-electron part contains $O(M^4)$ terms. After a Jordan–Wigner mapping, each fermionic term becomes a Pauli string whose implementation requires a circuit of length $O(M)$ in the worst case,

due primarily to the required CNOT ladder. Consequently, a single first-order Trotter step has cost

$$O(M^4) \times O(M) = O(M^5) \quad (1.16)$$

elementary logical gate operations. To resolve the energy to precision ϵ , quantum phase estimation requires total evolution time $T = O(1/\epsilon)$, giving the rough overall scaling

$$O\left(\frac{M^5}{\epsilon}\right). \quad (1.17)$$

For a concrete example, consider a molecular problem with $M = 100$ spin-orbitals. In the direct second-quantized mapping, this requires approximately 100 qubits to represent the occupations of the spin-orbitals, together with a small ancilla register for recursive phase estimation. Using the scaling above, one first-order Trotter step costs roughly

$$100^5 = 10^{10} \quad (1.18)$$

elementary logical gate operations. If the target precision is $\epsilon \sim 10^{-3}$ Hartree, then the total logical gate count is crudely estimated as

$$\frac{100^5}{10^{-3}} = 10^{13}. \quad (1.19)$$

Thus, for a problem of this size, the early quantum algorithm is polynomial rather than exponential, but still extremely resource intensive. Its advantage over full configuration interaction lies not in being small, but in avoiding the exponential growth of the exact classical Hilbert space.

1.2.2 Where Does the Speed Up Come From?

Now that we have examined a few quantum algorithms, we can broach the question of why a quantum computer can be faster than a classical one. Although this is a somewhat

controversial topic, I would argue that its origin lies in the unusual power of quantum correlations. Specifically, entangled systems have the ability to instantaneously rearrange the coherences of massive density matrices that describe the total state. The ability to prepare, manipulate, and control superpositions of such states, and to engineer their interference in a useful way, allows quantum algorithms to enhance the amplitudes of desired outcomes and thereby achieve improved scaling for certain problems.

The word instantaneous used above must be carefully qualified, and to do so accurately we must define entanglement. Classical correlations occur in everyday life; for instance, if one foot is found to be wearing a red sock, one immediately infers that the other foot has the matching red sock (Preskill, 1998). In this case, however, each sock possesses a definite color prior to observation, and the correlation reflects only incomplete knowledge of a pre-existing classical state.

Thus, a quantum correlation is fundamentally different. Suppose we have two entangled qubits in the state,

$$|\Phi^+\rangle_{12} = \frac{1}{\sqrt{2}} (|0\rangle_1 |0\rangle_2 + |1\rangle_1 |1\rangle_2). \quad (1.20)$$

Measuring qubit 1 (say the outcome is $|0\rangle$) gives a definite result and, at the same time, fixes the outcome of a subsequent measurement of qubit 2: it must also be $|0\rangle$. Prior to this measurement, however, if we had measured qubit 2 first, the outcomes $|0\rangle$ and $|1\rangle$ would have been equally likely. In this sense, the measurement of one affects the description of the other. This effectuation also happens faster than the speed of light, which is truly mind-blowing and was the subject of the 2022 Nobel prize (Aspect et al., 1982; Freedman and Clauser, 1972; Bouwmeester et al., 1997).

It should be noted that a simple proof shows that, despite these apparently instantaneous correlations, entanglement cannot be used to transmit information faster than light. This is called the no-signaling property and to prove it we must consider any local strategy Alice (a generic transmitter/receiver of information) might use to encode information.

Suppose Alice and Bob (her friend) share the Bell state

$$|\Phi^+\rangle_{AB} = \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B + |1\rangle_A |1\rangle_B). \quad (1.21)$$

The corresponding density matrix is

$$\rho_{AB} = |\Phi^+\rangle \langle \Phi^+|, \quad (1.22)$$

and Bob's local state is obtained by tracing over Alice's subsystem,

$$\rho_B = \text{Tr}_A(\rho_{AB}) = \frac{I}{2}. \quad (1.23)$$

Now suppose Alice attempts to send a classical bit by choosing between two different local operations on her qubit. For example, she may do nothing to send a 0, or apply some unitary U_A to send a 1. After Alice applies U_A , the joint state becomes

$$\rho'_{AB} = (U_A \otimes I) \rho_{AB} (U_A^\dagger \otimes I). \quad (1.24)$$

Bob's reduced state is then

$$\rho'_B = \text{Tr}_A(\rho'_{AB}) = \text{Tr}_A \left[(U_A \otimes I) \rho_{AB} (U_A^\dagger \otimes I) \right] = \rho_B = \frac{I}{2}. \quad (1.25)$$

Thus no local unitary applied by Alice can change Bob's local state. Therefore, although Alice's action changes the global state description of the entangled pair, it does not alter Bob's local measurement statistics. Since communication would require Bob's local statistics to depend on Alice's choice of operation, entanglement alone cannot be used to transmit information faster than light.

This entire conversation of "where does the speed up come from?" must be qualified by the important Gottesman-Knill Theorem. As it turns out, you can keep track of certain quantum algorithms, so long as the circuit that was used to make said algorithm is

contained in the Clifford group. This means that simply by keeping track of the stabilizer measurements which scale with n you can fully reproduce the 2^n quantum state that is made. However, once so-called T gates are added to the circuit, it is no longer classically simulatable (Gottesman, 1998). The T gate is special because, unlike Clifford phase gates such as Z and S, it rotates a stabilizer state into a state whose phase coherence lies between Pauli axes, so the state is no longer completely characterized by a finite set of Pauli stabilizer eigenvalue constraints.

1.3 How To Build a Quantum Computer

Quantum computing is presently in an exciting period of development. For the first time, major companies are investing heavily in a field that was once almost entirely academic. Nevertheless, despite nearly three decades of progress since the first experimental demonstrations of quantum gates, practical large-scale quantum computation remains distant. One of the leading recent platforms has demonstrated control of 98 qubits (Ransford et al., 2025), a remarkable milestone that places such devices beyond straightforward classical simulation. However, this system has not yet demonstrated a practically useful computation with a clear quantum advantage. As a result, we remain in what John Preskill termed the *Noisy Intermediate-Scale Quantum* (NISQ) era. This is defined as an era in which quantum processors can be built and controlled, but noise, decoherence, and gate errors still accumulate to the point that useful large-scale computation is not yet viable.

1.3.1 Is It Possible to Build a Quantum Computer?

Current proposals for useful quantum algorithms suggest that factoring a 2048-bit integer would require on the order of 10^6 physical qubits and 10^{12} operations (Gidney and Ekerå, 2021). Qubits will never be perfect, and with current gate schemes there are fundamental physical limits to the achievable fidelity of operations. At present, it is difficult to imagine average operational error rates below roughly 10^{-6} . At that error rate,

a computation involving 100 qubits and 1000 operations on each qubit would succeed with probability of about 90%, assuming independent errors. Increasing this to 10,000 operations per qubit would reduce the success probability to $e^{-1} \approx 37\%$. A natural question, then, is what level of success probability is actually required. For a problem with only two possible outputs, any success probability below 50% is indistinguishable from random guessing. For problems with a broader distribution of outputs, however, the computation may be repeated many times to build statistics and identify outcomes that occur with slightly higher probability than others.

At the end of the day (ATEOTD), this imposes a fundamental limit on quantum computers before they can become useful in comparison with classical devices. Errors obscure any opportunity to exploit quantum information for computational speedup, because quantum information is considerably more fragile than classical digital information. Moreover, in a quantum circuit, errors need not remain confined to the qubits on which they originate; through entangling gates, a single local fault can propagate into a correlated multi-qubit error, further diminishing the probability of successful computation. This line of reasoning was once seen as one of the strongest arguments against pursuing quantum computing as a serious scientific endeavor. However, Peter Shor demonstrated that quantum information can be protected, and that errors arising during a quantum computation can, in principle, be corrected (Shor, 1995).

This landmark discovery by Shor helped launch the field of quantum error correction. To simplify the discussion, we will focus on one of its most fundamental examples: the nine-qubit Shor code. It protects against bit-flip and phase-flip errors. The logical basis states are

$$|0_L\rangle = \frac{1}{2\sqrt{2}} (|000\rangle + |111\rangle) (|000\rangle + |111\rangle) (|000\rangle + |111\rangle), \quad (1.26)$$

$$|1_L\rangle = \frac{1}{2\sqrt{2}} (|000\rangle - |111\rangle) (|000\rangle - |111\rangle) (|000\rangle - |111\rangle). \quad (1.27)$$

Within each block of three qubits, the repetition structure protects against bit-flip errors, while the relative phase between the three blocks protects against phase-flip errors. Since

an arbitrary single-qubit error can be expanded in the Pauli basis as

$$E = aI + bX + cY + dZ, \tag{1.28}$$

the ability to detect and correct both X and Z errors implies protection against any single-qubit error, including $Y = iXZ$. In this way, the Shor code demonstrated for the first time that quantum information could be redundantly encoded and actively protected from decoherence without violating the no-cloning theorem.

In practice, the encoded state is protected by repeatedly measuring its stabilizer operators using ancillary qubits. These ancilla qubits are first prepared in a known state, then entangled with the data qubits in such a way that the ancilla acquires information about the error syndrome without revealing the logical amplitudes of the encoded state. For example, measuring a parity operator such as $Z_i Z_j$ reveals whether two qubits have the same or opposite bit value, while leaving the logical superposition intact. After the ancilla is measured, the outcome identifies whether an error has occurred and, if so, which correction should be applied. In this way, the ancilla serves as a temporary probe that extracts error information from the code space, allowing the computation to remove entropy and restore the encoded logical qubit without directly measuring or collapsing the quantum information itself.

The purpose of quantum error correction is to detect and correct errors before they propagate through a computation and corrupt the encoded information. In this sense, error correction involves a fundamental tradeoff between the overhead of additional qubits and operations in exchange for a higher probability of successful computation. The difficulty is that the encoding circuits, repeated syndrome measurements, and corrective operations are themselves imperfect and can introduce new faults. Moreover, once quantum information is distributed across many physical qubits to form a logical qubit, the algorithm must be implemented at the logical level, where new constraints arise that are absent for bare physical qubits. Logical T gates provide an important example. Although a physical

T gate is a simple single-qubit phase rotation, in many stabilizer-code architectures the corresponding logical T gate is not available transversally and must instead be implemented using a prepared non-stabilizer resource state, often called a magic state. This state is prepared and verified or distilled in a separate part of the computer and then consumed through gate teleportation to enact the logical T operation. As a result, logical T gates can require substantially more time, qubits, and circuit resources than their physical-gate counterparts.

Thus, it is crucial to evaluate the overhead associated with different quantum error-correcting codes in technical detail. In particular, one would like to determine the physical error-rate threshold below which a given code and fault-tolerant implementation actually suppresses errors at the logical level. A detailed threshold calculation of this kind was carried out by Aliferis, Gottesman, and Preskill for the $[[7, 1, 3]]$ Steane code (Aliferis et al., 2006). They showed that, if each physical circuit location fails with probability below approximately 2.7×10^{-5} , then arbitrarily long quantum computations can be performed.

1.3.2 Where Are We Now?

The preceding discussion gives a rough picture of the hardware requirements for fault-tolerant quantum computation. This thesis focuses mainly on improving qubit initialization, single-qubit gates, and measurement. Historically, the main fidelity bottleneck has often been assumed to be the two-qubit gate, since entangling gates are usually harder to implement with high fidelity and are required many times in large quantum algorithms. Without error correction, it is also natural to think that state preparation and measurement are less important, since initialization happens near the beginning of a circuit and measurement happens near the end.

Fault-tolerant error correction changes this picture. In an error-corrected computation, state preparation and measurement are not rare events. They occur repeatedly during syndrome extraction, ancilla preparation, verification, logical readout, and in many architectures, magic-state preparation. As a result, State Preparation and Measurement ()

errors directly affect the syndrome information used by the decoder. A faulty measurement can lead to an incorrect correction, while poor state preparation can inject errors before the computation has even begun.

The large number of repeated error-correction cycles also makes the speed of SPAM operations important. While the processor waits for initialization or measurement to finish, the data qubits continue to accumulate memory errors. Longer measurement and reset times also increase the chance that correlated errors occur during the correction cycle. For this reason, improving SPAM performance requires improvements in both fidelity and latency. Even when two-qubit gates dominate threshold estimates, fast and high-fidelity SPAM remains essential for reducing logical error rates and lowering the overhead of fault-tolerant quantum computation.

It should also be noted that the reported record SPAM infidelity of approximately 1×10^{-4} in Ref. (An et al., 2022a) is comparable to the reported record two-qubit gate error in Ref. (Hughes et al., 2025). However, the relevant timescales are very different. The two-qubit gate in Ref. (Hughes et al., 2025) was performed in $5 \mu\text{s}$, whereas the SPAM sequence in Ref. (An et al., 2022a) required a time on the order of 4 ms. Thus, even when the achieved error probabilities are comparable, the latency associated with SPAM represents a major source of overhead. Reducing this timescale is therefore an important target for improving the practicality of error-corrected quantum computation.

Finally, it is worth stepping back and asking what it would mean if a universal, fault-tolerant quantum computer ultimately proved impossible to build in practice. In my view, the endeavor would still be worth pursuing. Even if the original goal were not fully realized, the effort required to control, entangle, measure, and stabilize large quantum systems is likely to produce technologies and scientific insights with value beyond quantum computation itself. History suggests that ambitious scientific projects often generate their most important consequences indirectly, through new tools, techniques, and ways of thinking. In this sense, the attempt to build a quantum computer may be worthwhile not only because of the machine it seeks to create, but also because of the discoveries that

may emerge along the way.

CHAPTER 2

The Ion Trap

“We never experiment with just one electron or atom or (small) molecule. In thought-experiments we sometimes assume that we do; this invariably entails ridiculous consequences... we are not experimenting with single particles, any more than we can raise Ichthyosauria in the zoo” – Erwin Schrodinger

2.1 Introduction

Isolation is the ultimate tool of an experimentalist. The ability to probe, manipulate, and remove the constituent parts of the environment enables the fastest route to troubleshooting complex issues. The invention of the ion trap led to a new era of isolation by allowing for the suspension of material, even single atoms, almost completely ambivalent toward the matter surrounding them. For this new invention, Wolfgang Paul won the 1989 Nobel Prize. In his Nobel lecture, he explained that the idea grew out of early work on molecular beams and the desire to manipulate their trajectories using lenses (Paul, 1990). Before the ion trap existed, the quadrupole mass spectrometer confined charged particles in two dimensions. The goal of these devices was to focus or collimate atomic particles, in order to increase the sensitivity for spectroscopy experiments, as well as to separate the constituents of the beams based on their atomic properties. These goals of collimation and separation are two key steps on the road to complete confinement and isolation of single particles.

Stable equilibrium corresponds to a position of the system in which its potential energy $U(q)$ is a minimum. A movement away from this position results in the setting up of a force $-dU/dq$ which tends to return the system to equilibrium. Let the equilibrium value be at q_0 , for a small deviation from this position it is sufficient to retain only the first non-vanishing term in the expansion of $U(q) - U(q_0)$ in powers of $q - q_0 = x$ (Landau and Lifshitz, 1976). In general the second order term $\frac{1}{2}k(q - q_0)^2$ where k is a positive coefficient for the value of the second derivative $U''(q_0)$, thus

$$U(q) = \frac{1}{2}kq^2 \quad (2.1)$$

This simple derivation is the motivation for the potential we aim to design in trapping a particle. This can also be described as a binding force and written as

$$F = -kr \quad (2.2)$$

One can imagine that at the origin a particle under this force experiences zero force, but as it ventures farther from the trap center a linearly increasing restoring force is felt. From this we can arrive at a potential,

$$F = -\nabla U \implies \Phi = \alpha x^2 + \beta y^2 + \gamma z^2 \quad (2.3)$$

At this point our goal is to construct a realistic electric field that produces this potential. In some sense it would be ideal if we could construct this potential using only static charges, but Earnshaw's theorem specifically forbids this. The proof is quite elegant, so I will reproduce it here.

Consider a test charge q placed in a static electrostatic potential $\phi(\mathbf{r})$. Its potential energy is

$$U(\mathbf{r}) = q\phi(\mathbf{r}). \quad (2.4)$$

A stable equilibrium point would correspond to a local minimum of $U(\mathbf{r})$, and therefore of

$\phi(\mathbf{r})$ for $q > 0$. In electrostatics, the electric field is related to the potential by

$$\mathbf{E} = -\nabla\phi, \quad (2.5)$$

and Gauss's law gives

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon_0}. \quad (2.6)$$

In a charge-free region of space, $\rho = 0$, so the potential satisfies Laplace's equation,

$$\nabla^2\Phi = 0. \quad (2.7)$$

This statement implies that there cannot be a local minimum of $\phi(\mathbf{r})$ across this region because in order for this to be true the curvature of the field must be positive in every direction.

With this constraint, it becomes clear that a purely static electric field cannot produce a three-dimensional confining potential in free space. To obtain stable confinement, one must instead employ time-dependent fields, i.e., dynamic stabilization. A convenient and commonly used choice is a sinusoidal drive.

We therefore consider the lowest-order (quadratic) form of an electrostatic potential near the origin,

$$\Phi(\mathbf{r}) = \frac{\Phi_0}{2r_0^2} (\alpha x^2 + \beta y^2 + \gamma z^2). \quad (2.8)$$

In a charge-free region the potential satisfies Laplace's equation, $\nabla^2\Phi = 0$, which immediately implies

$$\alpha + \beta + \gamma = 0. \quad (2.9)$$

This equation is really just a restatement of Earnshaw's theorem. The coefficients here represent the curvatures of the field, which cannot all be positive. Thus, we cannot have a purely bowl-shaped potential. There are two unique solutions to this equation that correspond to different potential shapes. One focusing direction, one defocusing direction

and one flat direction, where $\alpha = 1 = -\gamma, \beta = 0$ up to permutation. This is a 2D saddle with no quadratic curvature along the remaining axis. This corresponds to a linear quadrupole trap with no endcaps. The other case is $\alpha = \beta = 1, \gamma = -2$. This corresponds to a rotationally symmetric quadrupole potential often seen in a hyperbolic ring trap with endcaps. For the purposes of this thesis we will focus on the first scenario, and assume we have the geometry of four hyperbolic electrodes. At this point we must add in a periodic element that causes the axis of defocusing and the axis of focusing to switch in time; thus, restraining the particle

$$\Phi_0 = U + V \cos \omega t. \quad (2.10)$$

2.2 Solving the Equations of Motion

Our next goal is to solve for the dynamics of a particle experiencing this potential. The spatial coordinates are decoupled and so without loss of generality only x is treated here. To begin we start with the familiar equation from Newton ([Leibfried et al., 2003](#))

$$\ddot{x} = -\frac{Z|e|}{m} \frac{\partial \Phi}{\partial x} = -\frac{Z|e|}{mr_0^2} [U\alpha + V \cos(\omega_{rf}t)\alpha']x. \quad (2.11)$$

This can then be transformed to the standard form of the Mathieu differential equation

$$\frac{d^2x}{d\tau^2} + [a_x - 2q_x \cos(2\tau)]x = 0, \quad (2.12)$$

using the substitution

$$a_x = \frac{4Z|e|U\alpha}{mr_0^2\omega_{rf}^2}, \quad q_x = \frac{2Z|e|V\alpha'}{mr_0^2\omega_{rf}^2}, \quad \tau = \frac{\omega_{rf}t}{2}. \quad (2.13)$$

To interject a bit of history, Mathieu investigated this equation in 1868 while studying the vibrations in an elliptical drum. The question then arises: for what values of a and q can we expect stable solutions to the Mathieu equations? To answer this question we turn

to Floquet Theory, which provides a very general approach to linear differential equations with periodic coefficients. Let x be an $n \times 1$ column vector and A be an $n \times n$ matrix with time varying coefficients that has a period T . Floquet theory is concerned with the following equation (Kovacic et al., 2018)

$$\frac{dx}{dt} = A(t)x, \quad A(t+T) = A(t). \quad (2.14)$$

Now let $X(t)$ be a fundamental solution matrix to the above equation, $X(t)$ is then a $n \times n$ matrix with each column consisting of a linearly independent solution vector. Without loss of generality we choose a basis such that $X(0) = I$. With this choice we can see

$$x(t) = X(t)x(0) \equiv C(t)x(0) \quad (2.15)$$

where we have defined C as the monodromy matrix. Because the function is periodic in time, the differential equation is invariant under transformation from t to $t+T$, thus this results in a new fundamental solution matrix that we can write in term of the other

$$x(mT) = C^m x(0) \quad m \in \mathbb{N}. \quad (2.16)$$

Let λ_i be the eigenvalues of C ,

$$Cv_i = \lambda_i v_i. \quad (2.17)$$

If $x(0) = v_i$ is chosen along an eigenvector, then after m periods

$$x(mT) = C^m v_i = \lambda_i^m v_i. \quad (2.18)$$

We can then rewrite our eigenvalues as

$$\lambda_i = e^{i\mu_i T}. \quad (2.19)$$

Next to find the independent linear solutions we can define

$$p_i(t) \equiv e^{-\mu_i t} x_i(t) = e^{-\mu_i t} X(t) v_i. \quad (2.20)$$

Then

$$p_i(t+T) = e^{-\mu_i(t+T)} x_i(t+T) = e^{-\mu_i t} e^{-\mu_i T} \lambda_i x_i(t) = e^{-\mu_i t} x_i(t) = p_i(t), \quad (2.21)$$

Hence $p_i(t)$ is periodic with period T , and the solution admits the Floquet form

$$x_i(t) = p_i(t) e^{\mu_i t}, \quad p_i(t+T) = p_i(t). \quad (2.22)$$

From this expression, we obtain the basic stability criterion: if there exists an i such that $|\lambda_i| > 1$, the solutions grow exponentially with the number of periods, and the motion is unstable. If $|\lambda_i| < 1$ for all i , the solutions decay toward the origin. Finally, if $|\lambda_i| = 1$ for all i , the solutions remain bounded.

Now we can return to the specific case of the Mathieu equations where we will define the state variables:

$$x_1 = x, \quad x_2 = \frac{dx}{dt}, \quad A(t) = \begin{pmatrix} 0 & 1 \\ -(-a_x - 2q_x \cos(\omega_{rf}t)) & 0 \end{pmatrix}. \quad (2.23)$$

Next, a fundamental solution matrix is constructed from two linearly independent solution vectors. Let these solutions be denoted by

$$x^{(1)}(t) = \begin{pmatrix} x_{11}(t) \\ x_{21}(t) \end{pmatrix}, \quad x^{(2)}(t) = \begin{pmatrix} x_{12}(t) \\ x_{22}(t) \end{pmatrix}, \quad (2.24)$$

and the monodromy matrix is

$$C = X(T) = X(\pi) = \begin{pmatrix} x_{11}(\pi) & x_{12}(\pi) \\ x_{21}(\pi) & x_{22}(\pi) \end{pmatrix}. \quad (2.25)$$

We can thus solve for the eigenvalues,

$$\lambda^2 - (\text{Tr}(C))\lambda + \det(C). \quad (2.26)$$

We can also note that C has the special property that $\det C = 1$. This is easily proven by noting $\frac{d(\det(c))}{dt} = 0$, and $W(0) = 1$. Therefore, the roots can be expressed as,

$$\lambda_{1,2} = \frac{\text{Tr}(C) \pm \sqrt{\text{Tr}(C^2) - 4}}{2}. \quad (2.27)$$

From here we can see that if $|\text{Tr}(C)| < 2$, the roots will be a pair of complex conjugates. The product of the roots must be unity because $\det(C) = 1$, thus the modulus of both is unity.

A general form in terms of the two unique linear Floquet solutions can be written as,

$$x(\tau) = A e^{i\beta_x \tau} \sum_{n=-\infty}^{\infty} C_{2n} e^{i2n\tau} + B e^{-i\beta_x \tau} \sum_{n=-\infty}^{\infty} C_{2n} e^{-i2n\tau}. \quad (2.28)$$

This is our guess solution which we have proven is the correct form. Now our goal is to solve for the coefficients C_n and β . We can do this by plugging in our general solution to the Mathieu equation and solving for a recursion relation. For simplicity we need only take one of the conjugate branches of the general solution, thus define

$$S(\tau) \equiv \sum_{n=-\infty}^{\infty} C_{2n} e^{i2n\tau}, \quad \text{so that} \quad y(\tau) = e^{i\beta_x \tau} S(\tau). \quad (2.29)$$

$$\frac{dy}{d\tau} = e^{i\beta_x\tau} \left(i\beta_x S(\tau) + \frac{dS}{d\tau} \right), \quad (2.30)$$

$$\frac{d^2y}{d\tau^2} = e^{i\beta_x\tau} \left(-\beta_x^2 S(\tau) + 2i\beta_x \frac{dS}{d\tau} + \frac{d^2S}{d\tau^2} \right). \quad (2.31)$$

$$\frac{dS}{d\tau} = \sum_{n=-\infty}^{\infty} (i2n) C_{2n} e^{i2n\tau}, \quad \frac{d^2S}{d\tau^2} = \sum_{n=-\infty}^{\infty} (i2n)^2 C_{2n} e^{i2n\tau}. \quad (2.32)$$

$$\begin{aligned} \frac{d^2y}{d\tau^2} &= e^{i\beta_x\tau} \sum_{n=-\infty}^{\infty} [-(\beta_x + 2n)^2] C_{2n} e^{i2n\tau} \\ &= - \sum_{n=-\infty}^{\infty} (\beta_x + 2n)^2 C_{2n} e^{i(\beta_x + 2n)\tau}. \end{aligned} \quad (2.33)$$

Then

$$\begin{aligned} &\left(a + 2q \cos(2\tau) \right) x(\tau) \\ &= e^{i\beta_x\tau} \left[a \sum_{n=-\infty}^{\infty} C_{2n} e^{i2n\tau} + q \left(e^{i2\tau} + e^{-i2\tau} \right) \sum_{n=-\infty}^{\infty} C_{2n} e^{i2n\tau} \right] \\ &= e^{i\beta_x\tau} \left[\sum_{n=-\infty}^{\infty} a C_{2n} e^{i2n\tau} + q \sum_{n=-\infty}^{\infty} C_{2n} e^{i2(n+1)\tau} + q \sum_{n=-\infty}^{\infty} C_{2n} e^{i2(n-1)\tau} \right]. \end{aligned} \quad (2.34)$$

We can then shift indices to express everything in the basis $e^{in\omega_{rf}t}$

$$\sum_{n=-\infty}^{\infty} C_{2n} e^{i2(n+1)\tau} = \sum_{n=-\infty}^{\infty} C_{2(n-1)} e^{i2n\tau}, \quad \sum_{n=-\infty}^{\infty} C_{2n} e^{i2(n-1)\tau} = \sum_{n=-\infty}^{\infty} C_{2(n+1)} e^{i2n\tau}. \quad (2.35)$$

And arrive at

$$\left(a + 2q \cos(2\tau) \right) x(\tau) = e^{i\beta_x\tau} \sum_{n=-\infty}^{\infty} \left[a C_{2n} + q(C_{2(n-1)} + C_{2(n+1)}) \right] e^{i2n\tau}. \quad (2.36)$$

$$\frac{d^2x}{d\tau^2} = e^{i\beta_x\tau} \sum_{n=-\infty}^{\infty} [-(\beta_x + 2n)^2] C_{2n} e^{i2n\tau}. \quad (2.37)$$

Substituting Eqs. (2.36) and (2.37) into Eq. (2.12) and cancelling the common factor $e^{i\beta_x\tau}$ yields

$$\sum_{n=-\infty}^{\infty} \{-(\beta_x + 2n)^2 C_{2n} + a C_{2n} + q(C_{2(n-1)} + C_{2(n+1)})\} e^{i2n\tau} = 0. \quad (2.38)$$

Since the functions $\{e^{i2n\tau}\}$ are linearly independent, each coefficient must vanish and we can thus write the recursion relation. For simplicity, let us define

$$D_{2n} \equiv \frac{(\beta_x + 2n)^2 - a}{q}. \quad (2.39)$$

Next we write the relationship between the coefficients

$$D_{2n} C_{2n} = C_{2(n-1)} + C_{2(n+1)}. \quad (2.40)$$

$$D_{2n} = \frac{C_{2(n-1)}}{C_{2n}} + \frac{C_{2(n+1)}}{C_{2n}} \quad D_{2(n+1)} = \frac{C_{2n}}{C_{2(n+1)}} + \frac{C_{2(n+2)}}{C_{2(n+1)}}. \quad (2.41)$$

$$\frac{C_{2(n+1)}}{C_{2n}} = \frac{1}{D_{2(n+1)} - \frac{C_{2(n+2)}}{C_{2(n+1)}}} \quad \frac{C_{2(n-1)}}{C_{2n}} = \frac{1}{D_{2(n-1)} - \frac{C_{2(n-2)}}{C_{2(n-1)}}} \quad (2.42)$$

$$\frac{C_{2(n+1)}}{C_{2n}} = \frac{1}{D_{2(n+1)} - \frac{1}{D_{2(n+2)} - \dots}} \quad \frac{C_{2(n-1)}}{C_{2n}} = \frac{1}{D_{2(n-1)} - \frac{1}{D_{2(n-2)} - \dots}} \quad (2.43)$$

$$D_{2n} = \frac{1}{D_{2(n-1)} - \frac{1}{D_{2(n-2)} - \dots}} + \frac{1}{D_{2(n+1)} - \frac{1}{D_{2(n+2)} - \dots}} \quad (2.44)$$

At this point we must choose what order to truncate our solution. To begin let's examine the first order solution. Let $C_{\pm 4} = 0$, and substitute in to solve for β

$$\begin{aligned}\frac{\beta_x^2 - a}{q} &= \frac{1}{\frac{(\beta_x + 2)^2 - a}{q}} + \frac{1}{\frac{(\beta_x - 2)^2 - a}{q}} \\ \beta_x^2 - a &= q^2 \left[\frac{1}{(\beta_x + 2)^2 - a} + \frac{1}{(\beta_x - 2)^2 - a} \right].\end{aligned}\tag{2.45}$$

$$(\beta_x \pm 2)^2 - a = \beta_x^2 \pm 4\beta_x + 4 - a \simeq 4 - a,\tag{2.46}$$

so Eq. (2.45) becomes

$$\beta_x^2 - a \simeq q^2 \left(\frac{1}{4 - a} + \frac{1}{4 - a} \right) = \frac{2q^2}{4 - a}.\tag{2.47}$$

$$\beta_x^2 \simeq a + \frac{2q^2}{4 - a} \simeq a + \frac{q^2}{2}, \quad \implies \quad \beta_x \simeq \sqrt{a + \frac{q^2}{2}}.\tag{2.48}$$

This is the general result given for the small a and q parameter approximation. Now we can note that,

$$C_{\pm 2} = \frac{q C_0}{(\beta_x \pm 2)^2 - a} \simeq \frac{q}{4} C_0.\tag{2.49}$$

This leads to an equation for the motion of a single ion,

$$x(t) \simeq 2AC_0 \cos\left(\beta_x \frac{\omega_{rf}}{2} t\right) \left[1 - \frac{q_x}{2} \cos(\omega_{rf} t) \right].\tag{2.50}$$

In this form it is clear there are three characteristic frequency terms. The first at $\beta\omega_{rf}$, which typically has the largest amplitude, is known as the secular motion. The other two motional frequencies occur at $(\omega_{rf}(1 \pm \beta))$. These are known as the intrinsic micromotion. We can observe that the ratio of the secular motion to the intrinsic micromotion is given by $\frac{q_x}{2}$.

In this work, the typical trapping parameters are $V = 200$ V, $\omega_{rf} = 2\pi \times 32$ MHz, and $r_0 = 200$ μm . Thus in the small parameter approximation this yields a secular




Figure 2.1: Fourier spectrum of the simulated ion trajectory $x(t)$ using the 5-term Floquet truncation.

frequency of $\omega_{sec} = \beta_x \frac{\omega_{rf}}{2} = 2.0302$ MHz, while in the non-approximate equation we arrive at $\beta_x \frac{\omega_{rf}}{2} = 2.042$ MHz, demonstrating that this is a $\approx .6\%$ effect. Higher order truncation results in more resonant frequencies of decreasing amplitude as shown (to 5th order terms) in Figure 2.1.

2.3 Determining Parameters of Stable Motion

The benefit of having an analytical expression for the ion motion, is that we can easily determine the conditions under which stable and unstable trapping occurs. This also allows for the fine tuning of parameters such that only specific charge to mass ratios are trapped while others are expelled. Utilizing our solution to the Mathieu equation, and noting that the general solution is made up of two linearly independent terms we see,

$$x(t) \simeq \left[1 - \frac{q_x}{2} \cos(\omega_{rf}t) \right] \left[A \cos(\omega_{sec}t) + B \sin(\omega_{sec}t) \right]. \quad (2.51)$$

From this we can use the initial conditions

$$u(0) = 1, \quad \dot{u}(0) = 0, \quad v(0) = 0, \quad \dot{v}(0) = 1, \quad (2.52)$$

Thus we have the C matrix defined in Eq. (2.25)

$$X(T) = X\left(\frac{2\pi}{\omega_{rf}}\right) = C = \begin{pmatrix} u(t) & v(t) \\ \dot{u}(t) & \dot{v}(t) \end{pmatrix} = \begin{pmatrix} \cos(\omega_{\text{sec}}T) & \frac{1}{\omega_{\text{sec}}} \sin(\omega_{\text{sec}}T) \\ -\omega_{\text{sec}} \sin(\omega_{\text{sec}}T) & \cos(\omega_{\text{sec}}T) \end{pmatrix}. \quad (2.53)$$

Now taking the trace of this matrix we can write the roots of the characteristic polynomial using Eq. (2.27)

$$\text{Tr } C = 2 \cos(\beta_x \pi). \quad (2.54)$$

$$\lambda_{\pm} = \cos(\beta\pi) \pm \sqrt{\cos^2(\beta\pi) - 1} = \cos(\beta\pi) \pm i \sin(\beta\pi) = e^{\pm i\beta\pi}. \quad (2.55)$$

The form of this solution makes it clear that unstable motion occurs when β has a nonzero imaginary component. While this equation is valid for small a and q parameters, if we desire to generate a stability diagram over a wide range of values, we must employ more complex mathematical methods. Instead of attempting to numerically integrate the differential equations, or algebraically solve the recursion relation to an arbitrarily high order, we can make some qualitative observations about the form of the Mathieu equation. We already have established that if the roots λ_i do not lie on the unit circle, unstable motion will result. We now want to find functions $a_i(q)$, that establish boundaries between stable and unstable motion on the a, q plane. As an example we can derive one such boundary, also known as a tongue. First let $q = 0$, and note the Mathieu equations take the form

$$x'' + ax = 0. \quad (2.56)$$

This implies the solution $A \cos(nt) + B \sin(nt)$, and thus $n = \sqrt{a}$. Now, recalling that our

driving term has frequency $\cos(2\tau)$, and thus a period of forcing of π , we can observe that if $a = 0, 1, 3, \dots, n^2$ for $n \in \mathbb{Z}$, the motion will undergo parametric resonance with either periodic or antiperiodic solutions. The characteristic value $a_0(q)$ is defined as the value of a for which there exists a π -periodic, even solution that reduces to a constant as $q \rightarrow 0$. We can use perturbation theory to solve for terms of $a_0(q)$, first we expand both the solution and the characteristic value as power series in q ,

$$x(\tau; q) = x_0(\tau) + q x_1(\tau) + q^2 x_2(\tau) + \dots, \quad (2.57)$$

$$a(q) = a_0 + a_1 q + a_2 q^2 + \dots. \quad (2.58)$$

For the lowest branch we require that, at $q = 0$, Eq. (2.56) reduces to $x_0'' + a_0 x_0 = 0$ and admits the constant solution. We therefore choose

$$x_0(\tau) = 1, \quad a_0 = 0. \quad (2.59)$$

We can then substitute these terms into our series expansion and plug them back in to the Mathieu equation, and collecting terms of order q gives

$$x_1(\tau) = -\frac{1}{2} \cos(2\tau), \quad (2.60)$$

To avoid a secular (nonperiodic) term proportional to τ^2 , the constant part of the forcing must vanish, which fixes a_1 . One can extend this method to further orders. The solution to three orders is

$$a_0(q) = -\frac{1}{2} q^2 + \frac{7}{128} q^4 + O(q^6). \quad (2.61)$$

In order to qualitatively understand this function recall by definition, $a_0(q)$ is the smallest value of a that admits a π -periodic even solution. A π -periodic solution satisfies $x(\tau + \pi) = x(\tau)$ and hence corresponds to a Floquet multiplier $\rho = +1$, implying

$$\text{Tr}(C) = \Delta(a_0(q), q) = 2. \quad (2.62)$$




Figure 2.2: Mathieu stability diagram for a range of parameters. The gold shaded regions indicate stable motion in the x-direction, while the blue shows stable z-direction motion. The zoom in region on the right panel shows the main zone of overlap that experimental ion traps choose to operate in.

To determine which side of $a_0(q)$ is stable, note that for $a \rightarrow -\infty$ the equation is dominated by $x'' - |a|x = 0$, whose solutions behave as $e^{\pm\sqrt{|a|}\tau}$. Over one period π , the corresponding multipliers behave as $e^{\pm\sqrt{|a|}\pi}$, so the discriminant satisfies

$$\Delta(a, q) \sim e^{\sqrt{|a|}\pi} + e^{-\sqrt{|a|}\pi} = 2 \cosh(\sqrt{|a|}\pi) > 2, \quad (a \rightarrow -\infty). \quad (2.63)$$

Therefore, all solutions are unstable for $a < a_0(q)$. Above this boundary line the solutions are stable until the next boundary line is met. It can be shown that periodic and antiperiodic solutions alternate and thus form the boundaries of stable regions. This summary is meant to shed light on the construction of the Mathieu stability diagrams as shown in Figure 2.2. The alternating boundaries of Mathieu characteristic functions $a_0(q)$ and $b_1(q)$ can clearly be seen. Due to the symmetry of the quadrupole potential, it is clear there is a stable trapping region along the line at $a = 0$.




Figure 2.3: A comparison of the solid angle allowed for light collection in a classical RF Paul trap with cylindrical electrodes and a blade trap. The radius of the cylindrical electrodes is optimally chosen to suppress higher order multipole terms, and the blade trap geometry is such that the angles between the blades are 120° and 60° .

2.4 Numerical Simulation of Real Trap Geometries

Ion trap technology has vastly improved in the 40 years since its inception. Modern traps strive to reach higher secular frequencies, which can enable faster two qubit gates, lower heating rates, and reduce the crowding of spectral modes. Pushing ion traps to this limit requires a delicate balancing of parameters. Our goal is to maximize the quantity

$$\omega_{sec} \approx \omega_{rf} \sqrt{a + \frac{q^2}{2}} \propto \frac{V}{r_0^2 \omega_{rf}^2} \quad (2.64)$$

while keeping q safely within the stable region of $0 < q \propto \frac{V}{r_0^2 \omega_{rf}^2} < .9$. One might desire to drastically increase the Voltage V , but there are typically technical limitations to this. Higher radio-frequency (RF) voltages can result in heating at resistive junctions, which can cause outgassing as well as dielectric breakdown, resulting in arcing. The general strategy of many has been to move the ion closer to the electrode, reducing r_0^2 , while simultaneously increasing ω_{rf} in order to maintain a fixed q parameter. This strategy has also proven to have complications; namely, increased heating rates as the ion is placed closer to the surface. Another key issue with smaller ion traps is that as the electrodes are

moved closer they begin to occlude the solid angle of scattered light from the ion, thereby reducing light collection efficiency. Due to these reasons, two types of trap geometries have found wide adoption. The first is a blade trap design, which uses much thinner electrodes than the typical cylindrical Paul traps used in the past, and allows for a much larger solid angle for light collection along two different access points as shown in Figure 2.3. The second is a lithographically etched surface electrode trap; these allow for arbitrary electrode mapping, and a large solid angle, but typically only from one access point. In this thesis, a cylindrical four-rod trap, a blade trap, and an integrated surface-electrode trap are used.

To understand how these differing geometries affect the ions we can numerically solve for the electrostatic potential, and we will focus on the blade trap design. In the charge free region between the trap electrodes the electrostatic potential satisfies Laplace's equation,

$$\nabla^2 \phi(x, y) = \frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} = 0. \quad (2.65)$$

To solve this numerically on a 2D cross section, we discretize the plane on a uniform Cartesian grid

$$x_i = x_0 + ih, \quad y_j = y_0 + jh, \quad (2.66)$$

and denote $\phi_{i,j} \equiv \phi(x_i, y_j)$. The second derivatives are approximated by standard central differences,

$$\left. \frac{\partial^2 \phi}{\partial x^2} \right|_{(i,j)} \approx \frac{\phi_{i+1,j} - 2\phi_{i,j} + \phi_{i-1,j}}{h^2}, \quad (2.67)$$

$$\left. \frac{\partial^2 \phi}{\partial y^2} \right|_{(i,j)} \approx \frac{\phi_{i,j+1} - 2\phi_{i,j} + \phi_{i,j-1}}{h^2}. \quad (2.68)$$

Substituting this into Laplace's equation allows us to solve for adjacent grid positions potential,

$$\phi_{i+1,j} + \phi_{i-1,j} + \phi_{i,j+1} + \phi_{i,j-1} - 4\phi_{i,j} = 0. \quad (2.69)$$




Figure 2.4: The numerical solutions to Laplace's equation for a blade trap

This along with the boundary conditions that the two of the blades are uniformly at a voltage V while the other two are at ground yields a static potential shown in Figure 2.4. We can compare this numerical solution to the case of the ideal hyperbolic trap.

To quantify the degree to which the numerically computed blade-trap potential resembles an ideal quadrupole near the RF null, we expand the solution $\phi(x, y)$ to second order about the point (x_*, y_*) where the RF field magnitude $|\nabla\phi|$ is minimized,

$$\phi(x, y) \approx \phi_* + \frac{1}{2} \begin{pmatrix} \Delta x & \Delta y \end{pmatrix} H \begin{pmatrix} \Delta x \\ \Delta y \end{pmatrix}, \quad \Delta x \equiv x - x_*, \quad \Delta y \equiv y - y_*, \quad (2.70)$$

where the Hessian matrix is

$$H \equiv \begin{pmatrix} \phi_{xx} & \phi_{xy} \\ \phi_{xy} & \phi_{yy} \end{pmatrix}_{(x_*, y_*)}. \quad (2.71)$$

One parameter used to evaluate the strength of the confinement is the quadrupole curvature,

$$\kappa_{\text{eff}} \equiv \frac{1}{4} (\phi_{xx} - \phi_{yy}), \quad (2.72)$$

For an ideal hyperbolic quadrupole with field radius r_0 and unit electrode voltage, the analytic solution is

$$\kappa_{\text{hyperbola}} = \frac{1}{2r_0^2}. \quad (2.73)$$

The ratio $\kappa_{\text{blade}}/\kappa_{\text{hyperbola}}$ therefore quantifies the quadrupole efficiency of the blade geometry for the chosen boundary conditions, or the curvature achieved near the RF null per applied volt. This solution yields a ratio of curvature of $\approx .02$, meaning the confinement is roughly 2% as efficient in the blade trap as compared to a perfect hyperbolic electrode trap. This can mostly be accounted for by realizing that the hyperbolic electrodes are infinitely large, while the effective electrode surface area in the blade trap is much smaller.

A related metric is the trap depth which can also be calculated from this numerical solution. The trap depth can be defined as the time averaged potential energy difference between the center of the trap and the lowest energy saddle point. We can calculate the potential energy of the ion by utilizing the pseudopotential approximation. The ion's motion can be modeled using

$$m\ddot{\mathbf{r}} = q\mathbf{E}(\mathbf{r}, t), \quad \mathbf{r}(t) = \mathbf{r}_{\text{sec}}(t) + \mathbf{r}_{\text{mm}}(t), \quad (2.74)$$

where we have assumed the trajectory can be separated into slow secular motion and fast micromotion. To leading order and treating $\mathbf{r}_{\text{sec}}$ as constant over one RF cycle, the

micromotion satisfies

$$m\ddot{\mathbf{r}}_{mm} \approx q \mathbf{E}_{rf}(\mathbf{r}_{sec}) \cos(\omega_{rf}t), \quad (2.75)$$

with solution

$$\mathbf{r}_{mm}(t) \approx -\frac{q}{m\omega_{rf}^2} \mathbf{E}_{rf}(\mathbf{r}_{sec}) \cos(\omega_{rf}t). \quad (2.76)$$

Substituting back and averaging over one RF period, the $\cos(\Omega t)$ terms vanish while $\langle \cos^2(\Omega t) \rangle_T = 1/2$ produces a nonzero slow force of the form,

$$m\ddot{\mathbf{r}}_{sec} = \frac{q^2}{4m\omega_{rf}^2} |\mathbf{E}_{rf}(\mathbf{r}_{sec})|^2 = -\nabla U(\mathbf{r}_{sec}). \quad (2.77)$$

This definition gives us the potential energy of the trapping field. Note, for an ideal hyperbolic quadrupole trap the pseudopotential, and thus, the confinement is symmetric about the trap null. In contrast, a blade trap with asymmetric angular gaps has an asymmetry in the strength of the electric field as is clearly seen in Figure 2.4. The depth of the trap is thus the energy required to escape in the direction of the weakest confinement (toward the 120° opening). Using the typical experimental trapping parameters, a trap depth in this direction is calculated to be 1.25 eV, and the distance to the trap center at which point the ion escapes is 170 μm . This is an important number to understand the probability of trapping hot atoms during the loading process. This corresponds to a Barium ion at $\approx 14,000$ K or traveling with a $v_{rms} = 1.3$ km/s.

Another common equation utilized for trap depth is

$$D = \frac{1}{2} \omega_{sec}^2 r_0^2$$

. The intuition for this equation is essentially the energy of the particle when its orbit is large enough to strike an electrode. This method yields a trap depth of 10 eV, a much larger and likely inaccurate increase. Note the secular frequency is also easily approximated by calculating the gradient of the potential energy at the trap null. This method yields a secular frequency of 3 MHz, which is remarkably close to the value of 2.9 MHz that is

experimentally observed.

Another important difference between the blade trap geometry and the hyperbolic square electrode trap is the orientation of the principal axes (Johnson, 2016). The principal axes are the orthogonal vectors that are completely decoupled from one another in the ion's motion. In a square hyperbolic electrode trap these axes are always along the direction of the center of the electrodes. This is typically convenient for cooling as a laser aligned through the plane of the middle of the trap can address both of these modes of motion. In the blade trap if the DC blades are held at ground the principal axes exactly bisect the gaps of the blades. This can cause issues in cooling and coupling to the motion of the axes that is perpendicular to the central plane which is the only practical access for lasers. To ameliorate this issue the principal axes must be rotated by applying a static DC dipole potential to two opposite blades.

2.5 Excess Micromotion

We have thus far defined the secular motion and the *intrinsic* micromotion of the ion. However, an important and unwanted component of the motion known as the excess micromotion can also play a role in the ion's motion and contribute significantly to the heating rate. Excess micromotion is caused by a uniform static electric field, which recasts the equations of motion into the form of the inhomogeneous Mathieu equations (Berkeland et al., 1998)

$$\frac{d^2x}{d\tau^2} + \left[a_x - 2q_x \cos(2\tau) \right] x = \frac{4Z|e|}{m\omega_{rf}^2} E_x. \quad (2.78)$$

One can imagine that this modifies the solution to our homogeneous equation by pushing the ion with a static field away from the stable equilibrium point and towards a region of the potential with higher curvature. Because of the fact that the ion's average position is now no longer at the point where the RF field cancels out, it experiences an oscillation at the frequency of the RF driving field. We have already found the general solution to the homogeneous Mathieu equations. Thus, to find the general solution to this new

inhomogenous version we need only one particular solution to the equation, such as

$$x_p(t) \approx x_0 + \frac{q_x}{2} x_0 \cos(\omega_{rf} t), \quad x_0 = \frac{Z|e|}{m\omega_{sec}^2} E_x, \quad (2.79)$$

so that a compact expression for the full trajectory is

$$x(t) \simeq A \cos(\omega_{sec} t) \left[1 - \frac{q_x}{2} \cos(\omega_{rf} t) \right] + x_0 \left[1 + \frac{q_x}{2} \cos(\omega_{rf} t) \right]. \quad (2.80)$$

From this solution it is clear there are two effects on the ions motion: one the equilibrium is shifted by x_0 ; and two a new frequency in the Fourier spectrum is introduced at ω_{rf} with amplitude $\propto x_0 q_x$. To understand the amplitude of this motion, we can apply 1 V to a single blade and numerically compute the field at the ions position, yielding a $E_{ex-mm} \approx 1,800$ V/m. To give some perspective, this would correspond to a micromotion amplitude of $\approx .7 \mu\text{m}$, while the typical secular amplitude at the Doppler cooling limit is ≈ 17 nm.

Excess micromotion can have numerous deleterious effects on an ion trap experiment. Spectroscopy experiments experience a large second order Doppler shift due to the increase in the ion's mean square velocity. In quantum computing applications, micromotion can increase the secular heating rate, reducing gate fidelity. In an ideal quadrupole potential, it is not immediately clear that the reservoir of excess micromotion energy could transfer to secular energy, and thus heat the ion's secular modes.

Unfortunately, in real-world experiments multiple avenues for this type of heating can occur. The first mechanism occurs if the RF voltage has noise at or near $2\omega_{sec}$, this results in a modulation of the trapping field and causes a parametrically driven excitation. The second mechanism can be understood by noting that as the micromotion amplitude increases, the ion samples regions of the trap potential where higher-order multipole terms become significant. Cross terms in this potential create many more frequencies of oscillation and increase the probability that secular modes will overlap with micromotion modes, causing heating. Finally, all experiments using Doppler cooling must take into

account how the first-order Doppler shift due to excess micromotion modifies the line shape of the cooling transition, as shown in the middle panel of Figure 2.5.

Understanding the excitation spectrum with the modification of excess micromotion is also important for most methods of micromotion detection. We can write the electric field in the reference frame of the ion (ignoring secular motion) as,

$$E(t) \propto e^{i(\omega_L t - \mathbf{k} \cdot \mathbf{r}_{mm}(t))} \quad \mathbf{k} \cdot \mathbf{r}_{mm}(t) = \beta_{mm} \cos(\omega_{rf} t), \quad \beta_{mm} \equiv \mathbf{k} \cdot \mathbf{x}_{mm}. \quad (2.81)$$

We can apply a fun math trick to expand this exponential called the Jacobi–Anger identity,

$$e^{-i\beta_{mm} \cos(\omega_{rf} t)} = \sum_{n=-\infty}^{\infty} (-i)^n J_n(\beta_{mm}) e^{-in\omega_{rf} t}, \quad (2.82)$$

where J_n are Bessel functions. The point of this expansion is to observe that the electric field now has added frequency components (in the ion’s reference frame) with modulation strength characterized by β_{mm} . We can use this to write the modified excitation spectrum at different magnitudes of micromotion, given the solution to the excited state optical Bloch equations

$$\bar{\rho}_{ee}(\Delta) = \frac{s_0}{2} \sum_{n=-\infty}^{\infty} \frac{J_n^2(\beta)}{1 + \left[\frac{2(\Delta + n\omega_{rf})}{\Gamma} \right]^2}. \quad (2.83)$$

We can see from Figure 2.5, if the micromotion is sufficiently large, the excited state population is depressed. This effect results in a lower scattering rate and, thus, less efficient detection of an ion’s internal state. Moreover, it can directly reduce the Doppler cooling limit, which is proportional to the slope of the cooling line shape.

Periodically, an experimental ion trap accumulates charges, either through the process of loading, where laser ablation causes plumes of plasma that coat the electrodes and dielectrics, or through the photoelectric effect, where lasers used to address the ion rip electrons off of the metal surface. This creates static electric fields that must be compensated for, requiring constant recalibration of micromotion for high-quality experiments.




Figure 2.5: The excitation spectra of a two level system with the addition of the ion’s micromotion. All graphs show a trap drive of 32 MHz, the three panels correspond to different transition linewidths commonly seen in this experiment. The middle panel is the typical regime of the Doppler cooling transition in Barium, while the right panel resembles the narrower quadrupole transition to the $D_{5/2}$ state. The left panel resembles a situation in which the linewidth is much larger than the RF drive frequency.

The first step in correcting this static electric field is to detect it, which can be done by observing changes in the excitation spectrum or the physical position of the ion. We will describe four methods to detect micromotion and discuss the pros and cons of each.

The first method utilizes the camera to image the ion as the RF drive voltage is lowered and raised, typically on the order of 4 dB (60 percent change). The principle behind this method is that as the trap voltage is relaxed the curvature of the trapping potential at a particular location is reduced. This change in curvature is also proportional to the distance the ion is from the RF null. Using this method, compensation electrodes can be adjusted to the point that when the RF is raised and lowered the ion does not move. The limits of this method are the resolution and magnification of the camera imaging system. This method also only works to compensate for micromotion orthogonal to the imaging plane.

The second method utilizes a narrow linewidth quadrupole transition and relies on measuring the Rabi rate while detuned to the RF micromotion sideband. This method is limited particularly in Barium by the low lamb-dicke factor of the quadrupole transition, which results in a small motional transition Rabi rate. It should be noted that this also only gives amplitude information about the static electric field and not direction. Finally,

it requires three orthogonal lasers to address all the modes of motion which is another significant experimental detractor. This methodology is compared to others in (Hogle et al., 2024).

The third method, known as the cross-correlation technique, measures the fluorescence modulation of the Doppler cooling transition. As we have seen the first order Doppler shift modulates the frequency spectrum of the ion. This analogously appears in the time domain when scattered photons from this excitation are collected. To visualize this, note that as the ion moves at the RF frequency, its velocity-dependent Doppler shift due to micromotion is given by

$$\mathbf{k} \cdot \mathbf{r}'_{mm}(t) = \omega_{rf} \beta_{mm} \sin(\omega_{rf} t). \quad (2.84)$$

Because the scattering rate is detuning dependent, we can simply write the fluorescence signal as

$$R(t) = \frac{\Gamma}{2} \frac{s_0}{1 + s_0 + \left[\frac{2}{\Gamma} (\Delta_0 - \omega_{rf} \beta_{mm} \sin(\omega_{rf} t)) \right]^2}. \quad (2.85)$$

In practice, the fluorescence is detected, the photons are time tagged and the data collection is triggered with the actual RF voltage signal. Thus successive experiments are started at the same phase. Importantly, this detection method is sensitive to the direction of the offset field. When the ion moves past the RF null, the sign of x_0 (and β) changes, this manifests as a 90° phase shift in the micromotion velocity, which can then be seen in the cross correlation measurement. This allows for a much faster binary search to determine the correct compensation voltages. Another advantage stemming from this phase sensitivity is that this method can detect a phase mismatch between the opposing RF blades. This will typically result in a continuous change of phase in the cross correlation as the RF null is swept across as opposed to a sharp change in phase. The downside of this method is the requirement to have three orthogonal laser beams in order to compensate micromotion in every direction.

The fourth and final method we will discuss is known as parametric excitation, which

has the major benefit of only requiring a single laser propagation path. However, this is assuming that the path has a projection along all three principal axes (Nadlinger et al., 2021). This is not a very stringent requirement as usually the Doppler cooling laser used in an experiment already has this quality. This feature is also particularly useful for surface electrode ion traps, which typically can not address the mode of motion perpendicular to the surface of the trap. The basic principle of this method is to modulate the amplitude of the trapping voltage. As we have noted in the above methods, when the RF voltage is changed, if the ion is not at the RF null it will experience a force. This is due to the position dependent change in curvature of the electric potential.

In practice, this modulation is created by splitting off the RF signal, and mixing in another signal that is swept over a particular mode of motions secular frequency. We can qualitatively understand what this does to the ion's motion by using the pseudopotential approximation and focusing only on the secular motion of the ion,

$$m\ddot{x} = -m\omega_{s0}^2(t)x + x_0. \quad (2.86)$$

If the RF amplitude is modulated at frequency ω_m , the secular frequency becomes time-dependent. For small modulation depth $h \ll 1$,

$$\omega_s^2(t) = \omega_{s0}^2 [1 + 2h \cos(\omega_m t)], \quad (2.87)$$

where ω_{s0} is the unmodulated secular frequency. Thus, we can see the secular frequency is now changing at our modulation frequency. Now noting that there will be a static displacement x_0 to the ion equilibrium position, we can substitute $x(t) = x_0 + u(t)$ and $\omega_s^2(t)$ into the equation of motion that gives

$$\ddot{u} + \omega_{s0}^2 u + 2h\omega_{s0}^2 \cos(\omega_m t) u = -2h\omega_{s0}^2 x_0 \cos(\omega_m t). \quad (2.88)$$

Equation (2.88) contains two oscillating driving terms. The first term, proportional to




Figure 2.6: A heat map showing the phase and intensity of the demodulated fluorescence signal. Each square shows intensity at the RF modulation frequency (x-axis) and the color indicates the phase of the demodulation signal. As the compensation voltages are swept and the frequency is scanned the row of data points with the smallest minimum value indicates the smallest DC offset electric field for this particular mode.

u , is a parametric drive which therefore has a resonance at $2\omega_{s_0}$. The second term is a regular driven oscillator force, and produces driven motion that is resonant at ω_{s_0} and importantly is proportional to the modulation index and x_0 .

From the above derivation we can see that if either the modulation strength or the micromotion amplitude is increased, the amplitude of the motion at the secular frequency will be increased proportionally. This motion is actively damped by Doppler cooling, during which the scattered light can be collected and demodulated to determine the micromotion amplitude and phase (see the discussion of the cross-correlation method above). It should be noted that when the driving term is too large, non-linear effects can take place causing unstable motion and reducing the signal.

Figure 2.6 shows an example of an experimental micromotion compensation scan. The optimal offset voltages can be determined by noting the point at which the phase flips and the row where the smallest minimum demodulation intensity is measured. One downside of this method can also be seen in this graph. As the micromotion offset x_0 changes, so does the secular frequency; thus, the modulation frequency must be scanned over a fairly

wide range. After taking a scan one must then compensate the orthogonal radial mode and iteratively return to the previous mode until a convergence of compensation voltages is found.

A simple algorithm can be used to speed up this process if compensation electrode voltages are applied to orthogonal directions. First one finds a voltage that minimizes the micromotion along a chosen mode x_1 ; then with this voltage set, a second minimum along the orthogonal mode is found y_1 ; and finally a third voltage x_2 is found with y_1 set. These three points, can then be used to find two orthogonal lines whose intersect is the optimal compensation voltage for both modes x_{opt}, y_{opt} according to the equations,

$$m_1 = \frac{y_1}{x_2 - x_1}, \quad b_1 = -m_1 x_1 \quad m_2 = -\frac{1}{m_1}, \quad b_2 = y_1 - m_2 x_2 \quad (2.89)$$

$$x_{opt} = \frac{b_2 - b_1}{m_1 - m_2}, \quad y_{opt} = m_1 x_{ad} + b_1 \quad (2.90)$$

2.6 Ion Heating Rates

Superficially, an ion isolated in a trap and suspended in an ultrahigh-vacuum environment appears to be completely isolated from the outside world. If this were actually true, it would be the ideal canvas for a quantum computer. One could imagine mapping a quantum state onto an ion, and leaving it undisturbed for hours without a single degree of freedom changing and thus disturbing a perfect superposition. Unfortunately, this is not the reality we live in. One key reason for this is the charge. The powerful lever that allows the ion to be so tightly confined and gives us the ability to isolate and manipulate its position in space, is also what makes it a sensor of any outside noise.

Heating of an ion occurs when unwanted noise appears on an electrode that has a frequency component that overlaps with one of the ion's motional modes. One could imagine taking the overlap integral of these two spectrum's and then arriving at a total rate of excitation from an ion's motional state to the next state (assuming the ion is not

moving by more than one excitation at a time). This qualitative description can be made rigorous through the following method (Turchette et al., 2000). First, we begin with a simple Hamiltonian modeling an ion as a simple harmonic oscillator with a perturbation due to an electric field

$$H = \frac{p^2}{2M} + \frac{1}{2}M\omega_z z^2/2 - q\epsilon(t)z \quad (2.91)$$

where $\epsilon(t)$ is the electric field drive noise. Thus our perturbation Hamiltonian is $H'(t) = -q\epsilon(t)z$. Second, we calculate the transition rate to the first order, beginning with Fermi's golden rule (first order perturbation theory)

$$R_{m \leftarrow n} = \frac{1}{T} \left| \frac{-i}{\hbar} \int_0^T dt' H'_{mn}(t') e^{i\omega_{mn}t'} \right|^2 \quad (2.92)$$

$$= \frac{1}{T} \frac{1}{\hbar^2} \left(\int_0^T dt_1 H'_{mn}(t_1) e^{i\omega_{mn}t_1} \right)^* \left(\int_0^T dt_2 H'_{mn}(t_2) e^{i\omega_{mn}t_2} \right) \quad (2.93)$$

$$= \frac{q^2}{T\hbar^2} |\langle m|z|n \rangle|^2 \int_0^T dt_1 \int_0^T dt_2 \epsilon(t_1)\epsilon(t_2) e^{i\omega_{mn}(t_2-t_1)}, \quad (2.94)$$

Third, we can note that because our electric field is being modeled as random noise, it has no preferred start time, so the correlation between functions of the two time variables only depends on the difference of the time $\tau = t_2 - t_1$, thus we can make a change of variables

$$= \frac{q^2}{T\hbar^2} |\langle m|z|n \rangle|^2 \int_0^T dt \int_{-t}^{T-t} d\tau \epsilon(t)\epsilon(t+\tau) e^{i\omega_{mn}\tau} \quad (2.95)$$

$$= \frac{q^2}{\hbar^2} |\langle m|z|n \rangle|^2 \int_{-T}^T d\tau \left(1 - \frac{|\tau|}{T} \right) e^{i\omega_{mn}\tau} \frac{1}{T} \int_0^{T-|\tau|} dt \epsilon(t)\epsilon(t+\tau). \quad (2.96)$$

Fourth, we take the limit as $T \rightarrow \infty$,

$$\frac{1}{T} \int_0^{T-|\tau|} dt \epsilon(t)\epsilon(t+\tau) \rightarrow \langle \epsilon(t)\epsilon(t+\tau) \rangle, \quad \left(1 - \frac{|\tau|}{T} \right) \rightarrow 1. \quad (2.97)$$

The above equation is simply a statement that we believe the noise is ergodic, meaning that if the measurement is taken for long enough, it is essentially equivalent to an ensemble average of many experiments. If the correlation decays, which is true under the assumption

that the noise is roughly broadband, then we can extend the limits. If we have extremely narrow band noise we must treat it as driving and this derivation does not apply. Thus our equation becomes,

$$\frac{q^2}{\hbar^2} |\langle m|z|n\rangle|^2 \int_{-\infty}^{\infty} d\tau e^{i\omega_{mn}\tau} \langle \epsilon(t)\epsilon(t+\tau) \rangle. \quad (2.98)$$

We will define the power spectral density (using the Wiener Kinchin theorem) as

$$S_E(\omega) = \int_{-\infty}^{\infty} d\tau e^{i\omega_{mn}\tau} \langle \epsilon(t)\epsilon(t+\tau) \rangle. \quad (2.99)$$

This equation holds so long as there is no long term drifts in the noise, because if that were true, there would not be a stationary power spectral density to define it. Next we can evaluate the matrix element

$$\hat{z} = \sqrt{\frac{\hbar}{2m\omega_z}}(a^\dagger + a). \quad (2.100)$$

This implies

$$R_{n\pm 1 \leftarrow n} = \frac{q^2}{4m\omega_z} S_E(\omega_z) (\delta_{n+1,n} (n+1) + \delta_{n-1,n} (n)) \quad (2.101)$$

Thus the heating rate can be written as

$$\dot{\bar{n}} = \frac{\langle \dot{E} \rangle}{\hbar\omega_z} = \sum_n P(n) (R_{n+1 \leftarrow n} - R_{n-1 \leftarrow n}) \quad (2.102)$$

$$= \frac{q^2}{4m\hbar\omega_z} \sum_n P(n) \hbar\omega_z S_E(\omega_z) (n+1 - n) = \frac{q^2}{4m\hbar\omega_z} S_E(\omega_z). \quad (2.103)$$

If the electric field noise is a gradient field the perturbation operator becomes

$$H' = m\omega_z^2 \epsilon(t) x^2. \quad (2.104)$$

This implies a transition rate

$$R_{m \leftarrow n} = \left(\frac{m^2 \omega_z^4}{2\hbar} \right) \int_{-\infty}^{\infty} d\tau e^{i\omega_{mn}\tau} \langle \epsilon(t) \epsilon(t+\tau) \rangle | \langle m | x^2 | n \rangle |^2 \quad (2.105)$$

$$= \frac{\omega_z^2 \hbar}{8} S_e(2\omega_z) (n+1 \pm 1) (n \pm 1). \quad (2.106)$$

Then the heating rate becomes

$$\dot{n} = \frac{\langle \dot{E} \rangle}{\hbar \omega_z} = \sum_n P(n) 2(R_{n+2 \leftarrow n} - R_{n-2 \leftarrow n}) \quad (2.107)$$

$$= \frac{\omega_z^2}{2} \sum_n P(n) S_E(2\omega_z) (2n+1) = \frac{\omega_z^2}{2} S_E(2\omega_z) (2\bar{n}+1). \quad (2.108)$$

At this point, we can ask: how does one actually measure the power spectral density of the electric field? One standard method is to measure the voltage-noise power spectral density on a trap electrode $S_V(\omega)$. This can be related to the electric field noise by

$$S_E(\omega) = |\beta|^2 S_V(\omega) \quad \beta \equiv \left. \frac{\partial E}{\partial V} \right|_{\text{ion}}. \quad (2.109)$$

The particular value of β for the blade trap used in this thesis can be extracted from the numerical electric potential simulations and was found to be $\beta = 3.9 \times 10^{-4} \frac{1}{m}$. In practice the measurement of $S_V(\omega)$ is typically done outside of vacuum through some sort of connection that is terminated in a 50Ω load. Thus a transfer function $H(\omega)$ should be measured before the trap is sealed, by simply sweeping a known input frequency and measuring the true voltage on the electrode at each relevant frequency. If this is not done, one can also assume a relatively flat transfer function and then measure $H(\omega_{rf})$. Figure 2.7 shows both the noise spectrum over the secular frequency range, and the power of the trap RF drive. By measuring the secular frequency we can extract the true trap $V_{RF} = 300V$. For a 50Ω load, a power P corresponds to an RMS voltage,

$$V_{\text{meas,rms}} = \sqrt{PR}, \quad R = 50 \Omega. \quad (2.110)$$




Figure 2.7: A spectrum analyzer measurement of the trap electrode. Left panel shows the relevant range of the ion's secular motion. Right panel shows the signal from the trap RF drive.

For the RF tone,

$$P_{\text{RF}} = 10^{3/10} \times 10^{-3} \text{ W} \simeq 1.995 \times 10^{-3} \text{ W}, \quad (2.111)$$

$$V_{\text{RF, meas}} = \sqrt{P_{\text{RF}} R} = \sqrt{(1.995 \times 10^{-3})(50)} \simeq 3.16 \times 10^{-1} \text{ V}. \quad (2.112)$$

For the noise level,

$$P_{\text{N}} = 10^{-80/10} \times 10^{-3} \text{ W} = 10^{-11} \text{ W}, \quad (2.113)$$

$$V_{\text{N, meas}} = \sqrt{P_{\text{N}} R} = \sqrt{(10^{-11})(50)} \simeq 2.24 \times 10^{-5} \text{ V}. \quad (2.114)$$

Assuming the same measurement-to-electrode transfer applies to nearby frequencies, the voltage scaling factor is

$$G \equiv \frac{V_{\text{elec}}}{V_{\text{meas}}} = \frac{V_{\text{RF, elec}}}{V_{\text{RF, meas}}} = \frac{300}{0.316} \simeq 9.49 \times 10^2. \quad (2.115)$$

Thus the inferred RMS electrode noise voltage is

$$V_{\text{N, elec}} = G V_{\text{N, meas}} \simeq (9.49 \times 10^2)(2.24 \times 10^{-5}) \simeq 2.13 \times 10^{-2} \text{ V}. \quad (2.116)$$

With $\beta = 3.9 \times 10^{-4} \text{ m}^{-1}$, the RMS electric-field noise amplitude is

$$E_{\text{N,rms}} = \beta V_{\text{N,elec}} \simeq (3.9 \times 10^{-4})(2.13 \times 10^{-2}) \simeq 8.3 \times 10^{-6} \text{ V/m}. \quad (2.117)$$

If the analyzer reading corresponds to a 1 Hz resolution bandwidth, then the electric-field noise PSD at the secular frequency is

$$S_E(\omega) \approx E_{\text{N,rms}}^2 \simeq (8.3 \times 10^{-6})^2 \simeq 6.9 \times 10^{-11} \frac{(\text{V/m})^2}{\text{Hz}}. \quad (2.118)$$

For a single motional mode, the heating rate is

$$\dot{\bar{n}} \simeq 1.2 \times 10^3 \text{ s}^{-1} \approx 1.2 \text{ quanta/ms}. \quad (2.119)$$

This is within an order of magnitude of the measured heating rate of the trap used in this experiment.

CHAPTER 3

Atomic Structure and Atom Laser Interactions

“The laser is a solution seeking a problem” – Theodore Maiman

The study of atomic spectra has played a central role in the development of modern quantum physics. The discovery of the spin of the electron was motivated by the need to explain fine structure in atomic spectra, including the doublet structure observed in sodium. (Uhlenbeck and Goudsmit, 1926, 1925). Rabi’s measurements of the hyperfine splitting in the ground state of hydrogen showed that the magnetic moment of the electron differed slightly from the value predicted by the simplest Dirac theory (Nafe et al., 1947). Schwinger explained this discrepancy by showing that the electron couples to the quantized electromagnetic field, which gives rise to a small radiative correction to its magnetic moment (Schwinger, 1948)

$$g_e = 2 \left(1 + \frac{\alpha}{2\pi} \right). \quad (3.1)$$

This result created the field of quantum electrodynamics.

The development of laser cooling transformed atomic physics from a field based primarily on spectroscopic observation into one based on precise quantum control. The idea of radiation pressure cooling was proposed independently for neutral atoms and trapped ions in 1975 (Hänsch and Schawlow, 1975; Wineland and Dehmelt, 1975). Later work by Wineland and others made it possible to cool and manipulate individual trapped ions with high precision (Wineland and Itano, 1979). These advances opened new directions in quantum computation, precision metrology, quantum simulation, and tests of fundamental physics.

Despite the development of laser cooling nearly five decades ago, atomic physicists have not yet conquered the entire periodic table. In practice, laser cooling has been most successful in atoms and ions with relatively simple electronic structure. These systems are often hydrogenic or alkali like, with a single valence electron outside a closed shell core. This structure produces a sparse spectrum with few states at low energy, allowing the atom to behave approximately as an effective two level or three level system.

This simplicity is essential for optical cycling. Laser cooling requires an atom to scatter many photons while remaining within a closed set of internal states. If the excited state can decay into many dark states, the atom quickly leaves the cooling cycle unless additional repump lasers are added. For atoms with simple level structures, optical cycling can often be achieved with only a small number of lasers, which is a crucial practical requirement for experiments.

Singly ionized barium belongs to this favorable class of systems. Ba^+ has one valence electron outside a closed shell core, giving it an electronic structure that is similar in many respects to that of an alkali atom. As a result, its relevant optical transitions form a manageable cooling and repumping scheme. This simple structure also makes Ba^+ well suited for reliable atomic structure calculations, since many of its low energy properties can be understood in terms of a single active electron interacting with a shielded core.

3.1 Barium Atomic Structure

Ba^+ can be approximated as a single valence electron outside a closed shell Ba^{2+} core. At large distances, the valence electron sees the Coulomb potential of a singly charged ion,

$$V(r) \rightarrow -\frac{1}{r}. \quad (3.2)$$

Thus, as mentioned above, this spectrum is hydrogenic when the electron is far from the core. However, at short distances the valence electron can penetrate the closed shell core and experience a potential that differs strongly from the pure Coulomb form. Quantum

defect theory provides a compact way to describe these shifts. The energy of a Rydberg series is written as

$$E_{n\ell j} = E_{\text{ion}} - \frac{R}{(n - \delta_{\ell j})^2}. \quad (3.3)$$

Here E_{ion} is the ionization limit, R is the reduced mass Rydberg constant, and $\delta_{\ell j}$ is the quantum defect. The quantum defect represents the phase shift acquired by the valence electron when it samples the short range core potential.

The Ba^+ valence electron wavefunction can be written as

$$\psi_{n\ell m}(\mathbf{r}) = \frac{u_{n\ell}(r)}{r} Y_{\ell m}(\theta, \phi). \quad (3.4)$$

The angular part is fixed by the spherical symmetry of the effective core potential, while the radial function contains the effect of core penetration. Quantum defect theory describes this effect by replacing the hydrogenic principal quantum number with an effective value,

$$\nu_{n\ell j} = n - \delta_{\ell j}. \quad (3.5)$$

The outer radial wavefunction is therefore Coulomb like, but with the energy and node structure set by $\nu_{n\ell j}$ rather than by the integer n . States with low angular momentum penetrate the core more strongly, so the s and p orbitals acquire large quantum defects and are pulled inward relative to hydrogenic orbitals with the same nominal n . This behavior is seen in Figure 3.1, where the Ba^+ $6s$, $6p$, and $5d$ radial probability densities are concentrated at smaller radii than the corresponding hydrogenic wavefunctions. The comparison illustrates why the Ba^+ spectrum is hydrogenic at long range, while its low lying optical transitions are set by short range interactions with the closed shell core.

An important artifact of this quantum defect theory is that the main doppler cooling transition in Ba^+ , the $S_{1/2} \rightarrow P_{1/2}$ transition, has a much longer wavelength than the same transition in hydrogen. In fact, of all of the laser-coolable ions, Ba^+ has the longest wavelength doppler cooling transition. If we examine the sequence Ca^+ , Sr^+ , and Ba^+ ,




Figure 3.1: Radial wavefunctions of Ba^+ and hydrogen for the first low lying excited states. The Ba^+ electron is more strongly localized near the nucleus than in hydrogen, reflecting the effect of the ionic core on the valence electron. The figure also shows that the radial separation between the s and p states is smaller in Ba^+ than in hydrogen.

the principal valence shell increases from $4s$ to $5s$ to $6s$. As the valence electron becomes more weakly bound and more spatially extended, the $S_{1/2} \rightarrow P_{1/2}$ transition shifts to longer wavelengths. This gives cooling transitions near 397 nm in Ca^+ , 422 nm in Sr^+ , and 493 nm in Ba^+ .

Interestingly, the heavier ion Ra^+ does not simply continue the same red shifting trend. Although it has the analogous $7S_{1/2} \rightarrow 7P_{1/2}$ transition, strong relativistic effects modify the $7s$, $7p$, and $6d$ level structure. In particular, relativistic stabilization of the $7S_{1/2}$ state increases the S to P energy gap, making the Ra^+ cooling transition bluer than that of Ba^+ . Thus, Ba^+ occupies a "Goldilocks" position in the alkaline earth ion sequence.

The metastable D manifold benefits from the same underlying physics. In Ba^+ , the $5D_{3/2}$ and $5D_{5/2}$ states lie lower in energy relative to the ground state than the corresponding D states in lighter alkaline earth ions. This reduced energy splitting extends the lifetime of the metastable manifold, since the electric quadrupole decay rate scales

strongly with transition frequency,

$$\Gamma_{E2} \propto \omega^5. \quad (3.6)$$

The long lifetime of these states can be used to improve electron shelving fidelity, since population transferred to the D manifold remains dark to the cooling transition for a long time. In addition, recent proposals have explored using states within the metastable D manifold as qubit states. Such qubits can be partially protected from laser fields that primarily address the $6S_{1/2}$ ground state, while remaining accessible through narrow optical transitions.

3.1.1 The Spin-Orbit Interaction

Thus far, the valence electron has been treated as moving in an effective electrostatic potential produced by the ionic core. This gives the gross structure of the Ba^+ spectrum. A more complete description must also include relativistic corrections. In the rest frame of the electron, its motion through the electric field of the ionic core produces an effective magnetic field. This magnetic field couples to the intrinsic magnetic moment associated with the electron spin. The resulting interaction shifts the energy levels according to the relative orientation of the spin angular momentum and the orbital angular momentum. This effect is known as spin-orbit coupling.

We can roughly derive this correction by noting, for an electron moving through the electric field of the ionic core, the Lorentz transformation of the field gives, to lowest order in v/c ,

$$\mathbf{B}' = -\frac{\mathbf{v} \times \mathbf{E}}{c^2}. \quad (3.7)$$

For a central potential $V(r)$, this field can be written in terms of the orbital angular momentum as

$$\mathbf{B}' = \frac{1}{em_e c^2 r} \frac{dV}{dr} \mathbf{L}. \quad (3.8)$$




Figure 3.2: Rough Ba⁺ level structure including the fine structure splitting.

The spin magnetic moment of the electron is

$$\boldsymbol{\mu}_s = -g_s \frac{e}{2m_e} \mathbf{S}. \quad (3.9)$$

The interaction energy $-\boldsymbol{\mu}_s \cdot \mathbf{B}'$ therefore gives a term proportional to $\mathbf{L} \cdot \mathbf{S}$. Since the electron rest frame is accelerating, the result must also include the Thomas precession correction, which reduces the coupling by a factor of two. The spin-orbit Hamiltonian is therefore

$$H_{\text{SO}} = \frac{1}{2m_e^2 c^2 r} \frac{dV}{dr} \mathbf{L} \cdot \mathbf{S}. \quad (3.10)$$

Using $\mathbf{J} = \mathbf{L} + \mathbf{S}$, we can write

$$\mathbf{J}^2 = \mathbf{L}^2 + \mathbf{S}^2 + 2\mathbf{L} \cdot \mathbf{S}. \quad (3.11)$$

Solving for the spin-orbit term gives

$$\mathbf{L} \cdot \mathbf{S} = \frac{1}{2} (\mathbf{J}^2 - \mathbf{L}^2 - \mathbf{S}^2). \quad (3.12)$$

For a state with definite n , ℓ , s , and j , the operator $\mathbf{L} \cdot \mathbf{S}$ has expectation value

$$\langle \mathbf{L} \cdot \mathbf{S} \rangle = \frac{\hbar^2}{2} [j(j+1) - \ell(\ell+1) - s(s+1)]. \quad (3.13)$$

Thus spin-orbit coupling splits each $\ell > 0$ orbital into fine structure levels with $j = \ell \pm 1/2$. In Ba^+ , this separates the $6P$ manifold into $6P_{1/2}$ and $6P_{3/2}$, and the $5D$ manifold into $5D_{3/2}$ and $5D_{5/2}$.

When the spin-orbit interaction is included in the Hamiltonian, the operators L_z and S_z no longer commute with the full Hamiltonian H . Consequently, m_l and m_s are no longer conserved quantities and cannot, in general, be used to label stationary eigenstates of H . In this sense, m_l and m_s are no longer *good* quantum numbers. Instead, the appropriate labels must be associated with operators that continue to commute with the Hamiltonian. Since the total angular momentum operator is defined by $\mathbf{J} = \mathbf{L} + \mathbf{S}$, its z -component is $J_z = L_z + S_z$. Although L_z and S_z do not individually commute with the spin-orbit term $\mathbf{L} \cdot \mathbf{S}$, their sum J_z does. Therefore, m_j , the eigenvalue associated with J_z , remains a good quantum number and can be used to label the stationary states of the spin-orbit coupled Hamiltonian.

3.1.2 The Hyperfine Interaction

The nucleus can also possess an intrinsic magnetic moment, which is determined by its nuclear spin. For the odd mass isotopes of Ba^+ , the nuclear spin is nonzero. In these

isotopes, the nuclear magnetic moment interacts with the magnetic field produced by the electrons. This interaction further splits each fine structure level according to the relative orientation of the nuclear spin $\mathbf{I}$ and the total electronic angular momentum $\mathbf{J}$. The leading contribution is the magnetic dipole hyperfine interaction,

$$H_{\text{M1}} = -\boldsymbol{\mu}_I \cdot \mathbf{B}_e, \quad \boldsymbol{\mu}_I = g_I \mu_N \frac{\mathbf{I}}{\hbar}, \quad (3.14)$$

where g_I is the nuclear g -factor and μ_N is the nuclear magneton. The electronic magnetic field $\mathbf{B}_e$, is the magnetic field due to the moving charge, L and the spin S ,

$$\mathbf{B}_e \rightarrow b_{n\ell j} \frac{\mathbf{J}}{\hbar}. \quad (3.15)$$

Substituting this projected field into the magnetic dipole interaction gives

$$H_{\text{M1}} = -g_I \mu_N b_{n\ell j} \frac{\mathbf{I} \cdot \mathbf{J}}{\hbar^2}. \quad (3.16)$$

All radial matrix elements, nuclear factors, and numerical constants are then absorbed into the hyperfine constant A . The interaction can therefore be written in the effective form

$$H_{\text{hfs}} = A \mathbf{I} \cdot \mathbf{J}. \quad (3.17)$$

The good angular momentum after including hyperfine coupling is the total atomic angular momentum,

$$\mathbf{F} = \mathbf{I} + \mathbf{J}. \quad (3.18)$$

Using

$$\mathbf{F}^2 = \mathbf{I}^2 + \mathbf{J}^2 + 2\mathbf{I} \cdot \mathbf{J}, \quad (3.19)$$

the magnetic dipole hyperfine energy shift is

$$\Delta E_F^{\text{M1}} = \frac{A}{2} [F(F+1) - I(I+1) - J(J+1)]. \quad (3.20)$$

It is conventional to define

$$K = F(F + 1) - I(I + 1) - J(J + 1), \quad (3.21)$$

For nuclei with $I \geq 1$, there can also be an electric quadrupole interaction between the nuclear quadrupole moment and the electric field gradient produced by the electrons. This gives an additional term for states with $J \geq 1$,

$$\Delta E_F^{\text{E2}} = B \frac{\frac{3}{4}K(K + 1) - I(I + 1)J(J + 1)}{2I(2I - 1)J(2J - 1)}. \quad (3.22)$$

The full hyperfine shift within a fine structure manifold is therefore

$$\Delta E_F = \frac{A}{2}K + B \frac{\frac{3}{4}K(K + 1) - I(I + 1)J(J + 1)}{2I(2I - 1)J(2J - 1)}. \quad (3.23)$$

3.2 The Zeeman Shift

The Zeeman interaction arises due to the coupling of an external magnetic field and the total magnetic moment of the atom

$$H_Z = -\boldsymbol{\mu} \cdot \mathbf{B}. \quad (3.24)$$

Thus the Zeeman Hamiltonian becomes

$$H_Z = g_J \mu_B B \frac{J_z}{\hbar} - g_I \mu_N B \frac{I_z}{\hbar}. \quad (3.25)$$

In the weak field limit, the hyperfine interaction is larger than the Zeeman interaction, so the states are well described by $|F, m_F\rangle$. For practical calculations at an arbitrary magnetic field, it is better to build and diagonalize the full Hamiltonian in the uncoupled




Figure 3.3: Frequency shifts due to the Zeeman interaction in the $D_{5/2}$ of $^{137}\text{Ba}^+$. Note that at low field the splitting the hyperfine interaction dominates, but at higher magnetic field, states from different F levels begin to mix.

basis $|m_J, m_I\rangle$. In this basis,


$$\mathbf{I} \cdot \mathbf{J} = I_z J_z + \frac{1}{2} (I_+ J_- + I_- J_+). \quad (3.26)$$

The diagonal matrix element is

$$\langle m_J, m_I | H | m_J, m_I \rangle = A m_I m_J + g_J \mu_B B m_J - g_I \mu_N B m_I. \quad (3.27)$$

The off diagonal hyperfine terms are

$$\begin{aligned} & \langle m_J - 1, m_I + 1 | H | m_J, m_I \rangle \\ &= \frac{A}{2} \sqrt{I(I+1) - m_I(m_I+1)} \sqrt{J(J+1) - m_J(m_J-1)} \end{aligned} \quad (3.28)$$



(a) Full sub structure of the $^{133}\text{Ba}^+$ isotope. (b) Full sub structure of the $^{137}\text{Ba}^+$ isotope.

Figure 3.4

and

$$\begin{aligned}
 & \langle m_J + 1, m_I - 1 | H | m_J, m_I \rangle \\
 &= \frac{A}{2} \sqrt{I(I+1) - m_I(m_I - 1)} \sqrt{J(J+1) - m_J(m_J + 1)}.
 \end{aligned} \tag{3.29}$$

Diagonalizing this matrix gives the Zeeman shifted energy levels. In the weak field limit, the result reduces to the linear shift $\Delta E_Z = g_F \mu_B B m_F$. In stronger fields, the same matrix captures the mixing of hyperfine states with the same value of $m_F = m_J + m_I$.

We can now examine the typical level structures of two promising barium isotopes, with all relevant terms in the Hamiltonian accounted for, as shown in Figure 3.4. One important distinction is that $^{133}\text{Ba}^+$ contains an isolated singlet $F = 0$ state, which is ideal for fast optical pumping and state preparation. In contrast, $^{137}\text{Ba}^+$ has a more complex level structure, which could be advantageous for encoding quantum information in more than two levels, as in qudit systems (Low et al., 2025). However, the larger nuclear spin of $^{137}\text{Ba}^+$ generally makes laser cooling more complicated and requires additional electro-optical-modulator (EOM) sidebands. The increased substructure may also make it more difficult to track leakage errors during state preparation or laser-driven Raman gates.




Figure 3.5: Clock states and the corresponding magnetic fields that they occur at in $^{133}\text{Ba}^+$ metastable $D_{5/2}$ state.

3.2.1 First Order Magnetic Field Insensitive Transitions

In a typical experimental apparatus, there is always some level of time-dependent magnetic field noise. Through the Zeeman shift, this magnetic field changes the frequency of the chosen qubit transition, producing a detuning error. For some trapped-ion qubit species, this can be the dominant source of decoherence (Ruster et al., 2016a). However, the magnetic-field sensitivity depends on both the choice of qubit states and the static magnetic field. It is therefore possible to choose particular transitions and magnetic fields for which the first-order magnetic-field shift is zero. These transitions are known as magnetic-field-insensitive, or “clock,” qubits. In general, all hyperfine isotopes of barium possess such clock transitions, although the required magnetic fields can be more or less experimentally accessible depending on the isotope and transition.

Magnetic field noise dephases a qubit encoded in two Zeeman states by converting fluctuations of the bias field into fluctuations of the qubit transition frequency,

$$\delta\omega_q(t) = \kappa \delta B(t), \quad \kappa \equiv \frac{\partial\omega_q}{\partial B}, \quad (3.30)$$

where κ is the differential Zeeman sensitivity of the two qubit states. In a Ramsey experiment, the accumulated phase is $\phi(t) = \kappa \int_0^t \delta B(t') dt'$, and the coherence is $C(t) = \langle e^{-i\phi(t)} \rangle$. For Gaussian quasi-static magnetic field noise with rms amplitude σ_B , this gives

$$C(t) = \exp \left[-\frac{1}{2} \kappa^2 \sigma_B^2 t^2 \right], \quad T_2^* = \frac{\sqrt{2}}{|\kappa| \sigma_B}. \quad (3.31)$$

Thus, for fixed magnetic noise, the Ramsey dephasing rate scales linearly with the magnitude of the differential Zeeman sensitivity. First order field-insensitive clock qubits suppress this leading term by operating near $\kappa \simeq 0$, which is reflected experimentally in coherence times ranging from tens of seconds to many minutes. For example, hyperfine clock qubits have demonstrated $T_2^* = 50$ s in $^{43}\text{Ca}^+$ (Harty et al., 2014), coherence exceeding 10 min in $^{171}\text{Yb}^+$ with sympathetic cooling and dynamical decoupling (Wang et al., 2017), and an estimated coherence time of approximately 5500 s in an optimized $^{171}\text{Yb}^+$ single ion memory (Wang et al., 2021). In contrast, Zeeman-sensitive electron spin qubits are typically limited to much shorter Ramsey times unless magnetic noise is strongly suppressed, with $T_2^* = 29.4 \pm 2.5$ ms measured for a ground-state $^{40}\text{Ca}^+$ Zeeman qubit (Haze et al., 2011), and $T_2^* = 300(50)$ ms, extendable to $T_2 = 2.1(1)$ s using spin echo, in a magnetically shielded $^{40}\text{Ca}^+$ Zeeman qubit (Ruster et al., 2016b).

Clock states can be identified by taking the derivative of the Zeeman-shifted transition frequency with respect to the applied magnetic field and locating points where this derivative vanishes, as shown in Figure 3.5. Although these transitions are first-order insensitive to magnetic-field noise, they still experience a second-order Zeeman shift, which can also contribute to decoherence. This sensitivity can be further reduced using techniques such as microwave dressing (Sárkány et al., 2014).

One practical limitation of very low-field clock states is that, below magnetic fields of order 1 G, Doppler cooling can become inefficient without polarization modulation. This occurs because coherent population trapping can form dark states in Zeeman-degenerate atomic systems. For this reason, a bias field of several gauss, often around 5 G, is typically

Table 3.1: First order field insensitive transitions considered in $^{133}\text{Ba}^+$. For $^{137}\text{Ba}^+$, only transitions with bias fields below 10 G are shown.

Isotope	State	Transition	B_0 (G)
$^{133}\text{Ba}^+$	$6S_{1/2}$	$ F = 0, m_F = 0\rangle \rightarrow F = 1, m_F = 0\rangle$	0.000
$^{133}\text{Ba}^+$	$5D_{5/2}$	$ F = 2, m_F = 0\rangle \rightarrow F = 3, m_F = 0\rangle$	0.000
$^{133}\text{Ba}^+$	$5D_{5/2}$	$ F = 2, m_F = 1\rangle \rightarrow F = 3, m_F = 1\rangle$	17.665
$^{133}\text{Ba}^+$	$5D_{5/2}$	$ F = 2, m_F = 2\rangle \rightarrow F = 3, m_F = 2\rangle$	35.330
$^{137}\text{Ba}^+$	$6S_{1/2}$	$ F = 1, m_F = -1\rangle \rightarrow F = 2, m_F = 1\rangle$	0.000
$^{137}\text{Ba}^+$	$6S_{1/2}$	$ F = 1, m_F = 0\rangle \rightarrow F = 2, m_F = 0\rangle$	0.000
$^{137}\text{Ba}^+$	$6S_{1/2}$	$ F = 1, m_F = 1\rangle \rightarrow F = 2, m_F = -1\rangle$	1.300
$^{137}\text{Ba}^+$	$5D_{5/2}$	$ F = 1, m_F = 0\rangle \rightarrow F = 2, m_F = 0\rangle$	0.000
$^{137}\text{Ba}^+$	$5D_{5/2}$	$ F = 2, m_F = 0\rangle \rightarrow F = 3, m_F = -1\rangle$	1.660
$^{137}\text{Ba}^+$	$5D_{5/2}$	$ F = 2, m_F = 2\rangle \rightarrow F = 3, m_F = 2\rangle$	3.340
$^{137}\text{Ba}^+$	$5D_{5/2}$	$ F = 3, m_F = -3\rangle \rightarrow F = 4, m_F = -1\rangle$	5.820
$^{137}\text{Ba}^+$	$5D_{5/2}$	$ F = 3, m_F = -1\rangle \rightarrow F = 4, m_F = 1\rangle$	8.620
$^{137}\text{Ba}^+$	$5D_{5/2}$	$ F = 3, m_F = 0\rangle \rightarrow F = 4, m_F = 2\rangle$	1.580

applied to destabilize dark states during cooling (Berkeland and Boshier, 2002). In principle, dressed microwave fields could also be used to shift the first-order field-insensitive point to a more experimentally convenient magnetic field.

3.3 The Light Shift

We have so far considered the effect of magnetic fields on the ion. A similar shift can also occur due to stray electric fields. In the static case, electric fields near the ion typically have only a small effect, since the relevant electric-dipole transitions are much farther detuned than the magnetic-dipole transitions considered above. However, when lasers address the ion, the AC Stark shift can produce appreciable shifts of the energy levels. These shifts can be especially deleterious during laser-driven gates, since amplitude noise on the laser is mapped onto the qubit as a detuning error through the light shift.

Often we will treat the AC light shift as a small perturbation, and work in the diagonalized hyperfine basis. We will see in later chapters that, for sufficiently large shifts, this approximation can break down. Nevertheless, the perturbative picture provides useful intuition for how laser fields shift the qubit frequency. Consider an atom interacting with

a classical electric field

$$\mathbf{E}(t) = \frac{1}{2}\mathbf{E}_0 e^{-i\omega_L t} + \frac{1}{2}\mathbf{E}_0^* e^{i\omega_L t}. \quad (3.32)$$

The interaction Hamiltonian is

$$H_{int} = -\mathbf{d} \cdot \mathbf{E}(t), \quad (3.33)$$

where $\mathbf{d}$ is the electric dipole operator. For a ground state $|g\rangle$ coupled off-resonantly to excited states $|e\rangle$, the second-order energy shift is

$$\Delta E_g = \sum_e \frac{|e - \mathbf{d} \cdot \mathbf{E}_0 / 2g|^2}{\hbar(\omega_g + \omega_L - \omega_e)}. \quad (3.34)$$

Defining the detuning

$$\Delta_{eg} = \omega_L - \omega_{eg}, \quad (3.35)$$

this becomes

$$\Delta E_g = \sum_e \frac{\hbar |\Omega_{eg}|^2}{4\Delta_{eg}}, \quad (3.36)$$

where

$$\Omega_{eg} = \frac{e\mathbf{d} \cdot \mathbf{E}_0 g}{\hbar} \quad (3.37)$$

is the resonant Rabi frequency for the transition $|g\rangle \leftrightarrow |e\rangle$. Thus the AC Stark shift of state $|g\rangle$ is

$$\delta_g = \frac{\Delta E_g}{\hbar} = \sum_e \frac{|\Omega_{eg}|^2}{4\Delta_{eg}}. \quad (3.38)$$

For a qubit encoded in two states $|0\rangle$ and $|1\rangle$, the laser shifts both states. The resulting shift of the qubit transition frequency is the differential AC Stark shift,

$$\delta_q = \delta_1 - \delta_0 = \sum_e \left(\frac{|\Omega_{e1}|^2}{4\Delta_{e1}} - \frac{|\Omega_{e0}|^2}{4\Delta_{e0}} \right). \quad (3.39)$$

This shift is proportional to the laser intensity through $|\Omega|^2 \propto |\mathbf{E}_0|^2$. Therefore, fluctuations

in laser power directly produce fluctuations in the qubit frequency,

$$\delta\omega_q \propto \delta I. \quad (3.40)$$

During a laser-driven gate, this appears as a time-dependent detuning error. If the intended qubit Hamiltonian is

$$H_q = \frac{\hbar\Omega_q}{2}\sigma_x, \quad (3.41)$$

then the differential light shift adds an unwanted term

$$H_{err} = \frac{\hbar\delta_q(t)}{2}\sigma_z. \quad (3.42)$$

Thus laser amplitude noise is converted into qubit phase noise through the AC Stark shift.

3.4 Doppler Cooling

One of the most critical tools in ion trapping is the ability to dissipate the kinetic energy of an ion by manipulating its internal degrees of freedom. This technique, known as Doppler cooling, is one of the crowning achievements of atomic physics. The basic principle is as follows. A laser is applied near resonance with an internal transition of an ion moving with velocity $\vec{v}$. In the ion's rest frame, the laser frequency is shifted by the Doppler effect, so the resonance condition depends on the ion's velocity. By tuning the laser slightly below resonance, ions moving toward the laser see the light shifted closer to resonance and are more likely to absorb photons.

Each absorption event transfers momentum from the laser field to the ion. If the excited state has a short lifetime, the ion then spontaneously emits a photon after a short time. The emitted photon is radiated in a nearly isotropic direction, so the recoil from spontaneous emission averages to zero over many scattering events. In contrast, the momentum transfer from absorption is directional, since photons are repeatedly absorbed from the cooling laser. As a result, the ion experiences a velocity-dependent damping force. This process removes

kinetic energy from the ion until the cooling from the velocity-dependent scattering force is balanced by heating from photon recoil. The resulting temperature is limited by the finite linewidth of the transition and is known as the Doppler cooling limit.

We can make the Doppler cooling picture quantitative by considering a two-level atom interacting with a laser of wavevector k and detuning

$$\Delta = \omega_L - \omega_0. \quad (3.43)$$

For an ion moving with velocity v along the laser direction, the laser detuning in the ion's rest frame is shifted by the Doppler effect,

$$\Delta_v = \Delta - kv. \quad (3.44)$$

The photon scattering rate is

$$\Gamma_{\text{sc}}(v) = \frac{\Gamma}{2} \frac{s}{1 + s + (2\Delta_v/\Gamma)^2}, \quad (3.45)$$

where Γ is the natural linewidth and $s = I/I_{\text{sat}}$ is the saturation parameter. Each absorbed photon transfers momentum $\hbar k$ to the ion. Therefore, the radiation-pressure force is

$$F(v) = \hbar k \Gamma_{\text{sc}}(v). \quad (3.46)$$

For small velocities, we expand the force about $v = 0$,

$$F(v) \simeq F(0) + \left. \frac{dF}{dv} \right|_{v=0} v. \quad (3.47)$$

For cooling, we want the velocity-dependent part of the force to oppose the velocity. We therefore define the damping coefficient α by

$$\alpha = - \left. \frac{dF}{dv} \right|_{v=0}. \quad (3.48)$$

Now we evaluate this derivative. The force depends on v only through the Doppler-shifted detuning Δ_v . Therefore, using the chain rule,

$$\frac{dF}{dv} = \hbar k \frac{d\Gamma_{\text{sc}}}{d\Delta_v} \frac{d\Delta_v}{dv} = -\hbar k^2 \frac{d\Gamma_{\text{sc}}}{d\Delta_v} = -\frac{4\hbar k^2 s \Delta / \Gamma}{[1 + s + (2\Delta / \Gamma)^2]^2}. \quad (3.49)$$

For red detuning, $\Delta < 0$, so $\alpha > 0$. The velocity-dependent force is then $-\alpha v$, which damps the ion motion. The cooling rate follows from the rate of change of kinetic energy. Since

$$E = \frac{1}{2}mv^2, \quad (3.50)$$

the cooling force gives

$$\left. \frac{dE}{dt} \right|_{\text{cool}} = F(v)v \simeq -\alpha v^2 = -\frac{2\alpha}{m}E. \quad (3.51)$$

Thus the Doppler cooling rate is

$$\gamma_{\text{cool}} = \frac{2\alpha}{m}. \quad (3.52)$$

The cooling force is ultimately balanced by recoil heating from photon scattering. Each scattering event consists of absorption from the cooling laser followed by spontaneous emission. The absorption process gives the ion a directed momentum kick of magnitude

$$\Delta p_{\text{abs}} = \hbar k. \quad (3.53)$$

Spontaneous emission also gives the ion a recoil kick, but the direction of this kick is random. Along the cooling axis, we denote this random kick by Δp_{em} . Since the emitted photon is radiated into random directions, the average kick along the cooling axis vanishes,

$$\langle \Delta p_{\text{em}} \rangle = 0. \quad (3.54)$$

Therefore, spontaneous emission does not produce a systematic force. However, the average

squared kick is nonzero,

$$\langle (\Delta p_{\text{em}})^2 \rangle \neq 0. \quad (3.55)$$

Thus, spontaneous emission causes diffusion in momentum space. This diffusion heats the ion because the kinetic energy depends on p^2 .

For the standard one-dimensional Doppler cooling treatment, the recoil heating rate is written as

$$\left. \frac{dE_z}{dt} \right|_{\text{heat}} = \frac{\hbar^2 k^2}{m} \Gamma_{\text{sc}}(0). \quad (3.56)$$

The total energy evolution is therefore

$$\frac{dE_z}{dt} = -\frac{2\alpha}{m} E_z + \frac{\hbar^2 k^2}{m} \Gamma_{\text{sc}}(0). \quad (3.57)$$

At steady state, the heating and cooling rates balance,

$$\frac{dE_z}{dt} = 0. \quad (3.58)$$

Therefore,

$$0 = -\frac{2\alpha}{m} E_z + \frac{\hbar^2 k^2}{m} \Gamma_{\text{sc}}(0). \quad (3.59)$$

Rearranging gives

$$\frac{2\alpha}{m} E_z = \frac{\hbar^2 k^2}{m} \Gamma_{\text{sc}}(0). \quad (3.60)$$

Solving for the steady-state energy gives

$$E_z = \frac{\hbar^2 k^2}{2\alpha} \Gamma_{\text{sc}}(0). \quad (3.61)$$

For one motional dimension,

$$E_z = \frac{1}{2} k_B T. \quad (3.62)$$

Thus,

$$k_B T = \frac{\hbar^2 k^2}{\alpha} \Gamma_{\text{sc}}(0). \quad (3.63)$$

We now substitute the scattering rate

$$\Gamma_{\text{sc}}(0) = \frac{\Gamma}{2} \frac{s}{1 + s + (2\Delta/\Gamma)^2} \quad (3.64)$$

and the damping coefficient

$$\alpha = -\frac{4\hbar k^2 s \Delta / \Gamma}{[1 + s + (2\Delta/\Gamma)^2]^2}. \quad (3.65)$$

This gives

$$k_B T = \frac{\hbar \Gamma}{4} \frac{1 + s + (2\Delta/\Gamma)^2}{-2\Delta/\Gamma}. \quad (3.66)$$

To find the minimum temperature, define

$$x = -\frac{2\Delta}{\Gamma}. \quad (3.67)$$

For red detuning, $x > 0$, and the temperature becomes

$$k_B T = \frac{\hbar \Gamma}{4} \frac{1 + s + x^2}{x}. \quad (3.68)$$

Thus we minimize

$$f(x) = \frac{1 + s + x^2}{x} = \frac{1 + s}{x} + x. \quad (3.69)$$

Taking the derivative,

$$\frac{df}{dx} = -\frac{1 + s}{x^2} + 1, \quad (3.70)$$

and setting this equal to zero gives

$$x^2 = 1 + s. \quad (3.71)$$

Therefore,

$$x = \sqrt{1 + s}, \quad (3.72)$$

and the optimal detuning is

$$\Delta_{\text{opt}} = -\frac{\Gamma}{2}\sqrt{1+s}. \quad (3.73)$$

In the low-saturation limit, $s \ll 1$, this becomes

$$\Delta_{\text{opt}} = -\frac{\Gamma}{2}. \quad (3.74)$$

Evaluating the temperature at $s \ll 1$ and $x = 1$ gives

$$k_B T_D = \frac{\hbar\Gamma}{4} \frac{1+1}{1} = \frac{\hbar\Gamma}{2}. \quad (3.75)$$

Thus the Doppler cooling limit is

$$T_D = \frac{\hbar\Gamma}{2k_B}. \quad (3.76)$$

3.5 Dipole Matrix Element Calculations

The most accurate experimental method to determine the dipole matrix element of an atomic transition is to extract it from the radiative lifetime. Therefore, we must derive the equation for the dipole matrix element in terms of the lifetime of the state. This is equivalent to the derivation of the rate of spontaneous emission. Qualitatively, we are asking: what is the coupling strength between the excited state of our atomic transition with one photon and the ground state with one fewer photon?

We consider an atomic transition $|e\rangle \rightarrow |g\rangle$ with energy splitting $\hbar\omega_0$. The electromagnetic field is quantized in a box of volume V . Initial and final states include both atom and field:

$$|i\rangle = |e\rangle |n_{\mathbf{k}\epsilon}\rangle, \quad |f\rangle = |g\rangle |n_{\mathbf{k}\epsilon} + 1\rangle, \quad (3.77)$$

corresponding to emission of one photon into mode $(\mathbf{k}, \epsilon)$. The differential transition rate into final photon states with energy in $[E, E + dE]$ and wavevectors within a differential

solid angle $d\Omega$ is

$$dW_{fi} = \frac{2\pi}{\hbar} |\langle f|H'|i\rangle|^2 \rho_f(E) P(E) dE, \quad (3.78)$$

where $\rho_f(E)$ is the density of final photon states per unit energy and $P(E)$ is the probability of absorption(emission) of a photon of a given energy (Budker et al., 2008). This function can be viewed as a restatement of the Heisenberg time-energy uncertainty principle, due to the finite lifetime of the excited state.

In order to explicitly write the density of emitted photon states, we begin by counting the possible modes in phase space. Suppose we have a finite volume $V = L^3$. Then the allowed wavevectors are discrete,

$$\mathbf{k} = \frac{2\pi}{L}(n_x, n_y, n_z), \quad n_i \in \mathbb{Z}, \quad (3.79)$$

so each mode occupies a k -space cell of volume

$$(\Delta k_x)(\Delta k_y)(\Delta k_z) = \left(\frac{2\pi}{L}\right)^3 = \frac{(2\pi)^3}{V}. \quad (3.80)$$

Hence the number of modes (for a fixed polarization) in a k -space volume element d^3k is

$$dN = \frac{d^3k}{(2\pi)^3/V} = \frac{V}{(2\pi)^3} d^3k. \quad (3.81)$$

In spherical coordinates $d^3k = k^2 dk d\Omega$, giving

$$dN = \frac{V}{(2\pi)^3} k^2 dk d\Omega, \quad (3.82)$$

and including both transverse photon polarizations multiplies this result by a factor of 2.

Taking

$$dN = \frac{V}{(2\pi)^3} k^2 dk d\Omega, \quad (3.83)$$

and using $E = \hbar\omega = \hbar ck$ yields

$$\rho_f(E) \equiv \frac{dN}{dE} = \frac{V}{(2\pi)^3} \frac{E^2}{\hbar^3 c^3} d\Omega. \quad (3.84)$$

If the excited state has decay rate γ , the allowed photon frequencies are distributed as a Lorentzian,

$$P(\omega) = \frac{\gamma/(2\pi)}{(\omega - \omega_0)^2 + (\gamma/2)^2}, \quad \int_{-\infty}^{\infty} P(\omega) d\omega = 1, \quad (3.85)$$

and in the limit $\gamma \rightarrow 0$ one recovers $P(\omega) \rightarrow \delta(\omega - \omega_0)$. This factor enforces the restrictions of energy and momentum conservation gives the accessible final states.

Starting from minimal coupling and keeping only terms linear in the field,

$$H' = -\frac{e}{m} \mathbf{p} \cdot \mathbf{A}(\mathbf{r}), \quad (3.86)$$

and expanding the quantized vector potential in modes (SI normalization),

$$\mathbf{A}(\mathbf{r}) = \sum_{\mathbf{k}, \epsilon} \sqrt{\frac{\hbar}{2\varepsilon_0 V \omega}} \left(a_{\mathbf{k}\epsilon} \boldsymbol{\epsilon} e^{i\mathbf{k} \cdot \mathbf{r}} + a_{\mathbf{k}\epsilon}^\dagger \boldsymbol{\epsilon}^* e^{-i\mathbf{k} \cdot \mathbf{r}} \right). \quad (3.87)$$

For emission $|n\rangle \rightarrow |n+1\rangle$ only the $a^\dagger$ term contributes, giving

$$\langle f | H' | i \rangle = -\frac{e}{m} \sqrt{\frac{\hbar}{2\varepsilon_0 V \omega}} \sqrt{n+1} \langle g | \mathbf{p} \cdot \boldsymbol{\epsilon}^* e^{-i\mathbf{k} \cdot \mathbf{r}} | e \rangle. \quad (3.88)$$

For optical transitions the electron cloud is small compared to the wavelength, thus we can make use of the dipole approximation,

$$\mathbf{k} \cdot \mathbf{r} \ll 1 \quad \Rightarrow \quad e^{-i\mathbf{k} \cdot \mathbf{r}} \simeq 1. \quad (3.89)$$

To convert $\mathbf{p}$ to $\mathbf{r}$, use the Heisenberg equation $\dot{\mathbf{r}} = (i/\hbar)[H_0, \mathbf{r}] = \mathbf{p}/m$ so that

$$\langle g | \mathbf{p} | e \rangle = \frac{im}{\hbar} \langle g | [H_0, \mathbf{r}] | e \rangle = \frac{im}{\hbar} \left(\langle g | H_0 \mathbf{r} | e \rangle - \langle g | \mathbf{r} H_0 | e \rangle \right) = -im\omega_0 \langle g | \mathbf{r} | e \rangle. \quad (3.90)$$

Define the electric dipole operator $\mathbf{d} \equiv -e\mathbf{r}$ to obtain

$$\langle f|H'|i\rangle = -i\sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}}\sqrt{n+1} \langle g|\mathbf{d} \cdot \boldsymbol{\epsilon}^*|e\rangle \quad (\text{set } \omega \simeq \omega_0). \quad (3.91)$$

Thus,

$$|\langle f|H'|i\rangle|^2 = \frac{\hbar\omega_0}{2\varepsilon_0 V}(n+1) |\langle g|\mathbf{d} \cdot \boldsymbol{\epsilon}^*|e\rangle|^2. \quad (3.92)$$

For spontaneous emission we set $n = 0$ ($n \neq 0$ terms are stimulated emission). Now we can insert the matrix element and $\rho_f(E)$ into Fermi's Golden Rule, and take $E = \hbar\omega_0$. Next we will integrate over $d\Omega$ and sum over the two polarizations.

$$\gamma = \frac{\omega_0^3}{8\pi^2\varepsilon_0\hbar c^3} \int d\Omega \sum_{\lambda=1}^2 \frac{1}{2J'+1} \sum_{M',M} |\boldsymbol{\epsilon}_{\mathbf{k}\lambda} \cdot \mathbf{d}_{gJM, eJ'M'}|^2, \quad (3.93)$$

where $\mathbf{d}_{gJM, eJ'M'} \equiv \langle g, J, M|\mathbf{d}|e, J', M'\rangle$. At this point we have made some practical considerations. First we have assumed that m_J, J are good quantum numbers. We have also assumed that we will only be measuring the total scattering rate from a particular J level to another J level, thus we will want to sum over the possible final states and account for an equal distribution of population in the initial m_J states. Hence, the factor $1/(2J'+1)$ averages over an unpolarized excited state.

Using the completeness relation for transverse polarization vectors,

$$\sum_{\lambda=1}^2 \epsilon_{\mathbf{k}\lambda, i} \epsilon_{\mathbf{k}\lambda, j}^* = \delta_{ij} - \hat{k}_i \hat{k}_j, \quad (3.94)$$

we obtain

$$\sum_{\lambda=1}^2 |\boldsymbol{\epsilon}_{\mathbf{k}\lambda} \cdot \mathbf{d}|^2 = d_i d_j^* (\delta_{ij} - \hat{k}_i \hat{k}_j) = |\mathbf{d}|^2 - |\hat{\mathbf{k}} \cdot \mathbf{d}|^2. \quad (3.95)$$

Choose the z axis along $\mathbf{d}$ so that $\hat{\mathbf{k}} \cdot \mathbf{d} = |\mathbf{d}| \cos \theta$. Then Eq. (3.95) becomes $|\mathbf{d}|^2 \sin^2 \theta$,

and

$$\int d\Omega \left(|\mathbf{d}|^2 - |\hat{\mathbf{k}} \cdot \mathbf{d}|^2 \right) = |\mathbf{d}|^2 \int d\Omega \sin^2 \theta = |\mathbf{d}|^2 \frac{8\pi}{3}. \quad (3.96)$$

Substituting Eq. (3.96) into Eq. (3.93) gives

$$\gamma = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} \frac{1}{2J'+1} \sum_{M',M} |\mathbf{d}_{gJM,eJ'M'}|^2. \quad (3.97)$$

Write the squared vector matrix element as a sum over spherical components, $|\mathbf{d}|^2 = \sum_{q=-1}^{+1} |d_q|^2$, so that

$$\sum_{M',M} |\mathbf{d}_{gJM,eJ'M'}|^2 = \sum_{M',M} \sum_{q=-1}^{+1} |\langle g, J, M | d_q | e, J', M' \rangle|^2. \quad (3.98)$$

Using the Wigner–Eckart theorem,

$$\langle g, J, M | d_q | e, J', M' \rangle = \frac{(g, J || \mathbf{d} || e, J')}{\sqrt{2J'+1}} C_{JM,1q}^{J'M'}, \quad (3.99)$$

we find

$$\sum_{M',M,q} |\langle g, J, M | d_q | e, J', M' \rangle|^2 = \frac{|(g, J || \mathbf{d} || e, J')|^2}{2J'+1} \sum_{M',M,q} \left| C_{JM,1q}^{J'M'} \right|^2. \quad (3.100)$$

By Clebsch–Gordan orthogonality,

$$\sum_{M,q} \left| C_{JM,1q}^{J'M'} \right|^2 = 1 \quad \text{for each fixed } M', \quad (3.101)$$

so summing over M' yields

$$\sum_{M',M,q} \left| C_{JM,1q}^{J'M'} \right|^2 = 2J'+1. \quad (3.102)$$

Therefore,

$$\sum_{M', M} |\mathbf{d}_{gJM, eJ'M'}|^2 = |(g, J \| \mathbf{d} \| e, J')|^2. \quad (3.103)$$

Substituting into Eq. (3.97) gives

$$\gamma = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} \frac{|(g, J \| \mathbf{d} \| e, J')|^2}{2J' + 1}. \quad (3.104)$$

Inverting for the reduced dipole matrix element,

$$|(g, J \| \mathbf{d} \| e, J')| = \sqrt{\frac{3\pi\epsilon_0\hbar c^3(2J' + 1)}{\omega_0^3}} \gamma. \quad (3.105)$$

Often, we will want to calculate the exact dipole matrix element between a specific, m_J or m_F sub-level. Accordingly, let's suppose we first have a specific $J_1, m_{J_1}, J_2, m_{J_2}$,

$$\begin{aligned} \langle Jm | d_q^{(1)} | J'm' \rangle &= (-1)^{J-m} \begin{pmatrix} J & 1 & J' \\ -m & q & m' \end{pmatrix} \langle J \| \mathbf{d} \| J' \rangle \\ &= \frac{1}{\sqrt{2J+1}} \langle J'm'; 1q | Jm \rangle \langle J \| \mathbf{d} \| J' \rangle, \quad q \in \{-1, 0, +1\}. \end{aligned} \quad (3.106)$$

Now suppose we are working in the F, m_F basis, we can expand our matrix elements in terms of the I, J basis,

$$\begin{aligned} \langle Fm_F | d_q^{(1)} | F'm'_F \rangle &= \sum_{m_J, m_I} \sum_{m'_J, m'_I} \langle Jm_J, Im_I | Fm_F \rangle \langle J'm'_J, Im'_I | F'm'_F \rangle \langle Jm_J; Im_I | d_q^{(1)} | J'm'_J; Im'_I \rangle \\ &= \sum_{m_J, m_I} \sum_{m'_J} \langle Jm_J, Im_I | Fm_F \rangle \langle J'm'_J, Im_I | F'm'_F \rangle \langle Jm_J | d_q^{(1)} | J'm'_J \rangle, \end{aligned} \quad (3.107)$$

where we used the fact that $d_q^{(1)}$ acts only on the electronic part, hence $\langle Im_I | Im'_I \rangle = \delta_{m_I, m'_I}$.

$$\langle Jm_J, Im_I | Fm_F \rangle = (-1)^{J-I+m_F} \sqrt{2F+1} \begin{pmatrix} J & I & F \\ m_J & m_I & -m_F \end{pmatrix}, \quad (3.108)$$

and similarly for $\langle J'm'_J, Im_I | F'm'_F \rangle$. Now we have,

$$\begin{aligned} & \langle Fm_F | d_q^{(1)} | F'm'_F \rangle \\ &= \langle J \| \mathbf{d} \| J' \rangle \sqrt{(2F+1)(2F'+1)} (-1)^{J-I+m_F} (-1)^{J'-I+m'_F} \\ & \times \sum_{m_J, m_I, m'_J} (-1)^{J-m_J} \begin{pmatrix} J & I & F \\ m_J & m_I & -m_F \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -m_J & q & m'_J \end{pmatrix} \begin{pmatrix} J' & I & F' \\ m'_J & m_I & -m'_F \end{pmatrix}. \end{aligned} \quad (3.109)$$

Now apply the Racah 3j-sum identity, which converts a sum of three 3j symbols into one 3j symbol times a 6j symbol,

$$\begin{aligned} & \sum_{m_J, m_I, m'_J} (-1)^{J-m_J} \begin{pmatrix} J & I & F \\ m_J & m_I & -m_F \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -m_J & q & m'_J \end{pmatrix} \begin{pmatrix} J' & I & F' \\ m'_J & m_I & -m'_F \end{pmatrix} \\ &= (-1)^{J'+I+F+1} \begin{Bmatrix} J & F & I \\ F' & J' & 1 \end{Bmatrix} \begin{pmatrix} F & 1 & F' \\ -m_F & q & m'_F \end{pmatrix}. \end{aligned} \quad (3.110)$$

Combining Eqs. (3.106) and (3.110) yields the fully explicit result:

$$\begin{aligned} \langle Fm_F | d_q^{(1)} | F'm'_F \rangle &= (-1)^{F-m_F} \begin{pmatrix} F & 1 & F' \\ -m_F & q & m'_F \end{pmatrix} \\ & \times (-1)^{J'+I+F+1} \sqrt{(2F+1)(2F'+1)} \begin{Bmatrix} J & F & I \\ F' & J' & 1 \end{Bmatrix} \langle J \| \mathbf{d} \| J' \rangle. \end{aligned} \quad (3.111)$$

It should be noted that our original definition of the reduced dipole matrix element, relies on the assumption that the initial state is the excited state. In order to take the

conjugate element we must be careful to note,

$$\left(d_q^{(1)}\right)^\dagger = (-1)^q d_{-q}^{(1)} \implies \langle a|d_q^{(1)}|b\rangle^* = (-1)^q \langle b|d_{-q}^{(1)}|a\rangle, \quad q = 0, \pm 1. \quad (3.112)$$

CHAPTER 4

The Experiment

“Einstein’s definition of insanity was doing the same thing over and over again and expecting different results. Clearly, he was not an experimentalist, because I do the same thing every day and still hope for better data.” – Me

4.1 First Generation $^{133}\text{Ba}^+$ Trap

Prior to the work of this thesis, Justin Christensen and Dave Hucul constructed the first of its kind $^{133}\text{Ba}^+$ trap (Christensen, 2020). The goal of this apparatus was to load and isolate single $^{133}\text{Ba}^+$ ions as a proof of principle. Therefore, design choices were made to optimize loading efficiency. This entailed creating a “floppy” trap with a large ion to electrode distance or r_0 of 3 mm, resulting in low secular frequencies and a large trapping volume. Initially, they set out to load ions via thermionic emission from an oven, but this method did not produce observable results. Following these initial efforts, they turned to laser ablation as an alternative method for ion creation. This was accompanied by the use of “trap-door” loading, in which the trap RF drive signal is gated with a transistor-transistor logic (TTL) and turned off in sync with the ablation laser pulse (Hendricks et al., 2007). The trap is then turned back on after a set delay that roughly corresponds to the transit time of the ions from the target.

Each ablation and loading attempt would result in hundreds of ions being loaded. The most abundant species, making up $\approx 80\%$ of the trap contents, was $^{132}\text{Ba}^+$. The rest of




Figure 4.1: Typical ion clouds in the first generation trap as distillation to single chains of $^{133}\text{Ba}^+$ occurs. The first picture is after ablation loading, the second is after laser heating of all isotopes except $^{132}\text{Ba}^+$ and $^{133}\text{Ba}^+$, the final picture shows the final step of distillation after tickling.

the loaded ions consisted of other isotopes of Ba^+ , as well as molecular ions such as BaCl^+ , BaO^+ , and BaH^+ . After initial trapping, the ion cloud was laser cooled and all species except for $^{133}\text{Ba}^+$ were ejected from the trap. This procedure for isolating a particular isotope was initially very stochastic and chaotic, and often resulted in the ejection of all ions. For this reason, loading $^{133}\text{Ba}^+$ could take hours to days and require hundreds of ablation attempts. The ejection of other Ba^+ isotopes was achieved by laser heating using a separate electro-optic modulator (EOM) tone on the 493 nm laser that was swept across the resonance of the Doppler cooling transition. After laser heating, the remaining contents of the trap were $^{133}\text{Ba}^+$ and dark molecular ions.

The initial procedure for expelling these molecular ions was to apply a large Mathieu a parameter for a brief period of time by ramping two orthogonal DC electrodes. The rationale for this method was that, by pushing the ion cloud to the edge of the stability region, any ion whose charge-to-mass ratio differed from that of $^{133}\text{Ba}^+$ would be expelled from the trap. While this method was sufficient for the first demonstrations of trapping, it had two significant limitations: first, it was difficult to determine precisely where the stability boundary lay for $^{133}\text{Ba}^+$; and second, it could not remove molecular ions with nearly the same mass as $^{133}\text{Ba}^+$, such as $^{132}\text{BaH}^+$. Due to these limitations, it was common to isolate a few $^{133}\text{Ba}^+$ ions together with roughly twice as many dark ions that were extremely difficult to “shake” out of the trap.

To accelerate the isotope purification process, we implemented an additional radio-frequency (RF) drive, commonly referred to as a “tickle,” which was applied to the trap electrodes and swept over frequency ranges corresponding to the secular motion of unwanted ions. The underlying principle of this technique is that ions in a Paul trap exhibit mass dependent secular frequencies; by driving the trap at or near one of these frequencies, resonant energy is transferred to ions with the corresponding charge-to-mass ratio, leading to increased motional amplitude and eventual loss from the trap.

The tickle drive was generated using an Agilent signal generator and attenuated through an adjustable filter chain to ensure precise control over the applied voltage. Careful reduction of the drive amplitude was essential to avoid excessive heating or unintentional loss of the desired $^{133}\text{Ba}^+$ ions. The drive frequency was initially swept over ranges corresponding to molecular ions and isotopes with masses significantly different from 133 amu, efficiently removing these species from the trap.

After this coarse purification step, the tickle amplitude was further reduced and the frequency sweep was refined to target ions with masses very close to $^{133}\text{Ba}^+$. In particular, molecular ions such as $^{132}\text{BaH}^+$ have nearly identical mass and therefore nearly identical secular frequencies, making them difficult to distinguish using static stability based methods alone. In this regime, the tickle drive was carefully calibrated by monitoring the displacement and heating of the $^{133}\text{Ba}^+$ ions, as observed through changes in fluorescence and spatial extent of the ion crystal.

The effectiveness of this method relies on the differential response of co-trapped species to the applied drive. While the tickle induces resonant heating of unwanted ions, the $^{133}\text{Ba}^+$ ions are simultaneously subjected to laser cooling. By detuning the cooling laser far to the red of the Doppler transition, we increase the damping rate for the desired ions, allowing them to remain trapped despite the applied drive. In contrast, molecular ions, which do not interact with the cooling light, experience uncompensated heating and are preferentially driven out of the trap. In this way, the technique effectively decouples the dynamics of $^{133}\text{Ba}^+$ and near-degenerate molecular ions, enabling selective removal of the

latter.

This resonant excitation approach significantly improved the efficiency and reproducibility of the loading process, reducing the time required to isolate pure samples of $^{133}\text{Ba}^+$ compared to methods based solely on static stability manipulation.

4.2 Second Generation $^{133}\text{Ba}^+$ Trap

The construction and design of the second $^{133}\text{Ba}^+$ trap began in January 2020. There were three guiding objectives in designing this trap: to improve SPAM measurements, to establish motional control of the ions, and to perform multi-ion entangling operations. To improve the state detection error, a blade trap geometry was chosen to maximize the solid angle for light collection. The trap is also significantly tighter than the first iteration, with $r_0 = 200\text{ }\mu\text{m}$. This allows for much stronger confinement and therefore higher secular frequencies at modest voltages. Higher secular frequencies enable faster sideband-cooling rates as well as faster motional entangling gates. Finally, the new trap is loaded via photoionization rather than “trap-door” loading. This allows ions to be loaded continuously without having to drop the trap and expel its contents.

The blade-trap assembly was sourced entirely from the Monroe group at the University of Maryland (now Duke), and we are grateful to them for providing it. The blade electrodes (See Figure 4.2) were fabricated using precision machining techniques, followed by surface polishing and gold coating to reduce patch potentials and ensure compatibility with ultra-high-vacuum operation (Kim et al., 2022). The electrodes were mounted on insulating spacers and aligned to achieve the desired trap geometry and ion–electrode distance. Additional details can be found in Refs. (Hucul, 2015; Johnson, 2020).

Manual assembly of the blades is a difficult process and likely introduces small imperfections in the trap geometry. For this reason, multiple groups have pursued alternative fabrication and assembly strategies, including the use of sapphire spacers to position the blades and more recently, fully monolithic trap fabrication (Menon et al., 2026).




Figure 4.2: Diagram of RF blades

After assembly, electrical connections to the blade electrodes were established using wire bonding, followed by attachment of external leads via spot welded constantan foil. Fine bond ribbons (99.99 Au size: .01 x .0005 Machine No:34W) were first used to connect each electrode to nearby contact pads, providing low-inductance and mechanically stable electrical connections at the trap structure. To interface these connections with the macroscopic wiring, strips of constantan foil were then spot welded to the bond pads. Constantan was chosen for its favorable mechanical properties and relatively low thermal conductivity, which helps minimize heat load while maintaining reliable electrical contact in ultra-high-vacuum conditions. These foil leads were subsequently attached to Kapton insulated wires, allowing connection to the external control electronics. This two stage approach provides both the precision of wire bonding at the electrode level and the robustness required for integration with the larger electrical system.

A possible improvement to the trap wiring would be to reduce the number of intermediate electrical junctions by implementing a more direct connection scheme, such as




Figure 4.3: Image of blades as the wiring bonding was done. Blades had to be assembled and aligned before the wire bonding due to the fact that the ribbons would easily detach during assembly.

a screw-terminal style interface between the trap leads and the external wiring. While the wire bonds and gold bond ribbons themselves are expected to have low resistance, the complete electrical path includes multiple material interfaces and mechanical junctions. Any poorly conducting connection in this chain can contribute additional series resistance, which in turn can degrade the resonator quality factor Q . A reduction in Q lowers the RF voltage obtainable for a given input power, thereby requiring higher power to achieve the same trapping potential. Minimizing the number of such junctions could therefore reduce resistive losses and improve the efficiency of RF delivery to the trap.

After the initial pump-down, a catastrophic electrical failure was discovered. The wire connecting the feedthrough to the constantan foil melted and caused a short once the RF power was increased above 0.1 W. The primary cause was an insufficient wire gauge,

which led to excessive resistive heating. After identifying this issue, the trap was reopened and the RF wiring was replaced with 18 AWG wire. To minimize the invasiveness of the repair, alligator clips were attached to the blade screw terminals rather than redoing the wire bonding.

A major motivation for this approach was to limit exposure to radioactive material by avoiding full removal of the trap from the vacuum chamber. An additional goal was to reduce water readsorption onto the stainless steel surfaces and thereby avoid the need for a second bakeout. To this end, nitrogen was backfilled through a cold trap consisting of a copper coil, and the chamber was brought to approximately 0.1 atm before being opened. Unfortunately, this procedure was not sufficient, and a second bakeout was required to recover ultra-high-vacuum conditions.

4.3 Trap Electronics

The trap drive is generated by a low-noise Rohde & Schwarz SML02 signal generator operated at 31.4 MHz with an output power of 1 dBm. This signal is then split and one branch is sent through a mixer, where it is combined with a DDS channel to enable amplitude modulation for micromotion compensation and secular-frequency diagnostics. The resulting signal is subsequently fed into a low-noise amplifier (Mini-Circuits ZX60-100VH+), which provides approximately 0.5 W of RF power. Several other trap amplifiers were also tested, including the ZHL-5W-2G-S+ and the LZY-22+, but both were found to produce ion heating rates approximately a factor of two larger. The amplified RF signal is then applied to the antenna coil of a bifilar helical resonator.

The trap cannot be driven directly from an RF amplifier, since a direct connection typically leads to impedance mismatch, reflected power, and increased injection of amplifier noise into the trap. For this reason, a helical resonator is placed between the amplifier and the ion trap. The resonator serves both as an impedance matching element and as a narrow band filter, thereby reducing off resonant noise while simultaneously enabling a




Figure 4.4: RF drive chain to generate the trap signal.

larger RF voltage for a given applied power.

The helical resonator may be understood as consisting of two coils (3 in the case of a bifilar can) with distinct roles. The small antenna coil is connected to the RF source and serves primarily as the coupling element that transfers power into the resonator through mutual inductance. The main helical coil, by contrast, is the element that stores magnetic energy and, together with the total capacitance of the trap, wiring, shield, and coil self-capacitance, forms the high- Q resonant circuit. In this sense, the antenna coil determines how efficiently RF power is delivered into the resonator, whereas the main coil is responsible for generating the large RF voltage applied to the trap electrodes on

resonance.

The outer shield, or “can,” is not merely mechanical housing but an active part of the resonant structure. In particular, it contributes a coil-to-shield capacitance and therefore affects both the resonant frequency and the quality factor of the resonator. For this reason, the main coil must be mounted centrally within the shield. If the coil is displaced from the center, the spacing between the coil and shield becomes asymmetric, increasing the effective shield capacitance on one side and introducing unwanted fringe-field effects. This can shift the resonant frequency, degrade the quality factor and reduce the reproducibility of the resonator response. Proper centering therefore helps minimize parasitic capacitance and ensures that the resonator behaves as intended. In order to properly align the main coil, 3D printed spacers were placed on four sides to center the coil.

Impedance matching is controlled primarily by the antenna coil. Since the antenna coil and the main coil are inductively coupled, the input impedance seen by the RF source depends on the strength of that coupling. Adjusting the antenna coil position changes the mutual inductance between the two coils and therefore changes the effective impedance presented to the source. This is why the antenna coil is placed on a precision micrometer and is adjusted experimentally to minimize reflected power and obtain a good match to the $50\ \Omega$ source.

We can model the trap system as an effective series RLC circuit operated on resonance. In that case, the resonant angular frequency is

$$\omega_0 = \frac{1}{\sqrt{LC}}, \quad (4.1)$$

and the quality factor is

$$Q = \frac{1}{R} \sqrt{\frac{L}{C}}. \quad (4.2)$$

The voltage across the trap is approximately the voltage across the effective capacitance

Table 4.1: Resonator parameters and estimated blade voltage.

Quantity	Symbol	Value
Resonance frequency	f	31.4 MHz
Effective capacitance	C	12 pF
Loaded quality factor	Q	211
Input power	P	0.5 W
Voltage (amp)	$V_{\text{blade,pk}}$	299 V

of the resonant circuit. Following the derivation from (Siverns et al., 2012),

$$V_{\text{pk}} \approx \kappa \sqrt{2PQ}, \quad (4.3)$$

where

$$\kappa = \left(\frac{L}{C} \right)^{1/4}, \quad (4.4)$$

so that the corresponding RMS voltage is

$$V_{\text{rms}} = \frac{V_{\text{pk}}}{\sqrt{2}} \approx \kappa \sqrt{PQ}. \quad (4.5)$$

Equivalently, using $\omega_0 = 1/\sqrt{LC}$, this may be written as

$$V_{\text{rms}} \approx \sqrt{\frac{PQ}{\omega_0 C}}, \quad V_{\text{pk}} \approx \sqrt{\frac{2PQ}{\omega_0 C}}. \quad (4.6)$$

The typical parameters used in this trap are found in Table 4.1. The calculated voltage is in good agreement with the measured voltage. Note that the method to measure the trap voltage is to reduce the power significantly, and then place a low capacitance probe on the RF feed through. Then after readjusting the antenna coil to maximize the Q, the voltage can be measured. Then a pick-off electrode, in this case a small piece of metal used as a filter for the ablation target, is also measured. The ratio of these two voltages is then used to calculate the trap voltage when the probe is removed. The pick off voltage at 1 dbm set point power is $3.8 V_{\text{amp}}$, and the measured ratio was 72, corresponding to a trap voltage of 280 V.




Figure 4.5: Inner view of the resonator can coil and black 3D printed spacers. Note the wire braid soldered to the ground of the can for DC input. Coil positioning directly affects the measured Q and must be carefully adjusted.

The resonator is implemented in a bifilar configuration, such that the two RF outputs are carried on separate but closely matched conductors. This allows a DC offset to be applied independently to each blade through a bias-tee while maintaining a common RF drive on both electrodes. In this way, the DC potentials on the two blades can be tuned separately; for example, for micromotion compensation without requiring a separate RF resonator for each electrode.

The two helices are interlaced so that both RF paths experience nearly identical inductance, capacitance, and mutual coupling to the shield. As a result, the electrical length and resonant response of the two outputs are closely matched, which minimizes relative phase and amplitude errors between the RF voltages delivered to the two blades. This symmetry is important, since any phase mismatch between the RF electrodes can distort the intended trapping pseudopotential and displace the RF null.

To further suppress differential RF motion between the two blades, a 350 pF shunt capacitor is placed directly between the two RF electrodes. Because this capacitance



Figure 4.6: Aluminum pieces on the front and bottom of the can machined for mechanical stabilization.

is large compared to the electrode to ground capacitance at the operating frequency, it strongly penalizes any voltage difference between the two blades. In the desired common mode operation, where both blades oscillate with the same RF phase and nearly the same amplitude, the voltage across this capacitor is small and it carries little current. By contrast, any phase or amplitude mismatch produces a differential voltage across the capacitor, which drives a compensating current and effectively shorts that differential RF component. In this sense, the shunt capacitor stabilizes the two blade voltages against relative phase error and helps enforce a phase matched RF drive.

In the initial resonator-can design, mechanical vibrations were found to have a deleterious effect on trap performance. In particular, clapping or knocking on the optical table would often cause ions to be ejected from the trap. It was also observed that the secular frequency drifted at the level of several kilohertz per minute, which is highly undesirable

for motional control experiments. This behavior was ultimately traced to mechanical motion affecting the grounding connection between the resonator can and the trap.

To mitigate this issue, additional mechanical components were machined to rigidly attach the resonator can to the octagon using screw-mounted connections, as shown in Figure 4.6. The can was also mounted on a machined angle block to provide reproducible and mechanically robust alignment to the RF feedthrough wires. Finally, the assembly was placed on vibration-damping rubber feet to further suppress mechanical perturbations. These modifications improved the passive stability of the secular frequency by at least three orders of magnitude. Following these improvements, the trap secular frequency was observed to be stable at the level of approximately 10 Hz/hour.

4.4 Vacuum Chamber

The vacuum chamber consists of an 8-inch stainless steel octagonal chamber manufactured by Kimball Physics. The chamber is equipped with two titanium sublimation pumps and one Agilent ion pump to maintain ultra-high-vacuum conditions. The top and bottom faces of the octagon are fitted with re-entrant windows, which reduce the working distance between the trapped ions and the external objective lens used for imaging. This geometry improves fluorescence collection efficiency by allowing the imaging optics to be positioned closer to the trap center. To mitigate charge accumulation on the glass surfaces, each re-entrant window is covered with a fine stainless steel mesh. This helps reduce stray electric fields near the trap region that could otherwise perturb the ion position and degrade trapping stability. The re-entrant windows and the 2.75 inch windows on the sides of the chamber all were coated with anti-reflected coatings optimized for 493 nm light.

The chamber components were first heat treated at 1000 °C in vacuum and then at 1500 °C in air, with the goal of forming an oxide layer to suppress hydrogen diffusion out of the bulk stainless steel. In order to preserve this oxide layer, the initial bakeout was performed at a relatively low temperature of 120 °C. The chamber was baked for

Table 4.2: Selected components of the trap and vacuum system.

Description	Vendor	Part Number
8-inch stainless steel octagon	Kimball Physics	MCF800-ExtOct-G2C8A16
Ion pump	Agilent	9191240
2x Titanium sublimation pump	Duniway	TSP-275-003
Re-entrant fused silica viewport, 6.00" CF	MPF Products, Inc.	Q22878-1 Rev A
2.69 fused silica viewport	MPFPI	A0650-4-CF
Ba oven	Alfa Vacuo	—
XGS-600 ion gauge cable	Ideal Vacuum	P104127

approximately 2.5 months before the base pressure was reached.

After the bakeout, the titanium sublimation pumps were outgassed for 15 minutes while the chamber remained connected to a turbo pump. The ion pump was then turned on and the chamber valve was closed. The chamber pressure during the bake remained in the 10^{-9} to 10^{-8} range, depending on gauge location, after which the system was allowed to return to room temperature. After cooling down, the final pressure reading measure with a cold cathode ion gauge was 4×10^{-10} .

Following the bakeout, the Ba and Sr ovens were activated by gradually increasing their drive currents while monitoring the chamber pressure. The Ba oven exhibited a large pressure rise during activation and showed the onset of emission at a current near 3 A, indicating an operating current at or below this value. The Sr oven required a higher current, showing a smaller activation-related pressure rise and an emission threshold near 3.5 A, indicating an operating current of approximately 3.5 A.

After the initial bakeout, the chamber had to be vented and was subsequently rebaked directly on the optical table using fiberglass insulation and resistive heaters. This second bake was carried out at 160 °C for approximately one month. At the conclusion of the bake, the pressure measured by the ion gauge was 2×10^{-10} . The titanium sublimation pumps were then fired once per day for the following several months and continue to be fired weekly to maintain low background pressure.




Figure 4.7: Pressure data from the first bake. The occasional spikes are outgassing from firing titanium sublimation pumps.

4.5 Ion Sources and Loading Procedures

A secondary goal of the second-generation trap was to demonstrate trapping of $^{133}\text{Ba}^+$ using source activities below $10\ \mu\text{Ci}$ ($\approx 40\ \text{ng}$). In the United States, the purchase and possession of radioactive materials are regulated by isotope-specific activity thresholds. For $^{133}\text{Ba}^+$, the U.S. Nuclear Regulatory Commission's (NRC) exempt quantity is $10\ \mu\text{Ci}$. Quantities at or below this limit are much easier to obtain because they can be possessed without a specific radioactive materials license. Above this threshold, acquisition and use generally must be covered by an institutional radioactive materials license. Demonstrating trapping with a source below the exempt limit was intended to reduce the barrier to adoption by other research laboratories.

The radioactive barium sample was purchased from Eckert & Ziegler. Oak Ridge National Laboratory was also considered as a supplier, but no material was obtained from them. The $^{133}\text{Ba}^+$ was produced by neutron irradiation of an enriched $^{132}\text{Ba}^+$ target, although the enrichment of the $^{132}\text{Ba}^+$ is not known. The material was supplied as BaCl_2






Figure 4.8: CAD rendering of the vacuum chamber with sliced view.

in 0.1 M HCl and stabilized with a barium carrier at a concentration of $10 \mu\text{g Ba/mL}$. The total amount of carrier added is not known. To preserve isotopic purity, it is preferable to avoid adding carrier barium after production. After extensive loading from targets prepared using only this radioactive sample, the isotopic composition was measured and is summarized in Table 4.3.

Table 4.3: Measured isotopic distribution of barium in the sample.

Isotope	Abundance (%)
$^{132}\text{Ba}^+$	12.8
$^{133}\text{Ba}^+$	2.1
$^{134}\text{Ba}^+$	4.3
$^{135}\text{Ba}^+$	3.2
$^{136}\text{Ba}^+$	8.5
$^{137}\text{Ba}^+$	2.1
$^{138}\text{Ba}^+$	67.0

Multiple ion sources with different chemical form factors and geometries were placed



(a) Position of the six-shooter target and barium ovens. (b) Six-shooter target with multiple radioactive BaCl_2 and naturally abundant BaCl_2 sources.

Figure 4.9: Views of the six-shooter target assembly and its position relative to the trap.

inside the chamber to increase the likelihood of successful loading. The primary target was a titanium “six-shooter” style cylinder containing six recessed wells, each loaded with a different source material, as shown in Figure 4.9b. This target was mounted to the blade holder and ablated along the trap axis, with its position shown in Figure 4.9a.

The target was positioned in front of four large platinum filter electrodes. These electrodes were intended to provide a substantial bias voltage so that any charged particles produced during ablation would be deflected away from the trap and collected on the oppositely biased electrode. In principle, this configuration should reduce unwanted charging caused by the ablation process without significantly affecting the production or transport of neutral atoms. However, the effectiveness of this charge-mitigation scheme has not yet been systematically characterized.

The chemical compositions of the ablation targets are listed in Table 4.4. A large

macroscopic target of naturally abundant barium was included to demonstrate ion loading and to characterize trap performance before attempting loading with the radioactive isotope. A microscopic naturally abundant barium target was also included as a proxy for the radioactive target, allowing tests of loading behavior and target depletion with a comparable amount of barium. Additional targets containing mixtures of radioactive and naturally abundant barium were prepared to dilute the $^{133}\text{Ba}^+$ and thereby reduce the risk of depleting the radioactive material too quickly. In practice, however, only naturally abundant barium was loaded from the six-shooter targets. A likely explanation is that naturally abundant barium migrated between neighboring wells, substantially diluting the radioactive material. An important lesson from this result is that microscopic quantities of $^{133}\text{Ba}^+$ should not be distributed throughout a shared target structure.

Fortunately, two additional targets consisting of thin metal foils loaded only with radioactive barium were also installed in the chamber. These targets were prepared by depositing droplets of the radioactive barium solution onto the center of each foil and allowing the solution to dry. During preparation, a small current was passed through the foil to provide mild heating, though not enough to approach the boiling point of the solution. The foils were mounted to vacuum feedthrough wires and could, in principle, be operated as ovens, although they have not been used in this way.

All $^{133}\text{Ba}^+$ loaded in this trap has come from these two targets, even though they were not originally intended to serve as the primary loading targets. As a result, they were not placed in an optimal position for ablation. The ablation beam must therefore be steered past the side of the blades and through the opposite window in order to reach them. In addition, no filter electrodes were placed around these targets. A view of these targets inside the trap is shown in Figure 4.10.

Ions are loaded into the second-generation trap by photoionization of neutral barium. This is accomplished using a two-step excitation process consisting of a resonant 413 nm photon followed by an autoionizing 530.8 nm photon. A detailed discussion of the photoionization schemes explored in this work is given in Chapter 7. Although it is

Table 4.4: Ablation Targets

Substrate	Chemical Compositions	Geometry
Titanium	Naturally Abundant $\text{BaCl}_2 > 1 \text{ g}$	Six-Shooter
Titanium	Naturally Abundant $\text{BaCl}_2 \approx 50 \text{ ng}$	Six-Shooter
Titanium	1.5 mg CaCl_2 and $1.66 \text{ } \mu\text{Ci } ^{133}\text{Ba}^+$	Six-Shooter
Titanium	$.5 \text{ mg BaCl}_2$ and $1.66 \text{ } \mu\text{Ci } ^{133}\text{Ba}^+$	Six-Shooter
Titanium	$.033 \text{ mg SrCa}_2$ and $1.66 \text{ } \mu\text{Ci } ^{133}\text{Ba}^+$	Six-Shooter
Titanium	$1.66 \text{ } \mu\text{Ci } ^{133}\text{Ba}^+$	Six-Shooter
Platinum	$1.66 \text{ } \mu\text{Ci } ^{133}\text{Ba}^+$	Foil-Filament
Rhenium	$1.66 \text{ } \mu\text{Ci } ^{133}\text{Ba}^+$	Foil-Filament

possible to selectively load only the odd barium isotopes by carefully controlling the power and frequency of the 413 nm laser, in this experiment all isotopes were intentionally loaded. One motivation for this choice was the preference to strongly drive the first step transition by saturating and power broadening it, thereby increasing the range of atomic velocities that could be photoionized. A second motivation was the possibility that other barium isotopes could sympathetically cool $^{133}\text{Ba}^+$, thereby improving the probability of successfully loading that isotope.

During ablation, the 493 nm Doppler cooling laser is red-detuned by approximately 80 MHz. We aim to cool as many barium isotopes as possible, with particular emphasis on $^{133}\text{Ba}^+$. The 493 nm laser is therefore parked red-detuned from the $F = 1 \rightarrow F = 0$ transition at approximately 607.429875 THz, while the 650 nm laser is tuned near resonance with the $F = 1 \rightarrow F = 0$ transition at approximately 461.312345 THz. This forms an almost closed Λ system and therefore requires very little power in the EOM sidebands to cool $^{133}\text{Ba}^+$. In practice, a single sideband at 11.727 GHz is typically sufficient to produce a bright $^{133}\text{Ba}^+$ ion.

To cool and image the even isotopes of barium, an additional EOM sideband at 4.280 GHz is applied to the 493 nm laser and a sideband at 990 MHz is applied to the 650 nm laser. Under these conditions, $^{138}\text{Ba}^+$ is typically very bright, while $^{136}\text{Ba}^+$ and $^{132}\text{Ba}^+$ are comparatively dim. During initial trapping, dark ions are also periodically checked for possible identification as $^{137}\text{Ba}^+$ by applying an additional EOM sideband at




Figure 4.10: View through the chamber window of the two pure radioactive targets.

6.350 GHz, or as $^{135}\text{Ba}^+$ by applying a sideband at 6.0 GHz. No non-barium ions have been observed immediately after loading. However, after trapping, ions occasionally turn dark, likely due to reactions with residual hydrogen.

Due to the low abundance of $^{133}\text{Ba}^+$ in the target, hundreds of ablation shots are typically required before a $^{133}\text{Ba}^+$ ion is observed. In practice, loading is more successful when the trap is not emptied after repeated attempts. Instead, ions are allowed to accumulate until the trap contains too many ions to crystallize. At that point, the trap RF is turned off, the trap is emptied, and loading is resumed. The motivation for this approach is that, if a $^{133}\text{Ba}^+$ ion is photoionized, it can be sympathetically cooled by the ions already present in the trap. Once a $^{133}\text{Ba}^+$ ion is identified, a distillation procedure is used to isolate it. This process has proven to be highly reliable. The ion cloud is first crystallized, after which the EOM sidebands are removed so that only $^{133}\text{Ba}^+$ remains Doppler cooled. The trap RF amplitude is then lowered slowly until all other ions are

ejected while the $^{133}\text{Ba}^+$ ion remains trapped.

Several techniques were found to improve the success of this procedure. First, the ion chain is micromotion compensated as well as possible so that the $^{133}\text{Ba}^+$ ion remains near the trap center as the RF is reduced. Second, the power in the Doppler cooling beams is lowered, and a second 493 nm beam detuned by 100 MHz from the main cooling beam is added. Finally, the $^{133}\text{Ba}^+$ ion can typically be positioned nearest the center of the trap and can often be encouraged to hop toward the center during distillation. One possible explanation is that a weak dipole force from the red-detuned 493 nm beam pulls the $^{133}\text{Ba}^+$ ion toward the beam focus. In this way, the non- $^{133}\text{Ba}^+$ isotopes are selectively ejected, leaving a single $^{133}\text{Ba}^+$ ion. To load multiple $^{133}\text{Ba}^+$ ions, the isolated ion is then shuttled away from the loading zone and the loading procedure is repeated.

4.6 Imaging

Maximizing light-collection efficiency is essential for high-fidelity state detection. The more rapidly enough photons can be collected to distinguish bright and dark states, the lower the probability that population shelved in the $D_{5/2}$ state decays during detection. The blade-trap geometry permits a large solid angle for fluorescence collection. To take advantage of this, a custom objective with numerical aperture 0.6 was designed by Photon Gear, Inc. for collection of 493 nm light. This objective is mounted below the trap, and the collected light is relayed and collimated before being focused onto either a Hamamatsu PMT (H10682-210) or a Thorlabs sCMOS camera.

The objective is mounted on a three-axis translation stage for precise alignment. Its working distance is 20 cm from the ion. Two irises placed at intermediate focal points along the imaging path serve as spatial filters. A full Zemax model of the light-collection system was used to determine the position of each optic and to model the point-spread function, as shown in Figure 4.11. A second objective is mounted above the trap on the top viewport. This objective has a numerical aperture of 0.42 and is an Edmund Optics lens




Figure 4.11: Zeemax of the light collection system for the .6 NA objective.

(Part No. 59-878). For cost reasons, only one high-NA objective was purchased, although light collection could be improved by replacing the upper objective with a second 0.6 NA objective.

The theoretical maximum geometric collection efficiency is approximately 10% for the 0.6 NA objective and 4.6% for the 0.42 NA objective. Assuming a PMT quantum efficiency of approximately 20%, the combined system would ideally detect about 1000 counts/ms for a maximally scattering ion. Experimentally, the largest count rates observed for $^{133}\text{Ba}^+$ are approximately 100 counts/ms from the lower objective and 50 counts/ms from the upper objective. For $^{138}\text{Ba}^+$, the corresponding values are approximately 180 counts/ms and 100 counts/ms. Since a maximally scattering ion emits light at a rate of approximately $4.7 \times 10^7 \text{ s}^{-1}$, the observed count rates are lower than the ideal estimate by roughly a factor of three. This reduction is likely due to transmission losses at the vacuum windows and additional losses from the mirrors and lenses in the imaging path.

4.7 Laser Systems

Laser beam paths into the vacuum chamber can be divided into two categories: loading and experimentation. Loading beams are typically made as large as possible to cover the full




Figure 4.12: Full diagram of light delivery to the blade trap.

trap region and are operated well above saturation intensity. In contrast, experimentation beams are designed to minimize scatter and are generally operated at low saturation intensity. A full diagram of the optical paths into the chamber is shown in Figure 4.12.

The beams used for loading, including the 493 nm, 650 nm, 614 nm, and 413 nm beams, are combined on a dichroic stack and then split into axial and radial paths. The radial path is enlarged with a 3:1 telescope and further expanded with a cylindrical lens to cover a larger trap volume. This radial beam is co-propagated with the 531 nm photoionization beam. Before the cylindrical lens, a portion of the radial beam is split off and used as the main experimentation, or “science,” beam. This beam is directed at a 45° angle to the trap axis in order to couple to all three motional modes of the ion. During loading, the axial, radial, and science beams are all turned on. During experimentation, the axial and radial beams are shuttered.

Another important beam path is the “protection” beam, which is turned on when an




Figure 4.13: Laser optical table diagram and AOM configuration.

ion remains dark for a specified period of time. This beam is detuned by 110 MHz from the main 493 nm cooling light. In practice, it can help recool hot ions and improve trap lifetime.

4.7.1 Cooling and Readout Lasers

The primary Doppler cooling lasers at 493 nm and 650 nm are Toptica DL Pro lasers. Each laser is typically stabilized in software by measuring its frequency on a HighFinesse wavemeter approximately every 500 ms and applying slow feedback to the piezo voltage. This lock generally provides a frequency stability on the order of 5 MHz. Shortly after purchase, the 493 nm laser was found to be sensitive to optical feedback from reflections, despite the presence of an internal 30 dB optical isolator. To mitigate this, a second optical isolator was installed immediately after the laser head. This substantially reduced the optical feedback and improved the stability of the laser. A small fraction of the light

from each laser is picked off with a polarizing beamsplitter cube and coupled into fibers that deliver the light to the wavemeter.

Both lasers are sent through ADVR fiber EOMs to generate sidebands and through double-passed Brimrose AOMs that serve as optical switches during experiments. The double-pass configuration suppresses optical leakage through the AOMs. For comparison, a single-pass AOM typically transmits approximately 50 nW of leakage from 50 mW of incident light, whereas the double-pass configuration reduces this leakage below the measurable level. Irises are also placed in both paths to block residual zeroth-order light from entering the fibers. Light is additionally picked off from the main beam path for locking the lasers to the home-built cavity. Finally, the zeroth-order light from the first pass through the AOM is sent to a second AOM and used as an off-resonant, red-detuned “protection” beam.

The shelving lasers at 455 nm and 585 nm are both home-built Littrow-configuration diode lasers. The 455 nm source is based on a nominal 450 nm diode purchased from Osram, which provides a special characterization service that allows selection of a diode with the desired center wavelength. The 585 nm source is generated from an Innolume gain chip operating at 1170 nm with a bandwidth of approximately 100 nm. This infrared light is sent to an ADVR frequency doubler, which produces approximately 100 μ W of 585 nm light from 50 mW of input power at 1170 nm.

The 614 nm deshelving laser is also produced by frequency doubling. In this case, 1228 nm light from a Toptica DL Pro is sent to an ADVR frequency doubler to generate the final wavelength.

4.7.2 Raman Laser System

The Raman laser system was designed to allow rapid switching between control of the 90 MHz metastable clock qubit and the 9.9 GHz ground-state qubit. To provide this flexibility, the required beat note can be generated either with a single double-passed AOM




Figure 4.14: Raman laser diagram for driving motionally dependent transition in $^{133}\text{Ba}^+$ or with an additional free-space resonant EOM manufactured by QUBIG. The same setup can also generate beat notes near 10 MHz, which are used for operations on the Zeeman qubit in $^{138}\text{Ba}^+$.

The delivery geometry of the Raman beams depends on the isotope and transition of interest. For $^{138}\text{Ba}^+$, the beams are propagated orthogonally in order to generate nearly pure π and σ^+ polarizations. The beam geometry for $^{133}\text{Ba}^+$, is shown in Figure 4.15. In this configuration, the two beams are counter-propagating, with one beam horizontally polarized and the other vertically polarized. This arrangement has the advantage that it does not require precise control of the relative σ -polarization content, which can vary with temperature-dependent drifts in the optics.

4.7.3 1762 Stable Laser Systems Cavity

In order to drive high fidelity pulses on the $S_{1/2} \rightarrow D_{5/2}$ optical transition a high power narrow laser is necessary. The 1762 nm laser is stabilized to a high finesse reference cavity housed in a temperature controlled vacuum enclosure inside an acoustic isolation chamber. A small fraction of the Toptica DL Pro output is phase modulated with an EOM, mode matched into the cavity, and used to generate a Pound-Drever-Hall error signal from the




Figure 4.15: Diagram of the internal components of the SLS 1762 nm cavity stabilization system.

reflected light. Fast feedback is applied through a Toptica FALC110 servo to the diode modulation input, while slower feedback is applied through the laser piezo and current controls. The system also includes a fiber noise cancellation for stable delivery through optical fiber and an offset locking module that allows the laser to be stabilized away from the cavity resonance. The measured cavity finesse is 3.68×10^5 , the fractional frequency instability is below 4×10^{-15} at 1 s and below 5×10^{-15} at 10 s, the long-term drift is approximately 11 kHz/day, and the zero-crossing temperature was 27.3°C.

After stabilization the laser light is sent to a two stage optical fiber amplifier manufactured by Cybel. This increases our available laser power from, 20 mW, to at maximum 5 W. We typically do not operate at the highest possible power settings, as multiple components including the EOM fiber, can be damaged at powers above 500 mW. The laser frequency is measured on a Bristol wave meter and set to $170.123220 \pm .000050$ THz. This carrier frequency is chosen as the lock point, such that an EOM sideband at ≈ 5740 MHz can drive the $F = 1 \rightarrow F = 3$ transition, while a tone at ≈ 4268 MHz, can drive the $F = 0 \rightarrow F = 2$ transition.

4.8 Doppler Cooling Cavities

In order to demonstrate high-fidelity SPAM, data must be acquired over many hours while keeping all experimental parameters as constant as possible. One of the first limitations to this long duty cycle was laser drift, particularly rapid drift caused by pressure fluctuations in the laboratory. It was observed that even doors opening in neighboring rooms could measurably shift the Doppler cooling laser frequencies. In practice, opening the lab door often produced a frequency shift of roughly 10 MHz in both Doppler cooling lasers. On an ion, this appears as a large change in the detected fluorescence and can easily be mistaken for an operational error. To make the experiment more robust against these environmental perturbations, both Doppler cooling lasers were locked to a home built optical cavity.

The reference cavity was constructed from Super Invar, a nickel-iron alloy with a very low coefficient of thermal expansion near room temperature. This makes it well suited for cavity construction because small temperature fluctuations produce only small changes in cavity length and therefore only small shifts in the cavity resonance frequency. Although materials such as Zerodur or ULE generally provide better dimensional stability for the highest performance reference cavities, Super Invar offered a practical balance of low thermal expansion, lower cost, and ease of machining. The cavity spacer was machined in house by drilling a 1 in. bore from either side and milling out the central region. The cavity mirrors were aligned while transmitting laser light to ensure propagation of the $\text{TEM}_{0,0}$ mode, then clamped in place and fixed with UV-cured adhesive. A single piezoelectric actuator controls the static cavity length, allowing the resonance to be tuned to the desired cooling frequency. The entire cavity is enclosed in a vacuum housing, which was initially pumped to 1×10^{-6} mbar, although the pressure has likely risen somewhat since the system is not actively pumped. The main purpose of the vacuum housing is to eliminate air currents and reduce vibration sensitivity.

The cavity length is 26.67 cm. It consists of one flat mirror and one concave mirror

with radius of curvature $r = -500$ mm. The mirrors were purchased from Laseroptik and are specified to have 99% reflectivity over the wavelength range 400–650 nm. Each mirror has a diameter of 12.7 mm. Proper mode matching is obtained when the input beam is focused to approximately $200\text{ }\mu\text{m}$ on the flat mirror. With 1 mW of incident light, the typical transmitted signal is 10 mV on a Thorlabs PDA10A2 photodiode. The laser current is modulated at 15 MHz with an RF power of -10 dBm. The photodiode output is then mixed internally in a Liquid Instruments Moku lock box to generate the error signal used for fast feedback to the laser current.

Although an error signal can be generated from either the transmitted or reflected light, locking to the reflected signal is generally preferable. The reflected signal is less sensitive to fluctuations in the incident optical power, whereas the transmitted signal depends directly on the intracavity power and is therefore more susceptible to amplitude noise. In addition, the transmitted signal is filtered by the cavity linewidth, which limits the achievable feedback bandwidth. The reflected signal contains prompt information about the laser detuning and therefore supports substantially higher bandwidth stabilization. At present, the cavity is locked using the transmitted signal, but future work will transition to reflected-signal locking.

The cavity mirrors were intentionally chosen to have relatively low reflectivity and a broad coating bandwidth so that the same cavity could be used to stabilize multiple lasers. The cavity finesse is approximately 200, and the measured free spectral range is 400 MHz. The 650 nm laser is offset locked to the cavity using a tunable 600 MHz EOM tone. The 650 nm and 493 nm lasers are combined and split using two dichroics and detected on separate photodiodes. The cavity transmission for both lasers is monitored digitally, and experimental sequences are automatically halted if either transmission signal drops. This prevents errors caused by laser drift during long data runs.




Figure 4.16: Full schematic of the current microwave drive chain to generate 9.925 GHz.

4.9 Microwave Drive Chain

The microwave drive chain is used to generate a stable 9.925 GHz signal for driving single-qubit gates. Several methods of signal generation were explored and tested. The current component sequence is shown in Figure 4.16. The final signal is produced by mixing a local oscillator set to a fixed frequency of 9.58 GHz with a DDS source that can rapidly control the frequency, phase, and amplitude. Mixing is used for two main reasons. First, the DDS provides phase control, which is important for calibration measurements and dynamical decoupling sequences. Second, it allows arbitrary pulse shaping in amplitude, which helps suppress off-resonant excitation of spectator states.

The main disadvantages of mixing are significant power loss and the introduction of additional spurs and sidebands. In an ideal system, the microwave signal would be generated directly by a high-quality digital oscillator while retaining full phase control. Two microwave switches are required to fully extinguish the signal, with each providing approximately 60 dB of isolation. With only a single switch, significant microwave leakage was observed. The microwaves are delivered with a 10 dB gain horn and are nominally polarized to drive π transitions. In practice, however, reflections within the chamber

reduce the polarization purity. Experimentally, a ratio of approximately 5:1 is observed between the π - and σ -transition Rabi rates.

4.10 Experimental Control Software

LabRAD is a distributed laboratory control framework based on a client-server architecture. Individual hardware devices are controlled by dedicated software servers, while a central manager coordinates communication between them. This design allows many experimental components to be integrated into a single control system, even when they are distributed across multiple computers. In AMO experiments, LabRAD provides a flexible and modular framework for hardware control, automation, data acquisition, and user interface development. It was originally developed in the Martinis group at UC Santa Barbara. The core LabRAD structure is written in Scala, while many of the user interfaces and device-control classes are implemented through the Python package `pylabrad`. Devices such as laser controllers, signal generators, and FPGAs are interfaced through LabRAD servers, which can run on any computer connected by LAN to the main experimental control computer.

Early in this work, the barium lab updated `pylabrad` for Python 3 and adopted ARTIQ as the primary system for FPGA, DDS, and TTL control in time-critical experiments. ARTIQ is a control system designed for precise timing and automation in quantum physics experiments. It combines a Python-based programming environment with dedicated hardware capable of executing deterministic sequences with sub-microsecond timing resolution. In trapped-ion experiments, ARTIQ is particularly useful for coordinating laser pulses, microwave pulses, digital triggers, and data acquisition in a reproducible way. A key feature of the system is that time-critical code is compiled and executed directly on the control hardware, allowing pulse sequences to be synchronized with much lower timing uncertainty than is possible with standard computer-based control alone. In this work, ARTIQ served as the real-time control layer for experimental sequences requiring precise

and repeatable timing.

CHAPTER 5

The Limits of Spontaneous Raman Scattering Errors

“The great modern achievements of applied mathematics have been in relativity and quantum mechanics, and these subjects are at present, at any rate, almost as ‘useless’ as the theory of numbers.” – G. H. Hardy

5.1 Raman Gates

Modern physics has introduced the concept of the truly forbidden. This is distinct from experimentally impractical, unfeasible, or preposterous. A new physical understanding of the limits of natural reality was irrefutably derived from axioms. Einstein’s proof that no mass could be accelerated past the speed of light with any finite force, and Heisenberg’s theorem bounding the uncertainty that could ever be achieved in measurement struck like a shiver through the spine of human philosophy. Despite the urge that every new generation has to buck the previous laws and challenge all aspects of the old guard, these truths are self evident.

The ability to put true physical bounds on what is naturally possible is actually a unique advantage of quantum systems. In macroscopic ensembles of bulk materials, isolating and quantifying individual error mechanisms is often challenging due to uncontrolled environmental coupling and complex many-body effects. However, working with single atoms affords an astonishing degree of theoretical predictability, coupled with a highly accurate means of measurement. The dynamics of individual atoms can be modeled from

first principles, and relevant error channels, such as spontaneous emission, magnetic field noise, and motional decoherence, can be independently characterized and constrained. We can harness this understanding to predict how fruitful certain methodologies for building a quantum computer will be. A central question that demands this level of scrutiny is which physical mechanisms are most suitable for implementing high fidelity single and two-qubit gates.

There are two possible choices to drive coherent Rabi oscillations between qubit states in an ion. The first method, is a direct excitation via a microwave field. The advantage of this method is the excellent electronic phase and amplitude control provided by RF technology. The challenge with microwave control is crosstalk. In view of the fact that typical microwave photons are on the order of centimeters, it is nearly impossible to target a single ion in a chain. In recent years, a tremendous amount of work has been done to ameliorate this issue. Solutions include, but are not limited to, using integrated microwave resonators on chips to address ions in the near field regime ([Warring et al., 2013](#)). This can be combined with the use of multiple microwave electrodes to create regions of constructive and destructive interference, allowing certain ions to be addressed while leaving nearby ions unaffected ([Tarlton et al., 2023](#)). Furthermore, dynamic decoupling schemes have also been demonstrated to prevent unwanted entanglement generation between neighboring ions ([Smith et al., 2024](#)).

The second choice is to use stimulated Raman transitions to drive single and two qubit gates. A principal advantage of this method is that it relies on optical fields, which can be tightly focused to micrometer scales, thereby enabling spatially selective addressing and significantly reducing crosstalk between neighboring ions. This approach is also particularly amenable to integrated photonic technologies, which could be a boon in scaling. In a stimulated Raman transition, two electronic states are coherently coupled through an intermediary excited state. This is achieved using two optical fields that are each detuned from the excited state, while their frequency difference is tuned to the energy splitting of the target states. The resulting two-photon process is often described as excitation

through a “virtual” intermediate state. Although in reality, a very small population is excited and de-excited from the off resonant excited state. This allows one to use strong dipole allowed transitions to couple two states that may be dipole forbidden.

The stimulated Raman effect was first seen by E. J. Woodbury and Gisela Eckhardt in 1962, at Hughes Research Lab in Malibu. During an experiment with a Q-switched ruby laser, an infrared emission line was observed, which could not be explained by ordinary fluorescence or linear optics. The exceptionally high intra-cavity intensities produced by Q-switching made the Raman process, which is normally extremely weak because of its small cross section, strong enough to become macroscopically observable, revealing what is now understood as stimulated Raman scattering. It was first used for coherent sideband cooling and single qubit gates by the Wineland group ([Monroe et al., 1995b,a](#)).

5.2 Single Qubit Gates

We can now derive the dynamics of stimulated Raman gates. First we begin with two optical fields,

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}_1(\mathbf{r}, t) + \mathbf{E}_2(\mathbf{r}, t), \quad (5.1)$$

$$\mathbf{E}_j(\mathbf{r}, t) = \frac{1}{2} \left[\boldsymbol{\mathcal{E}}_j e^{i(\mathbf{k}_j \cdot \mathbf{r} - \omega_{L_j} t + \phi_j)} + \boldsymbol{\mathcal{E}}_j^* e^{-i(\mathbf{k}_j \cdot \mathbf{r} - \omega_{L_j} t + \phi_j)} \right], \quad j \in \{1, 2\}. \quad (5.2)$$

Although we use complex notation for convenience, the physical electric field is obtained by taking the real part (equivalently, by adding the complex conjugate), ensuring that $\mathbf{E}(\mathbf{r}, t)$ is real. The atomic Hamiltonian can be modeled as a three level system,

$$H_{\text{atom}} = \hbar\omega_0 |0\rangle\langle 0| + \hbar\omega_1 |1\rangle\langle 1| + \hbar\omega_e |e\rangle\langle e|. \quad (5.3)$$

The interaction term is the familiar dipole operator separated into four terms, for the two optical fields and the levels they couple,

$$\mathbf{d} = \mathbf{d}_{e0} |e\rangle\langle 0| + \mathbf{d}_{e1} |e\rangle\langle 1| + \mathbf{d}_{0e} |0\rangle\langle e| + \mathbf{d}_{1e} |1\rangle\langle e|, \quad \mathbf{d}_{0e} = \mathbf{d}_{e0}^\dagger, \quad \mathbf{d}_{1e} = \mathbf{d}_{e1}^\dagger. \quad (5.4)$$

$$\begin{aligned} H_{\text{int}}(t) = - \sum_{j=1}^2 \frac{1}{2} \Big[& (\mathbf{d}_{e0} \cdot \boldsymbol{\mathcal{E}}_{L_j}) e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |e\rangle\langle 0| + \left(\mathbf{d}_{e0} \cdot \boldsymbol{\mathcal{E}}_{L_j}^* \right) e^{-i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |e\rangle\langle 0| \\ & + (\mathbf{d}_{e1} \cdot \boldsymbol{\mathcal{E}}_{L_j}) e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |e\rangle\langle 1| + \left(\mathbf{d}_{e1} \cdot \boldsymbol{\mathcal{E}}_{L_j}^* \right) e^{-i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |e\rangle\langle 1| \\ & + (\mathbf{d}_{0e} \cdot \boldsymbol{\mathcal{E}}_{L_j}) e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |0\rangle\langle e| + \left(\mathbf{d}_{0e} \cdot \boldsymbol{\mathcal{E}}_{L_j}^* \right) e^{-i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |0\rangle\langle e| \\ & + (\mathbf{d}_{1e} \cdot \boldsymbol{\mathcal{E}}_{L_j}) e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |1\rangle\langle e| + \left(\mathbf{d}_{1e} \cdot \boldsymbol{\mathcal{E}}_{L_j}^* \right) e^{-i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} |1\rangle\langle e| \Big]. \end{aligned} \quad (5.5)$$

From here we can move to the interaction picture. We define the unitary operator

$$U_0(t) = e^{-iH_{\text{atom}}t/\hbar}. \quad (5.6)$$

The interaction-picture Hamiltonian is given by

$$H_I(t) = U_0^\dagger(t) H_{\text{int}}(t) U_0(t). \quad (5.7)$$

The atomic basis states transform as

$$U_0(t)|e\rangle = e^{-i\omega_e t}|e\rangle. \quad (5.8)$$

Therefore, a generic operator $|e\rangle\langle 0|$ transforms as

$$U_0^\dagger(t) |e\rangle\langle 0| U_0(t) = e^{i(\omega_e - \omega_0)t} |e\rangle\langle 0|. \quad (5.9)$$

Thus,

$$\begin{aligned}
H_I(t) &= U_0^\dagger(t) H_{\text{int}}(t) U_0(t) \\
&= -\frac{1}{2} \sum_{j=1}^2 \left\{ \left(\mathbf{d}_{e0} \cdot \boldsymbol{\mathcal{E}}_j \right) e^{i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{+i(\omega_e - \omega_0)t} |e\rangle\langle 0| + \left(\mathbf{d}_{e0} \cdot \boldsymbol{\mathcal{E}}_j^* \right) e^{-i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{+i(\omega_e - \omega_0)t} |e\rangle\langle 0| \right. \\
&\quad + \left(\mathbf{d}_{e1} \cdot \boldsymbol{\mathcal{E}}_j \right) e^{i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{+i(\omega_e - \omega_1)t} |e\rangle\langle 1| + \left(\mathbf{d}_{e1} \cdot \boldsymbol{\mathcal{E}}_j^* \right) e^{-i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{+i(\omega_e - \omega_1)t} |e\rangle\langle 1| \\
&\quad + \left(\mathbf{d}_{0e} \cdot \boldsymbol{\mathcal{E}}_j \right) e^{i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{-i(\omega_e - \omega_0)t} |0\rangle\langle e| + \left(\mathbf{d}_{0e} \cdot \boldsymbol{\mathcal{E}}_j^* \right) e^{-i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{-i(\omega_e - \omega_0)t} |0\rangle\langle e| \\
&\quad \left. + \left(\mathbf{d}_{1e} \cdot \boldsymbol{\mathcal{E}}_j \right) e^{i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{-i(\omega_e - \omega_1)t} |1\rangle\langle e| + \left(\mathbf{d}_{1e} \cdot \boldsymbol{\mathcal{E}}_j^* \right) e^{-i(\mathbf{k}_j \cdot \mathbf{r} - \omega_j t + \phi_j)} e^{-i(\omega_e - \omega_1)t} |1\rangle\langle e| \right\}.
\end{aligned} \tag{5.10}$$

For convenience we can now collect terms,

$$e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} - \omega_{L_j} t + \phi_{L_j})} e^{+i(\omega_e - \omega_0)t} = e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} + \phi_{L_j})} e^{-i(\omega_{L_j} - \omega_{e0})t}, \quad \omega_{e0} \equiv \omega_e - \omega_0, \tag{5.11}$$

and similarly for the $|e\rangle\langle 1|$ coupling with $\omega_{e1} \equiv \omega_e - \omega_1$. We therefore define the single-photon detunings

$$\Delta_{j0} \equiv \omega_{L_j} - \omega_{e0}, \quad \Delta_{j1} \equiv \omega_{L_j} - \omega_{e1}, \quad j \in \{1, 2\}. \tag{5.12}$$

Next we introduce the single-photon Rabi frequencies for each beam,

$$\Omega_{j0} \equiv -\frac{\mathbf{d}_{e0} \cdot \boldsymbol{\mathcal{E}}_j}{\hbar}, \quad \Omega_{j1} \equiv -\frac{\mathbf{d}_{e1} \cdot \boldsymbol{\mathcal{E}}_j}{\hbar}. \tag{5.13}$$

To arrive at an analytical result we make the rotating-wave approximation (RWA), we retain only the near-resonant terms oscillating as $e^{-i\Delta t}$ and drop the counter-rotating contributions oscillating. It should be noted that these fast oscillating terms can have an




Figure 5.1: Energy level diagram of stimulated Raman transition.

effect and they can be treated analytically. The interaction-picture Hamiltonian becomes

$$H_I^{(\text{RWA})}(t) = \frac{\hbar}{2} \sum_{j=1}^2 \left[\Omega_{j0} e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} + \phi_{L_j})} e^{-i\Delta_{j0}t} |e\rangle\langle 0| + \Omega_{j1} e^{i(\mathbf{k}_{L_j} \cdot \mathbf{r} + \phi_{L_j})} e^{-i\Delta_{j1}t} |e\rangle\langle 1| + \text{h.c.} \right]. \quad (5.14)$$

We now employ a technique known as adiabatic elimination ([White, 2001](#)), where we make the assumption that the dynamics occurring in the excited state are very fast and generally do not affect the slower timescale population change in the target states.

Starting from the RWA interaction-picture Hamiltonian, we will assume beam 1 predominantly couples $|0\rangle \leftrightarrow |e\rangle$ and beam 2 predominantly couples $|1\rangle \leftrightarrow |e\rangle$. Writing $\Phi_j(\mathbf{r}) \equiv \mathbf{k}_j \cdot \mathbf{r} + \phi_j$, the Hamiltonian may be written as

$$H_I^{(\text{RWA})}(t) = \frac{\hbar}{2} \left[\Omega_{L_1} e^{i\Phi_{L_1}(\mathbf{r})} e^{-i\Delta t} |e\rangle\langle 0| + \Omega_{L_2} e^{i\Phi_{L_2}(\mathbf{r})} e^{-i(\Delta + \delta)t} |e\rangle\langle 1| + \text{h.c.} \right], \quad (5.15)$$

where Δ is the single-photon detuning and $\delta \equiv (\omega_1 - \omega_2) - \omega_{01}$ is the two-photon detuning

from the qubit splitting ω_{01} . We can write the state as

$$|\psi(t)\rangle = c_0(t) |0\rangle + c_1(t) |1\rangle + c_e(t) |e\rangle. \quad (5.16)$$

The Schrödinger equation $i\hbar\partial_t |\psi\rangle = H_I^{(\text{RWA})} |\psi\rangle$ yields the coupled differential equations for the probability amplitudes of our three states,

$$i\dot{c}_e = \frac{\Omega_{L1}}{2} e^{i\Phi_{L1}} e^{-i\Delta t} c_0 + \frac{\Omega_{L2}}{2} e^{i\Phi_{L2}} e^{-i(\Delta+\delta)t} c_1, \quad (5.17)$$

$$i\dot{c}_0 = \frac{\Omega_{L1}^*}{2} e^{-i\Phi_{L1}} e^{+i\Delta t} c_e, \quad (5.18)$$

$$i\dot{c}_1 = \frac{\Omega_{L2}^*}{2} e^{-i\Phi_{L2}} e^{+i(\Delta+\delta)t} c_e. \quad (5.19)$$

To remove the time dependence, we can redefine variables (this is equivalent to another unitary transformation),

$$\tilde{c}_e(t) \equiv c_e(t) e^{+i\Delta t}, \quad \tilde{c}_1(t) \equiv c_1(t) e^{-i\delta t}. \quad (5.20)$$

In terms of $(c_0, \tilde{c}_1, \tilde{c}_e)$, Eqs. (5.17)–(5.19) become

$$i\dot{\tilde{c}}_e + \Delta \tilde{c}_e = \frac{\Omega_{L1}}{2} e^{i\Phi_{L1}} c_0 + \frac{\Omega_{L2}}{2} e^{i\Phi_{L2}} \tilde{c}_1, \quad (5.21)$$

$$i\dot{c}_0 = \frac{\Omega_{L1}^*}{2} e^{-i\Phi_{L1}} \tilde{c}_e, \quad (5.22)$$

$$i\dot{\tilde{c}}_1 - \delta \tilde{c}_1 = \frac{\Omega_{L2}^*}{2} e^{-i\Phi_{L2}} \tilde{c}_e. \quad (5.23)$$

Next, we will justify the processes of adiabatically eliminating the excited state level. First lets analyze 5.21, this is a regular first order driven differential equation, we can define the drive term as

$$F(t) \equiv \frac{\Omega_{L1}}{2} e^{i\Phi_{L1}(\mathbf{r})} c_0(t) + \frac{\Omega_{L2}}{2} e^{i\Phi_{L2}(\mathbf{r})} \tilde{c}_1(t), \quad (5.24)$$

so that Eq. (5.21) can be written compactly as

$$i\dot{\tilde{c}}_e(t) + \Delta \tilde{c}_e(t) = F(t). \quad (5.25)$$

Dividing by i gives the standard driven first-order form

$$\dot{\tilde{c}}_e(t) - i\Delta \tilde{c}_e(t) = -iF(t). \quad (5.26)$$

Now we multiply Eq. (5.26) by the integrating factor $e^{-i\Delta t}$:

$$e^{-i\Delta t} \dot{\tilde{c}}_e(t) - i\Delta e^{-i\Delta t} \tilde{c}_e(t) = -iF(t)e^{-i\Delta t}. \quad (5.27)$$

Using the product rule, the left hand side becomes

$$\frac{d}{dt}(\tilde{c}_e(t)e^{-i\Delta t}) = -iF(t)e^{-i\Delta t}. \quad (5.28)$$

Integrating from 0 to t and multiplying by $e^{+i\Delta t}$ yields the exact expression

$$\tilde{c}_e(t) = e^{i\Delta t} \tilde{c}_e(0) - i e^{i\Delta t} \int_0^t F(t') e^{-i\Delta t'} dt'. \quad (5.29)$$

Assume $F(t)$ varies slowly compared to the rapid phase $e^{-i\Delta t}$, i.e.

$$\left| \frac{\dot{F}}{F} \right| \ll |\Delta|, \quad \text{which is ensured by } |\Delta| \gg |\Omega_{L_1}|, |\Omega_{L_2}|, |\delta|. \quad (5.30)$$

Then the oscillatory integral in Eq. (5.29) can be evaluated by integration by parts:

$$\begin{aligned} \int_0^t F(t') e^{-i\Delta t'} dt' &= \left[\frac{F(t') e^{-i\Delta t'}}{-i\Delta} \right]_0^t - \int_0^t \frac{\dot{F}(t')}{-i\Delta} e^{-i\Delta t'} dt' \\ &= \frac{F(t) e^{-i\Delta t} - F(0)}{-i\Delta} + \mathcal{O}\left(\frac{\dot{F}}{\Delta^2}\right). \end{aligned} \quad (5.31)$$

Substituting Eq. (5.31) into Eq. (5.29) gives

$$\tilde{c}_e(t) = \frac{F(t)}{\Delta} + e^{i\Delta t} \left[\tilde{c}_e(0) - \frac{F(0)}{\Delta} \right] + \mathcal{O}\left(\frac{\dot{F}}{\Delta^2}\right). \quad (5.32)$$

Now we can get a qualitative understanding about the population of the excited state in this equation. The second term is rapidly oscillating at frequency Δ . For typical initial conditions $\tilde{c}_e(0) = 0$, and on timescales relevant to the ground-state dynamics, this transient averages away when we integrate. Dropping the transient and keeping the leading term yields the adiabatic result

$$\tilde{c}_e(t) \simeq \frac{F(t)}{\Delta} + \mathcal{O}\left(\frac{\dot{F}}{\Delta^2}\right). \quad (5.33)$$

Finally, inserting the explicit form of $F(t)$ gives

$$\tilde{c}_e \simeq \frac{1}{2\Delta} (\Omega_{L_1} e^{i\Phi_{L_1} c_0} + \Omega_{L_2} e^{i\Phi_{L_2} \tilde{c}_1}). \quad (5.34)$$

Substituting this into Eqs. (5.22) and (5.23) yields the effective two-level dynamics

$$i\dot{c}_0 = \frac{|\Omega_{L_1}|^2}{4\Delta} c_0 + \frac{\Omega_{L_1}^* \Omega_{L_2}}{4\Delta} e^{i(\Phi_{L_2} - \Phi_{L_1})} \tilde{c}_1, \quad (5.35)$$

$$i\dot{\tilde{c}}_1 = \left(\delta + \frac{|\Omega_{L_2}|^2}{4\Delta} \right) \tilde{c}_1 + \frac{\Omega_{L_2}^* \Omega_{L_1}}{4\Delta} e^{i(\Phi_{L_1} - \Phi_{L_2})} c_0. \quad (5.36)$$

These equations are generated by the effective Hamiltonian (in the rotating frame where $c_1 \rightarrow \tilde{c}_1$)

$$H_{\text{eff}} = \hbar \frac{|\Omega_{L_1}|^2}{4\Delta} |0\rangle \langle 0| + \hbar \left(\delta + \frac{|\Omega_{L_2}|^2}{4\Delta} \right) |1\rangle \langle 1| + \hbar \left[\frac{\Omega_{L_1} \Omega_{L_2}^*}{4\Delta} e^{i(\Phi_{L_1} - \Phi_{L_2})} |0\rangle \langle 1| + \text{h.c.} \right]. \quad (5.37)$$

It is convenient to identify the AC Stark shifts

$$\Delta_0^{(\text{AC})} = \frac{|\Omega_{L_1}|^2}{4\Delta}, \quad \Delta_1^{(\text{AC})} = \frac{|\Omega_{L_2}|^2}{4\Delta}, \quad (5.38)$$




Figure 5.2: Raman gate Rabi rate at various laser detunings, the graph assumes 10 mW of laser power and 20 μm beam waist.

Now we can take the main result from this, the Raman Rabi frequency Ω_R

$$\Omega_R(\mathbf{r}) = \frac{\Omega_{L_1}\Omega_{L_2}^*}{4\Delta} \quad (5.39)$$

As an example, we will calculate the Rabi rate for a $^{133}\text{Ba}^+$ ion. We choose two co-propagating lasers each with σ^+ polarized electric fields. Using 3.111, we can calculate the relevant dipole matrix element. Using typical experimental parameters for driving gates ($\lambda_{laser} = 532 \text{ nm}$, $P_0 = 10 \text{ mW}$, $w_0 = 10 \mu\text{m}$) yields a Rabi rate of 10.4 KHz, and a corresponding gate time of 8 μs . One important note is that due to the angular momentum coefficients in the equation, the Rabi rate from the $P_{1/2}$ excited state can destructively interfere with the $P_{3/2}$ state. This behavior can be seen in Figure 5.2: between the two P states, the Rabi rates add constructively, while outside this region they interfere destructively. This interference also implies that if the detuning between the $P_{1/2}$ and $P_{3/2}$ states were 0, the Rabi rate would also be zero. This observation, sheds light on the qualitative reason that two lasers which only exert a force that causes a change in the

L angular momentum quantum number, can change the S spin quantum number of an atom. The spin orbit interaction, which is the term that occurs due to the coupling of the electrons spin with it's own angular momentum, allows for Raman gates to be possible.

5.3 Spontaneous Raman Scattering Error

In the above derivation, despite the fact that the excited state was adiabatically eliminated, in reality a small finite population is still present in this state. While the effect of this population on the dynamics is typically small, it can decay due to the short lifetime of the higher energy excited state. This exact mechanism can be the leading source of error in quantum gates.

Scattering is a two-photon process. One photon is absorbed from (or emitted into) the laser field, while the other is emitted into (or absorbed from) the vacuum field. This process can be modeled mathematically by beginning with perturbation theory, taking the limit of infinite time, and finally summing over all possible final states to arrive at Fermi's golden rule to second order. We begin with a review of the derivation of the first-order golden rule for pedagogical purposes (and fun). Fermi's golden rule is typically applied to a time-independent perturbation, for example, an atom coupled to the quantized modes of the electromagnetic field, resulting in spontaneous emission. We consider a Hamiltonian of the form

$$H(t) = H_0 + V(t), \quad (5.40)$$

where H_0 is the unperturbed Hamiltonian and $V(t)$ is a perturbation. Define the interaction-picture state as

$$|\psi_I(t)\rangle = e^{\frac{i}{\hbar}H_0t}|\psi_S(t)\rangle, \quad (5.41)$$

and the interaction-picture perturbation as

$$V_I(t) = e^{\frac{i}{\hbar}H_0t}V(t)e^{-\frac{i}{\hbar}H_0t}. \quad (5.42)$$

The Schrödinger equation becomes

$$i\hbar \frac{d}{dt} |\psi_I(t)\rangle = V_I(t) |\psi_I(t)\rangle. \quad (5.43)$$

Define the time-evolution operator $U_I(t, t_0)$ by

$$|\psi_I(t)\rangle = U_I(t, t_0) |\psi_I(t_0)\rangle. \quad (5.44)$$

It satisfies

$$i\hbar \frac{d}{dt} U_I(t, t_0) = V_I(t) U_I(t, t_0), \quad U_I(t_0, t_0) = \mathbb{I}. \quad (5.45)$$

Integrating from t_0 to t gives

$$U_I(t, t_0) = \mathbb{I} - \frac{i}{\hbar} \int_{t_0}^t dt_1 V_I(t_1) U_I(t_1, t_0). \quad (5.46)$$

We solve this equation iteratively. Substituting the same expression for $U_I(t_1, t_0)$ yields

$$\begin{aligned} U_I(t, t_0) = & \mathbb{I} - \frac{i}{\hbar} \int_{t_0}^t dt_1 V_I(t_1) \\ & + \left(-\frac{i}{\hbar} \right)^2 \int_{t_0}^t dt_1 \int_{t_0}^{t_1} dt_2 V_I(t_1) V_I(t_2) + \mathcal{O}(V^3). \end{aligned} \quad (5.47)$$

This is the Dyson series expanded to second order. Now, suppose the system is initially in the eigenstate $|i\rangle$ at $t_0 = 0$. The amplitude to find the system in a different eigenstate $|f\rangle$ at time t is

$$c_f^{(1)}(t) = \langle f | U_I(t, 0) | i \rangle = -\frac{i}{\hbar} \int_0^t dt_1 \langle f | V_I(t_1) | i \rangle. \quad (5.48)$$

Because we have assumed our unperturbed Hamiltonian is not time dependent we can see that

$$\langle f | V_I(t) | i \rangle = \langle f | e^{iH_0 t/\hbar} V e^{-iH_0 t/\hbar} | i \rangle = e^{i(E_f - E_i)t/\hbar} \langle f | V | i \rangle = -\frac{i}{\hbar} V_{fi} \int_0^t dt_1 e^{i\omega_{fi}t_1}. \quad (5.49)$$

Where we have defined

$$V_{fi} \equiv \langle f|V|i\rangle, \quad \omega_{fi} \equiv \frac{E_f - E_i}{\hbar}, \quad (5.50)$$

then evaluating the integral

$$\int_0^t dt_1 e^{i\omega t_1} = e^{i\omega t/2} \frac{2 \sin(\omega t/2)}{\omega}, \quad (5.51)$$

and substituting this into the equation for first-order transition probability, gives

$$P_{i \rightarrow f}(t) = |c_f^{(1)}(t)|^2 = \frac{|V_{fi}|^2}{\hbar^2} \frac{4 \sin^2(\omega_{fi} t/2)}{\omega_{fi}^2}. \quad (5.52)$$

We are interested in the long-time limit relative to the time scale of the oscillations $1/\omega_{fi}$.

We can make use of the Dirac delta identity,

$$\lim_{t \rightarrow \infty} \frac{4 \sin^2(\omega t/2)}{\omega^2} = 2\pi t \delta(\omega). \quad (5.53)$$

Thus, for large t ,

$$P_{i \rightarrow f}(t) \approx \frac{2\pi t}{\hbar} |V_{fi}|^2 \delta(E_f - E_i). \quad (5.54)$$

The transition can occur to more than one final state and if our goal is to find the total rate to any of the levels we can integrate over possible final states,

$$P_i(t) = \int dE_f \rho(E_f) \frac{2\pi t}{\hbar} |V_{fi}|^2 \delta(E_f - E_i) = \frac{2\pi t}{\hbar} |V_{fi}|^2 \rho(E_f), \quad (5.55)$$

where $\rho(E_f)$ is defined as the number of states in the energy interval $(E, E + dE)$. The transition rate is therefore the derivative of this formula, which is the standard form of Fermi's golden rule,

$$\Gamma_{i \rightarrow f} = \frac{dP_i}{dt} = \frac{2\pi}{\hbar} |V_{fi}|^2 \rho(E_i). \quad (5.56)$$

It is important to note, that the state vector under consideration also contain informa-

tion about the quantized field. This is the reason we can enforce energy conservation while the atomic state is changing. Because we are in the Schrodinger picture, there is no explicit time dependence in the field. Let the atomic levels be $|a\rangle$, $|b\rangle$, and $|c\rangle$. Let the laser mode have occupation number n_L , and let $|0_s\rangle$ denote the vacuum of all scattered-photon modes, then

$$\begin{aligned} |i\rangle &= |a\rangle \otimes |n_L\rangle \otimes |0_s\rangle \\ , |f\rangle &= |b\rangle \otimes |n_L\rangle \otimes |0_s\rangle \\ , |k\rangle &= |c\rangle \otimes |n_L\rangle \otimes |0_s\rangle. \end{aligned} \quad (5.57)$$

Now we can proceed to Fermi's golden rule to the second order. We begin with the second order term from the Dyson series,

$$c_f^{(2)}(t) = \left(-\frac{i}{\hbar}\right)^2 \int_0^t dt_1 \int_0^{t_1} dt_2 \langle f | V_I(t_1) V_I(t_2) | i \rangle. \quad (5.58)$$

Insert a complete set of intermediate states,

$$\sum_k |k\rangle \langle k| = \mathbb{I}, \quad (5.59)$$

to obtain

$$c_f^{(2)}(t) = -\frac{1}{\hbar^2} \sum_k \int_0^t dt_1 \int_0^{t_1} dt_2 \langle f | V_I(t_1) | k \rangle \langle k | V_I(t_2) | i \rangle. \quad (5.60)$$

As in the first order derivation, we can pull the non time-dependent interaction terms out of the integral

$$c_f^{(2)}(t) = -\frac{1}{\hbar^2} \sum_k V_{fk} V_{ki} \int_0^t dt_1 \int_0^{t_1} dt_2 e^{i\omega_{fk}t_1} e^{i\omega_{ki}t_2}. \quad (5.61)$$

Performing the t_2 integral first,

$$\int_0^{t_1} dt_2 e^{i\omega_{ki}t_2} = \frac{e^{i\omega_{ki}t_1} - 1}{i\omega_{ki}}, \quad (5.62)$$

so that

$$\begin{aligned} c_f^{(2)}(t) &= -\frac{1}{\hbar^2} \sum_k \frac{V_{fk} V_{ki}}{i\omega_{ki}} \int_0^t dt_1 e^{i\omega_{fk} t_1} (e^{i\omega_{ki} t_1} - 1) \\ &= -\frac{1}{\hbar^2} \sum_k \frac{V_{fk} V_{ki}}{i\omega_{ki}} \int_0^t dt_1 (e^{i\omega_{fi} t_1} - e^{i\omega_{fk} t_1}), \end{aligned} \quad (5.63)$$

where we used

$$\omega_{fk} + \omega_{ki} = \omega_{fi}. \quad (5.64)$$

Therefore,

$$c_f^{(2)}(t) = -\frac{1}{\hbar^2} \sum_k \frac{V_{fk} V_{ki}}{i\omega_{ki}} \left[\frac{e^{i\omega_{fi} t} - 1}{i\omega_{fi}} - \frac{e^{i\omega_{fk} t} - 1}{i\omega_{fk}} \right]. \quad (5.65)$$

It is useful to define

$$A_{fi}(t) \equiv \frac{e^{i\omega_{fi} t} - 1}{\omega_{fi}}, \quad A_{fk}(t) \equiv \frac{e^{i\omega_{fk} t} - 1}{\omega_{fk}}, \quad (5.66)$$

so that

$$c_f^{(2)}(t) = \frac{1}{\hbar^2} \sum_k \frac{V_{fk} V_{ki}}{\omega_{ki}} [A_{fi}(t) - A_{fk}(t)]. \quad (5.67)$$

The transition probability is then

$$P_{i \rightarrow f}^{(2)}(t) = \frac{1}{\hbar^4} \sum_{k, k'} \frac{V_{fk} V_{ki} V_{fk'}^* V_{k'i}^*}{\omega_{ki} \omega_{k'i}} [A_{fi}(t) - A_{fk}(t)] [A_{fi}(t) - A_{fk'}(t)]^*. \quad (5.68)$$

We can now examine the terms in this equation, only one enforces overall energy conservation between the initial and final state,

$$|A_{fi}(t)|^2 = \left| \frac{e^{i\omega_{fi} t} - 1}{\omega_{fi}} \right|^2 = \frac{4 \sin^2(\omega_{fi} t/2)}{\omega_{fi}^2}. \quad (5.69)$$

The other terms are oscillatory in time and contribute transient or nonsecular corrections. They do not produce a term proportional to $t \delta(\omega_{fi})$ unless an intermediate state is itself resonant. This would correspond to using a resonant Raman laser and is obviously not a

realistic use case.

$$A_{fi}(t)A_{fk'}(t)^*, \quad A_{fk}(t)A_{fi}(t)^*, \quad A_{fk}(t)A_{fk'}(t)^*, \quad (5.70)$$

Thus, for the long-time transition, one keeps the secular part of the probability,

$$P_{i \rightarrow f}^{(2)}(t) = \frac{1}{\hbar^4} \left| \sum_k \frac{V_{fk} V_{ki}}{\omega_{ki}} \right|^2 \frac{4 \sin^2(\omega_{fi} t/2)}{\omega_{fi}^2}. \quad (5.71)$$

Taking the long-time limit,

$$\Gamma_{i \rightarrow f}^{(2)} \equiv \lim_{t \rightarrow \infty} \frac{P_{i \rightarrow f}^{(2)}(t)}{t} = \frac{2\pi}{\hbar^4} \left| \sum_k \frac{V_{fk} V_{ki}}{\omega_{ki}} \right|^2 \delta(\omega_{fi}). \quad (5.72)$$

With the above formula, we can now attempt to write the equation for the total scattering rate, but we must first sum over all possible final states. The equation derived thus far is not yet the full story, because we have not accounted for the fact that the electric field operator in the quantized picture has two parts,

$$\hat{\mathbf{E}}(\mathbf{r}) = i \sum_{\mathbf{k}, \lambda} \sqrt{\frac{\hbar \omega_k}{2\epsilon_0 V}} \left[\hat{a}_{\mathbf{k}\lambda} \boldsymbol{\epsilon}_{\mathbf{k}\lambda} e^{i\mathbf{k} \cdot \mathbf{r}} - \hat{a}_{\mathbf{k}\lambda}^\dagger \boldsymbol{\epsilon}_{\mathbf{k}\lambda}^* e^{-i\mathbf{k} \cdot \mathbf{r}} \right]. \quad (5.73)$$

Because of the explicit time ordering of the Dyson series, we must distinguish between the processes where the photon is created versus destroyed first. There will be four terms in the expansion and we will go through them individually. First, we consider the process in which the atom absorbs a laser photon of frequency ω_1 and then emits a photon of

frequency ω_2 . The matrix element for the laser-absorption event is

$$\langle b|V|a\rangle \rightarrow \langle b, n_L - 1, 0|\vec{d} \cdot \vec{E}_L|a, n_L, 0\rangle \quad (5.74)$$

$$= i\sqrt{\frac{\hbar\omega_1}{2V\epsilon_0}} \left\langle b, n_L - 1, 0 \left| \vec{d} \cdot \left(\hat{\epsilon}_L a_{\vec{k}_L} e^{i\vec{k}_L \cdot \vec{r}} + \hat{\epsilon}_L^* a_{\vec{k}_L}^\dagger e^{-i\vec{k}_L \cdot \vec{r}} \right) \right| a, n_L, 0 \right\rangle \quad (5.75)$$

$$= i\sqrt{\frac{\hbar\omega_1}{2V\epsilon_0}} \langle b|\vec{d} e^{i\vec{k}_L \cdot \vec{r}} \cdot \hat{\epsilon}_L|a\rangle \langle n_L - 1|a_{\vec{k}_L}|n_L\rangle \quad (5.76)$$

$$= i\sqrt{\frac{\hbar\omega_1}{2V\epsilon_0}} \sqrt{n_L} \vec{d}_{ba} \cdot \hat{\epsilon}_L. \quad (5.77)$$

The matrix element for the emission event is

$$\langle c|V|b\rangle \rightarrow \langle c, n_L - 1, 1_{\vec{k}_s}|\vec{d} \cdot \vec{E}_s|b, n_L - 1, 0\rangle \quad (5.78)$$

$$= i\sqrt{\frac{\hbar\omega_2}{2V\epsilon_0}} \left\langle c, n_L - 1, 1_{\vec{k}_s} \left| \vec{d} \cdot \left(\hat{\epsilon}_s a_{\vec{k}_s} e^{i\vec{k}_s \cdot \vec{r}} + \hat{\epsilon}_s^* a_{\vec{k}_s}^\dagger e^{-i\vec{k}_s \cdot \vec{r}} \right) \right| b, n_L - 1, 0 \right\rangle \quad (5.79)$$

$$= i\sqrt{\frac{\hbar\omega_2}{2V\epsilon_0}} \langle c|\vec{d} e^{-i\vec{k}_s \cdot \vec{r}} \cdot \hat{\epsilon}_s^*|b\rangle \langle 1_{\vec{k}_s}|a_{\vec{k}_s}^\dagger|0\rangle \quad (5.80)$$

$$= i\sqrt{\frac{\hbar\omega_2}{2V\epsilon_0}} \vec{d}_{cb} \cdot \hat{\epsilon}_s^*. \quad (5.81)$$

The atom can also emit the photon before absorbing the laser photon. In this case, the relevant matrix element for spontaneous emission into the mode $(\vec{k}_s, \hat{\epsilon}_s)$ is

$$\langle b, n_L, 1_{\vec{k}_s}|\vec{d} \cdot \vec{E}_s|a, n_L, 0\rangle = i\sqrt{\frac{\hbar\omega_2}{2V\epsilon_0}} \vec{d}_{ba} \cdot \hat{\epsilon}_s^*. \quad (5.82)$$

The atom then absorbs a photon from the laser beam, and the relevant matrix element is

$$\langle c, n_L - 1, 1_{\vec{k}_s}|\vec{d} \cdot \vec{E}_L|b, n_L, 1_{\vec{k}_s}\rangle = i\sqrt{\frac{\hbar\omega_1}{2V\epsilon_0}} \sqrt{n_L} \vec{d}_{cb} \cdot \hat{\epsilon}_L. \quad (5.83)$$

There are two more terms, the emission (absorption) first from the laser, followed by another emission (absorption) into the field. Both of these processes do not conserve energy typically, but in the case of metastable qubits, these processes can contribute to the scattering rate. For a detailed discussion of this, see (Moore, 2023) and (Moore et al.,

2023a).

Putting all these terms together, we arrive at the Raman-scattering rate into a specific scattered mode $(\vec{k}_s, \lambda)$,

$$\Gamma_{i \rightarrow f}^{(\vec{k}_s, \lambda)} = \frac{2\pi}{\hbar^2} \frac{n_L \omega_L \omega_s}{4V^2 \epsilon_0^2} \left| \sum_k \left[\frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)(\vec{d}_{ki} \cdot \hat{\epsilon}_L)}{\omega_{ki} - \omega_L} + \frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_L)(\vec{d}_{ki} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)}{\omega_{ki} + \omega_s} \right] \right|^2 \delta(\omega_f - \omega_i + \omega_s - \omega_L), \quad (5.84)$$

where $\omega_{ki} \equiv \omega_k - \omega_i$. The total Raman-scattering rate is obtained by summing over all scattered wavevectors and both polarizations:

$$\Gamma_{i \rightarrow f}^{Ram} = \sum_{\lambda} \sum_{\vec{k}_s} \Gamma_{i \rightarrow f}^{(\vec{k}_s, \lambda)}. \quad (5.85)$$

Using

$$\sum_{\vec{k}_s} \rightarrow \frac{V}{(2\pi)^3} \int d^3 k_s, \quad d^3 k_s = k_s^2 dk_s d\Omega = \frac{\omega_s^2}{c^3} d\omega_s d\Omega, \quad (5.86)$$

this becomes

$$\begin{aligned} \Gamma_{i \rightarrow f}^{Ram} &= \frac{2\pi}{\hbar^2} \frac{n_L \omega_L}{4V \epsilon_0^2 (2\pi)^3 c^3} \sum_{\lambda} \int d\Omega \int d\omega_s \omega_s^3 \left| \sum_k \left[\frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)(\vec{d}_{ki} \cdot \hat{\epsilon}_L)}{\omega_{ki} - \omega_L} + \frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_L)(\vec{d}_{ki} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)}{\omega_{ki} + \omega_s} \right] \right|^2 \\ &\quad \times \delta(\omega_f - \omega_i + \omega_s - \omega_L). \end{aligned} \quad (5.87)$$

The δ -function fixes the scattered frequency to

$$\omega_s = \omega_L - (\omega_f - \omega_i) \equiv \omega_{sc}, \quad (5.88)$$

so that

$$\Gamma_{i \rightarrow f}^{Ram} = \frac{n_L \omega_L \omega_{sc}^3}{4\hbar^2 \epsilon_0^2 (2\pi)^2 c^3 V} \sum_{\lambda} \int d\Omega \left| \sum_k \left[\frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)(\vec{d}_{ki} \cdot \hat{\epsilon}_L)}{\omega_{ki} - \omega_L} + \frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_L)(\vec{d}_{ki} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)}{\omega_{ki} + \omega_{sc}} \right] \right|^2. \quad (5.89)$$

The laser mode occupation factor can be given in terms of the classical laser electric field,

$$\vec{E}_L^{(+)}(\vec{r}, t) = i\sqrt{\frac{\hbar\omega_L}{2\epsilon_0 V}} \hat{\epsilon}_L a_{\vec{k}_L} e^{i(\vec{k}_L \cdot \vec{r} - \omega_L t)}, \quad \vec{E}_L^{(-)} = (\vec{E}_L^{(+)})^\dagger. \quad (5.90)$$

Therefore,

$$\langle \vec{E}_L^{(-)} \cdot \vec{E}_L^{(+)} \rangle = \frac{\hbar\omega_L}{2\epsilon_0 V} \langle a_{\vec{k}_L}^\dagger a_{\vec{k}_L} \rangle = \frac{\hbar\omega_L}{2\epsilon_0 V} n_L. \quad (5.91)$$

If we identify the corresponding classical monochromatic field as

$$\vec{E}_{cl}(\vec{r}, t) = \frac{1}{2} \mathcal{E}_L \hat{\epsilon}_L e^{i(\vec{k}_L \cdot \vec{r} - \omega_L t)} + \frac{1}{2} \mathcal{E}_L^* \hat{\epsilon}_L^* e^{-i(\vec{k}_L \cdot \vec{r} - \omega_L t)}, \quad (5.92)$$

then

$$\vec{E}_{cl}^{(+)} = \frac{1}{2} \mathcal{E}_L \hat{\epsilon}_L e^{i(\vec{k}_L \cdot \vec{r} - \omega_L t)}, \quad \vec{E}_{cl}^{(-)} = \frac{1}{2} \mathcal{E}_L^* \hat{\epsilon}_L^* e^{-i(\vec{k}_L \cdot \vec{r} - \omega_L t)}, \quad (5.93)$$

and hence

$$\vec{E}_{cl}^{(-)} \cdot \vec{E}_{cl}^{(+)} = \frac{|\mathcal{E}_L|^2}{4}. \quad (5.94)$$

Equating the classical and quantum field intensities gives

$$\frac{|\mathcal{E}_L|^2}{4} = \langle \vec{E}_L^{(-)} \cdot \vec{E}_L^{(+)} \rangle = \frac{\hbar\omega_L}{2\epsilon_0 V} n_L, \quad (5.95)$$

or

$$|\mathcal{E}_L|^2 = \frac{2\hbar\omega_L n_L}{\epsilon_0 V}. \quad (5.96)$$

Thus

$$\frac{n_L \omega_L}{V \epsilon_0} = \frac{|\mathcal{E}_L|^2}{2\hbar}, \quad (5.97)$$

and the Raman-scattering rate can be written as

$$\Gamma_{i \rightarrow f} = \frac{|\mathcal{E}_L|^2 \omega_{sc}^3}{8\hbar^3 \epsilon_0 (2\pi)^2 c^3} \sum_{\lambda} \int d\Omega \left| \sum_k \left[\frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)(\vec{d}_{ki} \cdot \hat{\epsilon}_L)}{\omega_{ki} - \omega_L} + \frac{(\vec{d}_{fk} \cdot \hat{\epsilon}_L)(\vec{d}_{ki} \cdot \hat{\epsilon}_{\vec{k}_s \lambda}^*)}{\omega_{ki} + \omega_{sc}} \right] \right|^2 \quad (5.98)$$

Now to handle the angular dependence of the integral define,

$$\vec{\mathcal{M}} = \sum_k \left[\frac{\vec{d}_{fk} (\vec{d}_{ki} \cdot \hat{\epsilon}_L)}{\omega_{ki} - \omega_L} + \frac{\vec{d}_{ki} (\vec{d}_{fk} \cdot \hat{\epsilon}_L)}{\omega_{ki} + \omega_{sc}} \right]. \quad (5.99)$$

Then

$$\Gamma_{i \rightarrow f}^{Ram} = \frac{|\mathcal{E}_L|^2 \omega_{sc}^3}{8\hbar^3 \epsilon_0 (2\pi)^2 c^3} \sum_{\lambda} \int d\Omega \left| \vec{\mathcal{M}} \cdot \hat{\epsilon}_{k_s \lambda}^* \right|^2. \quad (5.100)$$

Next, we note that $\{k, \epsilon_{\vec{k}_{s1}}^-, \epsilon_{\vec{k}_{s2}}^-\}$ represent a complete orthonormal basis, thus we can use the identity,

$$\sum_{\lambda} \epsilon_{\vec{k}_s \lambda, i}^- \epsilon_{\vec{k}_s \lambda, j}^{*-} = \delta_{ij} - \hat{k}_{s,i} \hat{k}_{s,j}, \quad (5.101)$$

and obtain

$$\sum_{\lambda} \left| \vec{\mathcal{M}} \cdot \hat{\epsilon}_{k_s \lambda}^* \right|^2 = |\vec{\mathcal{M}}|^2 - |\vec{\mathcal{M}} \cdot \hat{k}_s|^2. \quad (5.102)$$

Therefore,

$$\Gamma_{i \rightarrow f}^{Ram} = \frac{|\mathcal{E}_L|^2 \omega_{sc}^3}{8\hbar^3 \epsilon_0 (2\pi)^2 c^3} \int d\Omega \left(|\vec{\mathcal{M}}|^2 - |\vec{\mathcal{M}} \cdot \hat{k}_s|^2 \right). \quad (5.103)$$

The second term in this integral is simple to evaluate,

$$\int d\Omega \hat{k}_{s,i} \hat{k}_{s,j} = \frac{4\pi}{3} \delta_{ij}. \quad (5.104)$$

Finally, we can write the full Raman scattering equation for a given initial and final state as

$$\Gamma_{i \rightarrow f}^{Raman} = \frac{\mathcal{E}^2 e^4 \omega_{sc}^3}{12\pi c^3 \epsilon_0 \hbar^3} \sum_q \left| \sum_k \left[\frac{\langle f | \vec{r} \cdot \hat{\epsilon}_q | k \rangle \langle k | \vec{r} \cdot \hat{\epsilon} | i \rangle}{\omega_{ki} - \omega} + \frac{\langle f | \vec{r} \cdot \hat{\epsilon} | k \rangle \langle k | \vec{r} \cdot \hat{\epsilon}_q | i \rangle}{\omega_{ki} + \omega_{sc}} \right] \right|^2, \quad (5.105)$$

which is also known as the Kramers-Heisenberg formula. It is important to note that the sum over intermediate states is inside the modulus, meaning interference between scattering channels through different intermediate states can occur. On the other hand, the polarization sum is outside of the modulus; thus, if the scattered photons have different polarizations the terms do not interfere.

5.4 Experimental Verification of the Raman Scattering Rate in $^{133}\text{Ba}^+$

Previous models attempted to estimate the Raman-scattering gate error across many trapped ion candidates. However, due largely to an incorrect assumption, the large-detuning behavior was not treated correctly (Ozeri et al., 2007). Specifically, these models employed what is now referred to as a constant-density of states approximation, in which the scattered photon was effectively assumed to have a fixed frequency, taken to be the frequency difference between the intermediate and final states. However, as the derivation above shows, the scattered photon frequency only appears after summing over all possible final photon modes, and it depends on the actual initial and final atomic states. This has important physical consequences. For example, if the scattered photon does not have enough energy to reach a given final state, then that scattering channel is simply closed. Neglecting this effect causes the theory to fail in the large detuning limit.

The main conclusion of Ref. (Ozeri et al., 2007) was that the optimal detuning for Raman gates lies midway between the $P_{1/2}$ and $P_{3/2}$ states. Within the constant density of states approximation, it was natural to treat the scattered photon frequency as being set by the intermediate state energy, since its true detuning dependence was ignored. In the end, these compounded assumptions led to an incomplete description of the Raman gate error, particularly at large detuning.

The incomplete model previously available in the literature motivated a careful reexamination of the Raman scattering rate in $^{133}\text{Ba}^+$, both theoretically and experimentally. In the calculation presented here, the assumptions responsible for the failure of the earlier large-detuning treatment are removed. In particular, the scattered photon frequency is determined self consistently by energy conservation after summing over all allowed final photon modes, rather than being fixed by the intermediate state energy. This automatically accounts for the closure of energetically forbidden scattering channels at large detuning. In addition, the full set of relevant scattering pathways is retained, including the terms

neglected in earlier treatments, so that interference effects and the correct asymptotic detuning dependence are captured. As a result, the present calculation provides a more complete description of Raman scattering errors in $^{133}\text{Ba}^+$, especially in the regime of large detuning where the previous model breaks down.

Part of the confusion in calculating the Raman scattering rate arises from the fact that it is often written in more than one form. To make the relationship between these expressions clear, both forms are explicitly presented below. The two scattering models used in this work, denoted the “CDA” and “ ω^3 -theory” models, are written using a common notation applicable to ions with similar atomic structure. Throughout, the conventions shown in Figure 5.3 and Table 5.1 are adopted. Einstein A coefficients, electronic angular momentum quantum numbers, and transition frequencies are labeled with subscripts corresponding to the relevant states (e.g., $A_{e',g''}$, J_g , and $\omega_{e,g'}$).

We begin with the constant-density-of-states (CDA) model, which is based on the earlier treatment of Ref. (Ozeri et al., 2007). As discussed above, this approximation is most accurate when the laser detuning from an intermediate state is small compared with the optical frequency, so that the frequency dependence of the scattered photon can be neglected. For a g -type zero field hyperfine clock state qubit, the total spontaneous scattering rate is

$$\Gamma_{\text{total,CDA}} \equiv (A_{e,g} + A_{e,g'}) \frac{g^2}{3} \frac{1}{\Delta^2} + (A_{e',g} + A_{e',g'} + A_{e',g''}) \frac{g^2}{3} \frac{2}{(\Delta - \omega_{\text{FS}})^2}, \quad (5.106)$$

where

$$\Delta = \omega_\ell - \frac{E_{2P_{1/2}^o}}{\hbar} \quad (5.107)$$

is the detuning of the laser from the $2P_{1/2}^o \leftrightarrow 2S_{1/2}$ resonance. Within this model, the spontaneous Rayleigh scattering rate is

$$\Gamma_{\text{Rayleigh,CDA}} \equiv A_{e',g} \frac{g^2}{3} \left(\frac{3\Delta^2 - 2\Delta\omega_{\text{FS}} + \omega_{\text{FS}}^2/3}{\Delta^2(\Delta - \omega_{\text{FS}})^2} \right), \quad (5.108)$$




Figure 5.3: A standardized energy level diagram for driving a stimulated Raman transition in a clock state half integer-spin isotope of an alkaline earth ion.

and the spontaneous Raman scattering rate is obtained from the difference

$$\Gamma_{\text{Raman,CDA}} = \Gamma_{\text{total,CDA}} - \Gamma_{\text{Rayleigh,CDA}}. \quad (5.109)$$

The branching fraction for spontaneous Raman scattering into the $^2\text{D}_{5/2}$ manifold is then

$$\eta_{\text{D}_{5/2},\text{CDA}} \equiv \frac{\Gamma_{\text{D}_{5/2},\text{CDA}}}{\Gamma_{\text{Raman,CDA}}} = \frac{A_{e',g''}}{3} \frac{g^2}{(\Delta - \omega_{\text{FS}})^2}. \quad (5.110)$$

The ω^3 model closely resembles the CDA model near resonance, but differs significantly at large detuning. The essential difference is that the spontaneous scattering rate retains its dependence on the frequency of the scattered photon. As a consequence, this model




Figure 5.4: $^{133}\text{Ba}^+$ level structure. (a) The SRS photon frequency, ω_{sc} , varies depending on the laser frequency, ω_ℓ , and the decay channel (wavy lines). (b) As the gate laser is detuned further to the red of the intermediate resonance ($\omega_{\ell i}$, $i = 1 \rightarrow 3$), the scattering rate to a metastable state (Γ_i , proportional to ω_{sc}^3) decreases until the error channel is closed (see $\omega_{sc}^3 \Theta(\omega_{sc})$ in Eq. 5.111).

must also include step functions that enforce the closure of scattering channels when the laser photons do not carry enough energy to populate a given final state, as illustrated in Figure 1(b). For this model, the total spontaneous scattering rate is

$$\begin{aligned}
 \Gamma_{\text{total}, \omega^3} = & \left(\frac{\omega_\ell}{\omega_{e,g}} \right)^3 A_{e,g} \frac{g^2}{3} \left(\frac{1}{\Delta^2} \right) \\
 & + \Theta(\omega_\ell - \omega_{g,g'}) \left(\frac{\omega_\ell - \omega_{g,g'}}{\omega_{e,g'}} \right)^3 A_{e,g'} \frac{g^2}{3} \left(\frac{1}{\Delta^2} \right) \\
 & + \left(\frac{\omega_\ell}{\omega_{e',g}} \right)^3 A_{e',g} \frac{g^2}{3} \left(\frac{2}{(\Delta - \omega_{FS})^2} \right) \\
 & + \Theta(\omega_\ell - \omega_{g,g'}) \left(\frac{\omega_\ell - \omega_{g,g'}}{\omega_{e',g'}} \right)^3 A_{e',g'} \frac{g^2}{3} \left(\frac{2}{(\Delta - \omega_{FS})^2} \right) \\
 & + \Theta(\omega_\ell - \omega_{g,g''}) \left(\frac{\omega_\ell - \omega_{g,g''}}{\omega_{e',g''}} \right)^3 A_{e',g''} \frac{g^2}{3} \left(\frac{2}{(\Delta - \omega_{FS})^2} \right). \quad (5.111)
 \end{aligned}$$

State	Shorthand	J	Energy ($/\hbar$)
$^2S_{1/2}$	g	$J_g = \frac{1}{2}$	$\equiv 0$
$^2P_{1/2}^o$	e	$J_e = \frac{1}{2}$	$\omega_{e,g}$
$^2P_{3/2}^o$	e'	$J_{e'} = \frac{3}{2}$	$\omega_{e',g}$
$^2D_{3/2}$	g'	$J_{g'} = \frac{3}{2}$	$\omega_{g',g}$
$^2D_{5/2}$	g''	$J_{g''} = \frac{5}{2}$	$\omega_{g'',g}$

Table 5.1: Shorthand notation for the five manifolds considered here.

The spontaneous Rayleigh scattering rate takes a form similar to that in the CDA model,

$$\Gamma_{\text{Rayleigh}, \omega^3} = \left(\frac{\omega_\ell}{\omega_{e',g}} \right)^3 A_{e',g} \frac{g^2}{3} \left(\frac{3\Delta^2 - 2\Delta\omega_{\text{FS}} + \omega_{\text{FS}}^2/3}{\Delta^2(\Delta - \omega_{\text{FS}})^2} \right). \quad (5.112)$$

The spontaneous Raman scattering rate is again given by the difference between the total and Rayleigh rates:

$$\begin{aligned} \Gamma_{\text{Raman}, \omega^3} &= \Gamma_{\text{total}, \omega^3} - \Gamma_{\text{Rayleigh}, \omega^3} \\ &= \left(\frac{\omega_\ell}{\omega_{e',g}} \right)^3 A_{e',g} \frac{2g^2}{9} \left(\frac{\omega_{\text{FS}}}{\Delta(\Delta - \omega_{\text{FS}})} \right)^2 \\ &\quad + \Theta(\omega_\ell - \omega_{g,g'}) \left(\frac{\omega_\ell - \omega_{g,g'}}{\omega_{e',g'}} \right)^3 A_{e',g'} \frac{g^2}{3} \left(\frac{1}{\Delta^2} \right) \\ &\quad + \Theta(\omega_\ell - \omega_{g,g'}) \left(\frac{\omega_\ell - \omega_{g,g'}}{\omega_{e',g'}} \right)^3 A_{e',g'} \frac{g^2}{3} \left(\frac{2}{(\Delta - \omega_{\text{FS}})^2} \right) \\ &\quad + \Theta(\omega_\ell - \omega_{g,g''}) \left(\frac{\omega_\ell - \omega_{g,g''}}{\omega_{e',g''}} \right)^3 A_{e',g''} \frac{g^2}{3} \left(\frac{2}{(\Delta - \omega_{\text{FS}})^2} \right), \end{aligned} \quad (5.113)$$

where we have used the relation

$$\frac{A_{e,g}}{\omega_{e,g}^3} = \frac{A_{e,g'}}{\omega_{e,g'}^3}. \quad (5.114)$$

The branching fraction for spontaneous Raman scattering into the $^2D_{5/2}$ manifold is

$$\eta_{D_{5/2}, \omega^3} = \frac{\Theta(\omega_\ell - \omega_{g,g''}) \left(\frac{\omega_\ell - \omega_{g,g''}}{\omega_{e',g''}} \right)^3 A_{e',g''} \frac{g^2}{3} \left(\frac{2}{(\Delta - \omega_{\text{FS}})^2} \right)}{\Gamma_{\text{Raman}, \omega^3}}. \quad (5.115)$$

Although the ω^3 model represents a clear improvement over the CDA model in the



Figure 5.5: Pulse Sequence Diagram outlining the 6 steps to each experimental cycle. I: The atom is laser cooled for 5ms. II: Using an EOM to phase modulate the 493 light, the ion is optically pumped to the $|S_{1/2}, F = 0\rangle$ state for $150\ \mu s$. III: The atom is then illuminated by the high power 532nm laser when SRS can occur. IV: For $300\ \mu s$, the SRS error states of the atom are electron shelved to the $D_{5/2}$ metastable manifold. V: The cooling lasers are turned back on for 8ms and photon counts are collected to determine whether or not SRS has occurred. VI: A 614nm laser is used at the end to deshelve population back to the cooling manifold to repeat the experiment.

regime explored experimentally in this work, it still neglects several effects that become relevant in other parameter regimes. As discussed in Ref. (Moore et al., 2023a) (Hereafter referred to as the Moore et al. model), additional contributions can arise from emission-first processes, as well as from coupling to excited states beyond the five manifolds shown in Figure 5.3. These effects are not included in the simplified treatment presented here.

A key conceptual point is that spontaneous Raman scattering is not merely a technical nuisance, but an architecture level source of gate infidelity. Earlier CDA based treatments predicted a nonzero lower bound on this error, arising from scattering into metastable D manifolds. By contrast, once the frequency dependence of the scattered photon is retained, the scattering rate acquires the correct ω_{sc}^3 behavior, and energetically forbidden channels close through the condition $\Theta(\omega_{sc})$, as shown in Figure 5.4. This changes the large detuning limit qualitatively and removes the previously predicted lower bound on the Raman gate error.

To experimentally characterize scattering in the far detuned regime, we prepare a single trapped ion in a well defined initial state, expose it to far detuned light, and detect internal state changes produced by spontaneous Raman scattering (SRS). Because such scattering events are intrinsically rare, the measurement must clearly distinguish genuine SRS events from state preparation and measurement (SPAM) errors. We therefore perform the experiment using a single $^{133}\text{Ba}^+$ ion, which enables high fidelity state preparation and readout.

A single $^{133}\text{Ba}^+$ ion is confined in a linear Paul trap with a minimum ion electrode spacing of 3 mm and an RF drive frequency of $\Omega = 2\pi \times 2.6$ MHz, yielding a radial secular frequency of $\omega_r = 2\pi \times 230$ kHz. The ion state is detected by imaging laser-induced fluorescence (LIF) from the $2S_{1/2} \rightarrow 2P_{1/2}$ transition at 493 nm through an objective with numerical aperture 0.28.

To induce off resonant scattering, the ion is illuminated with a Continuum Verdi-V10 laser at 532 nm, focused to a $1/e^2$ intensity radius of approximately $w_0 \approx 40$ μm at the ion position. The optical power is varied between 0.3 W and 1.4 W. This wavelength is suitable for gate operations involving both the $2S_{1/2}$ g -type qubits and the $2D_{5/2}$ m -type qubits in $^{133}\text{Ba}^+$. The laser polarization is set to drive σ^+ transitions, and a 0.5 mT magnetic field is applied at the ion, aligned antiparallel to the laser wavevector $\mathbf{k}$. The laser intensity at the ion is calibrated using a differential AC Stark shift measurement on the o -type clock-state qubit defined by $|2S_{1/2}, F = 1, m_F = 0\rangle$ and $|2D_{5/2}, F = 3, m_F = 0\rangle$ via narrow-linewidth spectroscopy at 1762 nm.

The formula used for conversion from optical intensity to Stark shift is,

$$\delta E_{nFMF}^{(2)} = \frac{1}{4\hbar^2} \sum_i \frac{|\langle m_s | \mathbf{d} \cdot \mathcal{E} \hat{\epsilon} | i \rangle|^2}{\Delta_i} \quad (5.116)$$

Note $\mathcal{E}$ is the electric field amplitude. As stated above, the laser illumination power ranged from .3-1.4 W. Using our measured beam waist of 40 μ we can calculate an electric field amplitude and plug it into equation (5.116). This electric field amplitude range corresponds to optical Stark shifts of .32-1.5 Mhz. This range of Stark shift values directly agree with the measured Stark shift data points in the horizontal axis of Figure 2. The experimental sequence used to determine the SRS rate is shown schematically in Figure 5.5. The ion is first prepared in the $|\uparrow\rangle \equiv |F = 0, m_F = 0\rangle$ g -type clock qubit state in the $2S_{1/2}$ manifold using optical pumping. The ion is then exposed to the 532 nm beam for a controlled duration between 4 ms and 11 ms, during which spontaneous Raman scattering may transfer population out of the initial state.




Figure 5.6: Experimental pulse sequence to measure the total Raman scattering rate.

Following this interaction, a 300 μs pulse of resonant light at 455 nm, 585 nm, and 650 nm is applied to transfer population in the $2S_{1/2}(F=1)$ and $2D_{3/2}$ manifolds to the metastable $2D_{5/2}$ manifold with high fidelity. This electron shelving process is illustrated in Figure 5.6 and converts internal-state changes caused by scattering into a measurable bright or dark outcome.

We then determine whether the ion remains in the initial $|\uparrow\rangle$ state by monitoring fluorescence during 8 ms of Doppler cooling with 493 nm and 650 nm light. An experimental trial that yields no fluorescence is classified as a dark event, indicating that an SRS process has occurred. To correct for background events unrelated to the 532 nm illumination, the same experimental sequence is repeated with the 532 nm beam blocked, and the resulting dark event probability is subtracted from the laser on data to obtain a background corrected SRS probability.

Because the probability of scattering during a single exposure window is small, each




Figure 5.7: SRS rate, Γ , from 532 nm laser illumination (green) versus differential o -type Stark shift, δ . The upper grouping (a) shows the total SRS rate, while the lower grouping (b) shows rate for SRS into only $2D_{5/2}$. (Black circles) Measurements of the scattering rate. (Red, dotted) CDA theory prediction. (Black, solid) Moore *et al.* treatment. (Blue, dashed) ω^3 -theory prediction.

data point is averaged over more than 5×10^4 experimental repetitions to achieve sufficient statistical precision. Under these conditions, the probability of multiple SRS events occurring within a single illumination interval is negligible. The SRS rate into only the $2D_{5/2}$ manifold is measured using the same sequence, except that the shelving step is omitted.

For the total SRS-rate measurement, the observed background event probability is $6(2) \times 10^{-4}$, arising primarily from imperfections in the shelving process. In contrast, when measuring scattering into the $2D_{5/2}$ manifold directly (without shelving), background events are dominated by state-preparation errors, resulting in a reduced background probability of $4(4) \times 10^{-5}$.

Figure 5.7(a) presents the measured total spontaneous Raman scattering (SRS) rate as a function of the differential AC Stark shift of the o -type qubit, while Figure 5.7(b)

Table 5.2: Comparison of theoretical predictions with the measured SRS rate versus laser intensity. Values are given in units of $10^{-9} \text{ s}^{-1}/(\text{W}/\text{m}^2)$.

SRS Type	CDA	ω^3	Moore <i>et al.</i>	Measured
Total (Γ_{SRS})	1.65	1.37	1.52	1.52(5)
$^2\text{D}_{5/2}$ ($\Gamma_{D_{5/2}}$)	0.48	0.25	0.29	0.28(2)

shows the corresponding measurements for scattering into only the $^2\text{D}_{5/2}$ manifold. The scattering rate, Γ_{meas} , is determined from the laser exposure time, τ_ℓ , and the background corrected probability of scattering, P_{meas} , according to

$$\Gamma_{\text{meas}} = -\frac{\ln(1 - P_{\text{meas}})}{\tau_\ell}. \quad (5.117)$$

This expression accounts for the probabilistic nature of rare scattering events over a finite interaction time. The vertical error bars shown in Figure 5.7 are calculated using the Wilson score interval for binomial statistics (Wilson, 1927), providing a reliable estimate of the uncertainty in the measured scattering probabilities. We use the Wilson interval rather than the standard Gaussian approximation because it provides more reliable coverage for small event probabilities and finite sample sizes, which is particularly important in the present rare scattering regime.

The horizontal uncertainties arise from the determination of the differential Stark shift and are given by the standard error in the fitted spectral peak centers, with and without the applied Stark shift, combined in quadrature. These uncertainties are smaller than the plotted data markers.

To quantify the dependence of the scattering rate on the Stark shift, we perform a linear fit using orthogonal distance regression (ODR), which incorporates uncertainties in both the measured scattering rate and the Stark shift. The Stark shift is related to the laser intensity through a calibrated model, allowing the horizontal axis to be interpreted in terms of optical intensity. The resulting fit parameters are summarized in Table 5.2.

The measured data are in agreement, within experimental uncertainty and with the




Figure 5.8: The SRS branching fraction to the $D_{5/2}$ manifold, $\eta_{D_{5/2}}$ as a function of laser frequency. (Black, solid) Moore *et al.* treatment. (Red, dotted) CDA model. (Blue, dashed) ω^3 -theory prediction. The black point shows the result of 21 measurements with the standard error displayed as error bars. The inset shows a zoom-in near the gate-laser frequency.

theoretical treatment of spontaneous Raman scattering presented by Moore *et al.* In contrast, the ω^3 scaling model predicts a systematic offset of approximately 10% in the total SRS rate, which can be attributed to the omission of emission-first contributions included in the Moore *et al.* framework. The experimental results also deviate from the predictions of the CDA model, particularly for scattering into the $^2D_{5/2}$ manifold. This discrepancy is consistent with the fact that transitions to the $^2D_{5/2}$ state involve the emission of a lower-energy photon $\hbar\omega_{sc}$ compared to scattering processes that return the ion to the $^2S_{1/2}$ manifold, thereby amplifying the deviation of the CDA model from the observed data in this channel.

To further investigate the discrepancy between the measured scattering rates and the

predictions of different theoretical models, we examine not only the total SRS rate but also the branching fraction into specific final states. In particular, we measure the fraction of spontaneous Raman scattering events that populate the $^2D_{5/2}$ manifold, defined as

$$\eta_{D_{5/2}} \equiv \frac{\Gamma_{D_{5/2}}}{\Gamma_{\text{SRS}}}, \quad (5.118)$$

where $\Gamma_{D_{5/2}}$ is the scattering rate into the $^2D_{5/2}$ manifold and Γ_{SRS} is the total SRS rate.

This branching fraction provides a particularly useful observable because, unlike the absolute scattering rate, it is largely independent of the laser intensity. Although the differential AC Stark shift is used experimentally to monitor and calibrate the laser intensity at the ion, the ratio $\eta_{D_{5/2}}$ depends only on the relative transition amplitudes and the emission spectrum, and is therefore primarily determined by the laser wavelength and the atomic structure. As a result, it offers a more direct probe of the underlying scattering physics and is less sensitive to systematic uncertainties in intensity calibration.

Figure 5.8 shows the measured branching fraction together with theoretical predictions. A key feature highlighted in the figure is the role of the photonic density of states, which introduces a factor proportional to ω_{sc}^3 in the scattering rate, where ω_{sc} is the frequency of the emitted photon. The inclusion of this factor (black curve) leads to a strong suppression of scattering into channels that require emission of low frequency photons.

The consequences of this frequency dependence are particularly evident when comparing to the CDA model (red dashed curve), which neglects this ω_{sc}^3 scaling. In that model, the branching fraction approaches an asymptotic value of

$$\lim_{\omega_{\ell} \rightarrow 0} \eta_{D_{5/2}} = 0.73, \quad (5.119)$$

implying that in the limit of very low laser frequency (i.e., nearly static fields), a substantial fraction of scattering events would populate the $^2D_{5/2}$ manifold. Such behavior is unphysical,

as it would require the emission of photons with insufficient energy to bridge the level spacing, effectively violating energy conservation.

In contrast, the ω^3 -theory model correctly incorporates the frequency dependence of the electromagnetic density of states. In this framework, scattering into the 2D_J manifolds is smoothly suppressed as the laser frequency decreases, and becomes energetically forbidden once the laser frequency falls below the transition frequency required to reach the 2D_J states. This behavior is physically consistent, as it reflects the requirement that the emitted photon carry away the appropriate energy difference.

Experimentally, we measure a branching fraction of $\eta_{D_{5/2}} = 0.19(1)$, obtained from repeated measurements as described above. This value is in good agreement with the ω^3 -theory prediction, differing by only $0.7\sigma_E$, where σ_E denotes the standard error of the measurement. By contrast, the CDA model prediction is strongly inconsistent with the data, differing by $8.7\sigma_E$. This clear disagreement provides strong evidence that inclusion of the ω_{sc}^3 density-of-states factor is essential for accurately describing spontaneous Raman scattering in this regime.

For applications in quantum information processing, it is crucial to understand the fundamental limits on gate fidelity imposed by spontaneous Raman scattering (SRS) in the far-detuned regime. Although far-detuned laser fields are used to suppress scattering, residual SRS events remain an intrinsic source of error during gate operations and ultimately limit achievable fidelities.

For a stimulated Raman transition, the probability that an ion undergoes at least one SRS event during a gate of duration τ_ℓ can be modeled as a Poisson process. If multiple laser beams indexed by j contribute to the scattering, this probability is given by

$$\epsilon \approx 1 - e^{-\sum_j \tau_\ell \Gamma_j}, \quad (5.120)$$

where Γ_j is the SRS rate associated with beam j . This expression reflects the cumulative probability of scattering from all beams during the gate time.




Figure 5.9: Calculated two-qubit gate error versus gate-laser optical frequency. (Black, solid) Moore *et al.* theory using a three-beam configuration, $K = 1$ and motional sideband frequency of $2\pi \times 5$ MHz. (Red, dotted) CDA theory for three beams, which reproduces limit $\varepsilon_{D\infty} = (1/2)1.46 \cdot 10^{-4}$ for this configuration. (Blue, dashed) ω^3 -theory prediction. (Black point) Inferred achievable two-qubit gate error, calculated from the measured scattering rate and Stark shifts at 532 nm.

For a single-qubit gate implemented via a stimulated Raman transition, which we model as a π pulse, the gate duration is set by the resonant Raman Rabi frequency Ω_R :

$$\tau_{1q} = \frac{\pi}{2|\Omega_R|}. \quad (5.121)$$

In this case, two Raman beams contribute to both the coherent drive and the scattering, and the total error scales with both the scattering rate and the gate duration.

For a two-qubit entangling gate, such as a Mølmer–Sørensen gate, the situation is more complex. The gate couples the internal states of the ions to a shared motional mode,

and the duration depends on both the Raman coupling strength and the Lamb–Dicke parameter η , which characterizes the coupling between motion and internal states. For a gate that traces K loops in motional phase space, the gate duration is given by

$$\tau_{2q} = \frac{\pi\sqrt{K}}{2\eta|\Omega'_R|}, \quad (5.122)$$

where Ω'_R is the effective Raman coupling for the sideband transition.

Experimentally, multi-beam configurations are often used to optimize the Raman coupling while minimizing technical constraints. In particular, a three-beam configuration can be employed to generate two Raman couplings, with the most efficient arrangement achieved when two beams propagate in one direction and a third beam, with twice the intensity, propagates in the opposite direction. This configuration increases the effective Rabi frequency Ω'_R by a factor of $\sqrt{2}$ compared to a four-beam configuration with equal intensities.

Using the scattering rates Γ_j , the resulting gate errors for single- and two-qubit operations can be approximated as

$$\begin{aligned} \epsilon_{1q} &\approx 2\tau_{1q}\Gamma = \frac{\pi\Gamma}{|\Omega_R|}, \\ \epsilon_{2q} &\approx 3\tau_{2q}\tilde{\Gamma} = \frac{3\pi\sqrt{K}\tilde{\Gamma}}{\eta|\Omega'_R|}. \end{aligned} \quad (5.123)$$

Here, Γ is the scattering rate per beam for the single-qubit case, while $\tilde{\Gamma}$ represents the scattering rate corresponding to the average beam intensity in the multi-beam configuration used for the two-qubit gate. The prefactor of 3 in ϵ_{2q} reflects the contributions from all beams, and the factor of 2 in ϵ_{1q} accounts for the two Raman beams used in the single-qubit gate. Additionally, for two-qubit gates, scattering from either ion results in gate failure, further increasing the effective error probability.

A key connection between scattering and gate performance arises from the fact that the stimulated Raman Rabi frequency Ω_R is proportional to the differential AC Stark

shift induced by the laser. As a result, the experimentally measured relationship between SRS rate and Stark shift (Figure 5.7) can be used to directly infer the achievable gate error for a given laser wavelength. Using this relationship, we determine an empirical estimate of the two-qubit gate error at 532 nm. The slope of the orthogonal distance regression (ODR) fit to the total SRS data in Figure 5.7 provides the ratio Γ/δ , where δ is the differential Stark shift. The ratio δ/Ω_R is calculated using Eq. 5.116. Combining these gives $\epsilon_{2q} = (\Gamma/\Omega_R)(3\pi\sqrt{K}/\eta)$, as in Eq. 5.123.

Figure 5.9 shows the calculated two-qubit gate error as a function of laser frequency for a representative gate with $K = 1$ and a motional mode frequency of $2\pi \times 5$ MHz. From the experimentally measured scattering rates and Stark shifts at 532 nm, we infer a two-qubit gate error of $6.4(2) \times 10^{-5}$, which is in excellent agreement with the prediction of the Moore *et al.* model, 6.6×10^{-5} .

This result stands in contrast to earlier predictions based on the CDA model. In particular, Table III of Ref. (Ozeri et al., 2007) suggests a lower bound on the achievable error of $\epsilon_{D\infty} = (1/2) \times 1.46 \times 10^{-4} = 7.3 \times 10^{-5}$ (for a three-beam configuration), implying that errors below this value are fundamentally unattainable when detuned outside the $^2P_J^o$ levels. Furthermore, the CDA model predicts an error of 1.05×10^{-4} at 532 nm, both of which are inconsistent with our experimentally inferred value.

Importantly, both the updated theoretical treatments and the experimentally inferred error rates fall below the 10^{-4} threshold, which is commonly associated with fault-tolerant quantum error correction. The absence of the asymptotic error floor predicted by the CDA model indicates that further reductions in SRS-induced gate error are achievable. In particular, moving to gate laser wavelengths that are more strongly red-detuned from the $^2P_J^o$ manifold provides a clear path toward reducing scattering and improving gate fidelity.

In summary, we have experimentally characterized the spontaneous Raman scattering (SRS) rate in $^{133}\text{Ba}^+$ under illumination with far-detuned laser light at 532 nm and compared the results with several theoretical models. We find that the measured scattering rate is significantly lower than earlier predictions based on the CDA model, and is instead

in quantitative agreement with more recent and complete treatments of SRS. The primary distinction between these models arises from the inclusion of the photonic density of states, which introduces a ω_{sc}^3 dependence on the emitted photon frequency and plays a crucial role in determining both the total scattering rate and the branching fractions.

By separately measuring the branching fraction into the $^2\text{D}_{5/2}$ manifold, we further demonstrate that models neglecting this frequency dependence lead to qualitatively incorrect predictions, including unphysical behavior in the low-frequency limit. In contrast, models incorporating the correct density-of-states scaling accurately capture both the magnitude of the scattering rate and the distribution of scattering among final states. These results provide strong experimental validation of the underlying atomic physics and clarify the mechanisms governing off-resonant scattering in this system.

Importantly, these findings have direct implications for quantum information processing with trapped ions. Using the measured scattering rates together with the experimentally calibrated Stark shifts, we infer achievable gate errors for both single- and two-qubit operations. At 532 nm, a commonly available and technologically convenient wavelength, we determine a two-qubit gate error of $6.4(2) \times 10^{-5}$ arising from spontaneous Raman scattering. This value is consistent with the updated theoretical models and lies below the $\sim 10^{-4}$ threshold typically associated with fault-tolerant quantum error correction.

More broadly, our results revise the previously understood limits on gate fidelity imposed by atomic physics. In contrast to earlier predictions, which suggested the existence of a fundamental lower bound on SRS-induced gate error, we find no such asymptotic limit. Instead, when the correct frequency dependence of the scattering process is taken into account, the error probability decreases continuously with increasing detuning, such that $\epsilon_{2q} \rightarrow 0$ as $\omega_\ell \rightarrow 0$. This implies that spontaneous Raman scattering does not impose a fundamental barrier to achieving arbitrarily low gate errors in trapped-ion systems, provided sufficient laser power and detuning are available.

These results not only resolve discrepancies between theory and experiment, but also identify clear pathways for further reducing scattering-induced errors. In particular,

operations at longer wavelengths, further red-detuned from the $^2P_J^o$ manifold, offer a straightforward route to improved performance. Together, these findings strengthen the case for $^{133}\text{Ba}^+$ as a promising platform for high-fidelity quantum information processing and provide a more accurate framework for evaluating and optimizing gate performance in trapped-ion systems.

CHAPTER 6

Magic Angle Raman Gates in Magnetic Clock Qubits

“Although I am primarily an experimental physicist, theoretical physics is my hobby.” – Norman Ramsey

6.1 The AC Light Shift of a Clock Qubit

While measuring the scattering rate from a far detuned laser, we first needed an accurate calibration of the laser intensity at the ion. This is commonly done by measuring the AC Stark shift of a well characterized transition. In our experiment, we were able to measure this shift in two independent ways. We first used the optical transition between the $S_{1/2}$ and $D_{5/2}$ states, where the Stark shift was relatively large and on the order of hundreds of kHz. We then repeated the measurement using the clock state qubit in $^{133}\text{Ba}^+$ as an independent check. Although the expected shift on the clock transition was much smaller, on the order of tens of Hz, the long coherence time of the clock state and the low phase noise of the microwave drive made this measurement feasible.

This resulted in a mismatch between the two intensity measurements, the ratio of which seemed easy to calculate theoretically, or so we believed. Thus, we kept measuring and kept checking our calculation. It is often the case in experimental physics that the theoretical calculations one performs do not immediately agree with the real world. When this happens, an experimentalist immediately begins to worry about, and to try to quantify, every assumption and experimental variable. We questioned whether the electric field

of the trap, the astigmatism of the laser beam, or some other imperfection in the setup might explain the discrepancy. Finally, we noticed a truly unexpected trend. We placed a quarter-wave plate in the laser path and saw very clearly that the differential clock Stark shift varied with the helicity of the beam. At this point, our bosses severely doubted our experimental methods, as all PIs do when faced with a surprising measurement. Thus, we tested the effect extensively and, without a doubt, mapped out a clear and reproducible trend that our theory did not account for at all.

The original expectation can be understood from the usual first-order treatment of the AC Stark shift in a fixed hyperfine manifold. For a far detuned laser, the light shift of a hyperfine state $|F, m_F\rangle$ is commonly decomposed into scalar, vector, and tensor contributions. In this approximation, the vector part of the light shift can be written as

$$H_v^{(F)} = \chi_v^{(F)} [i (\boldsymbol{\epsilon}^* \times \boldsymbol{\epsilon})] \cdot \mathbf{F}, \quad (6.1)$$

where $\boldsymbol{\epsilon}$ is the complex polarization vector of the light, $\mathbf{F}$ is the total angular momentum operator, and $\chi_v^{(F)}$ contains the laser intensity, detuning, and vector polarizability. The quantity

$$i (\boldsymbol{\epsilon}^* \times \boldsymbol{\epsilon}) = \xi \hat{z}, \quad (6.2)$$

where $\xi = +1$ for σ^+ light and $\xi = -1$ for σ^- light, up to the convention used to define the propagation direction. Taking the expectation value in the state $|F, m_F\rangle$ gives

$$\Delta E_v(F, m_F) = \langle F, m_F | H_v^{(F)} | F, m_F \rangle = \chi_v^{(F)} \xi m_F. \quad (6.3)$$

Since both clock states have $m_F = 0$, this theory predicts no helicity-dependent differential Stark shift of the clock transition.

This formalism can sometimes be rather confusing and seeing as I am an experimentalist "laser-jock", I find it simpler to explain this by calculating the light shift directly from

second order perturbation theory. The AC Stark shift from a far detuned laser is

$$\Delta E_g^{(q)} = \frac{E_0^2}{4} \sum_e \frac{|\langle e | d_q | g \rangle|^2}{E_g + \hbar\omega_L - E_e}, \quad (6.4)$$

where E_0 is the electric field amplitude of the laser, ω_L is the laser frequency, and q labels the polarization of the light. Here $q = +1$ corresponds to σ^+ light, while $q = -1$ corresponds to σ^- light. The sum is over the excited states coupled to the ground state by the electric dipole operator. For the $6S_{1/2}$ ground state these are primarily excited P states. For the clock qubit in $^{133}\text{Ba}^+$, the two states are

$$|0\rangle = |6S_{1/2}, F = 0, m_F = 0\rangle \quad |1\rangle = |6S_{1/2}, F = 1, m_F = 0\rangle. \quad (6.5)$$

The important point is that both states have $m_F = 0$. For pure σ^+ light, the dipole operator only couples the $m_F = 0$ state to the $m_F = 1$, or for σ^- to the $m_F = -1$. The matrix elements are the same up to the Clebsch-Gordan coefficients. Thus the two coefficients can differ by a phase, but their magnitudes are the same:

$$\left| C_{F,0;1,+1}^{F',+1} \right|^2 = \left| C_{F,0;1,-1}^{F',-1} \right|^2. \quad (6.6)$$

The key point here is that the light shift perturbation squares these coefficients so their respective signs don't contribute. Thus, the light shift for the $|0\rangle$ and $|1\rangle$ will be the same with σ^+ or σ^- polarization.

The limitation of the treatment above is that it applies perturbation theory to the light shift after the clock-state basis has already been chosen by diagonalizing the hyperfine and Zeeman Hamiltonian. In this picture, the AC Stark shift is treated as a small correction to states that have already been dressed by the magnetic field. However, the Zeeman interaction and the helicity-dependent AC Stark shift affect the same two-state subspace and can change the clock splitting through the same mixing mechanism. They should therefore not be treated as two independent perturbations applied one after the other.

Instead, the hyperfine interaction, the Zeeman interaction, and the AC Stark shift must be included in the Hamiltonian at the same time and diagonalized together. This is most transparent when the problem is written in the uncoupled basis.

For $^{133}\text{Ba}^+$ in the $6S_{1/2}$ ground state, both the electronic and nuclear angular momenta are spin one half, The two $m_F = 0$ states are

$$|F = 1, m_F = 0\rangle = \frac{1}{\sqrt{2}}(|+\rangle + |-\rangle), \quad |F = 0, m_F = 0\rangle = \frac{1}{\sqrt{2}}(|+\rangle - |-\rangle). \quad (6.7)$$

where

$$|+\rangle = |m_J = +1/2, m_I = -1/2\rangle, \quad |-\rangle = |m_J = -1/2, m_I = +1/2\rangle, \quad (6.8)$$

In this basis, the hyperfine Hamiltonian is

$$H_{\text{hfs}} = hA \mathbf{I} \cdot \mathbf{J}. \quad (6.9)$$

Using

$$\mathbf{I} \cdot \mathbf{J} = I_z J_z + \frac{1}{2} (I_+ J_- + I_- J_+), \quad (6.10)$$

the Hamiltonian in the basis $\{|+\rangle, |-\rangle\}$ is

$$\frac{H_{\text{hfs}}}{h} = \begin{pmatrix} -A/4 & A/2 \\ A/2 & -A/4 \end{pmatrix}. \quad (6.11)$$

The static magnetic field gives a diagonal shift in this same basis. For a magnetic field along the quantization axis,

$$\frac{H_B}{h} = \frac{\mu_B}{h} g_J B J_z - \frac{\mu_N}{h} g_I B I_z. \quad (6.12)$$

For the two states $|+\rangle$ and $|-\rangle$, this becomes

$$\frac{H_B}{h} = \frac{1}{2}\beta B \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \beta = \frac{\mu_B}{h}g_J + \frac{\mu_N}{h}g_I. \quad (6.13)$$

Now consider the AC Stark shift from pure circular light. In this basis, the light shift is diagonal because the electric dipole coupling depends on the electronic angular momentum projection m_J , not directly on m_F . For σ^+ light, let the Stark shifts of the two basis states be $\Delta_+^{(\sigma^+)}$ and $\Delta_-^{(\sigma^+)}$, where these quantities are written as frequencies, so that the energy shifts are $h\Delta_\pm$. These shifts come from ordinary second order perturbation theory,

$$\Delta_\pm^{(\sigma^+)} = \frac{E_0^2}{4h} \sum_e \frac{|\langle e | d_{+1} | m_J = \pm 1/2 \rangle|^2}{E_g + \hbar\omega_L - E_e}. \quad (6.14)$$

For σ^- light, the corresponding expression uses d_{-1} instead of d_{+1} . By symmetry, changing from σ^+ to σ^- swaps the two m_J shifts,

$$\Delta_+^{(\sigma^-)} = \Delta_-^{(\sigma^+)}, \quad \Delta_-^{(\sigma^-)} = \Delta_+^{(\sigma^+)}. \quad (6.15)$$

It is useful to define the average Stark shift

$$\overline{\Delta} = \frac{1}{2} \left(\Delta_+^{(\sigma^+)} + \Delta_-^{(\sigma^+)} \right), \quad (6.16)$$

and the differential Stark shift between the two uncoupled basis states as

$$\Delta_v = \Delta_+^{(\sigma^+)} - \Delta_-^{(\sigma^+)}. \quad (6.17)$$

Then the AC Stark Hamiltonian can be written as

$$\frac{H_{ac}^{(\eta)}}{h} = \begin{pmatrix} \overline{\Delta} + \eta\Delta_v/2 & 0 \\ 0 & \overline{\Delta} - \eta\Delta_v/2 \end{pmatrix}, \quad (6.18)$$

where $\eta = +1$ for σ^+ light and $\eta = -1$ for σ^- light. Combining the hyperfine interaction,

the static magnetic field, and the AC Stark shift gives

$$\frac{H^{(\eta)}}{h} = \begin{pmatrix} -A/4 + \bar{\Delta} + (\beta B + \eta \Delta_v)/2 & A/2 \\ A/2 & -A/4 + \bar{\Delta} - (\beta B + \eta \Delta_v)/2 \end{pmatrix}. \quad (6.19)$$

Equivalently,

$$\frac{H^{(\eta)}}{h} = \left(-\frac{A}{4} + \bar{\Delta} \right) \mathbb{1} + \frac{A}{2} \sigma_x + \frac{1}{2} (\beta B + \eta \Delta_v) \sigma_z. \quad (6.20)$$

The eigenvalues are

$$\frac{E_{\pm}^{(\eta)}}{h} = -\frac{A}{4} + \bar{\Delta} \pm \frac{1}{2} \sqrt{A^2 + (\beta B + \eta \Delta_v)^2}. \quad (6.21)$$

Therefore the clock splitting is

$$\nu_{\text{clock}}^{(\eta)} = \sqrt{A^2 + (\beta B + \eta \Delta_v)^2}. \quad (6.22)$$

Without the Stark beam, the clock splitting is

$$\nu_{\text{clock}}^{(0)} = \sqrt{A^2 + (\beta B)^2}. \quad (6.23)$$

Thus, the Stark-induced change in the clock frequency is

$$\delta \nu_{\text{clock}}^{(\eta)} = \sqrt{A^2 + (\beta B + \eta \Delta_v)^2} - \sqrt{A^2 + (\beta B)^2}. \quad (6.24)$$

We can expand this expression for small Δ_v , expanding to second order gives

$$\delta \nu_{\text{clock}}^{(\eta)} \simeq \eta \frac{\beta B}{\sqrt{A^2 + (\beta B)^2}} \Delta_v + \frac{A^2}{2 [A^2 + (\beta B)^2]^{3/2}} \Delta_v^2. \quad (6.25)$$

This expression finally explains the strange effect we saw in the lab. There is a linear dependence on the Stark shift, that is equal and opposite for $\eta = \pm 1$. When the laser is off, $\Delta_v = 0$ and the equation becomes the familiar Breit-Rabi formula. At $B = 0$, we also

see the naive second order perturbation theory result for the differential Stark shift, that retains no helicity dependence. The Stark shift adds to the Zeeman shift term inside the square and for this reason it is sometimes called a fictitious magnetic field.

Again it can be useful to illustrate this through another avenue. The helicity-dependent clock shift can be viewed as a third-order perturbation involving one Zeeman interaction and two interactions with the optical field. Schematically, for a clock state $|a\rangle$, the relevant contribution has the form

$$\Delta E_a^{(3)} = \sum_{b,e} \frac{\langle a| H_Z |b\rangle \langle b| V_{\text{opt}} |e\rangle \langle e| V_{\text{opt}} |a\rangle}{(E_a - E_b)(E_a + \hbar\omega_L - E_e)} + \text{perm.} \quad (6.26)$$

Here $|b\rangle$ is the other hyperfine clock state in the ground-state manifold, while $|e\rangle$ is an excited state coupled by the optical field. The term “perm” represents the other possible orderings of the same three interactions. For circularly polarized light, the optical interaction is

$$V_{\text{opt}} = -\mathbf{d} \cdot \mathbf{E} \propto d_q, \quad (6.27)$$

where $q = +1$ for σ^+ light and $q = -1$ for σ^- light. The corresponding contribution can therefore be written more explicitly as

$$\Delta E_a^{(3,q)} = \sum_{b,e} \frac{\langle a| H_Z |b\rangle \langle b| d_q^\dagger |e\rangle \langle e| d_q |a\rangle}{(E_a - E_b)(E_a + \hbar\omega_L - E_e)} + \text{perm.} \quad (6.28)$$

The differential shift of the clock transition is then

$$h\Delta\nu_{\text{clock}}^{(3,q)} = \Delta E_1^{(3,q)} - \Delta E_0^{(3,q)}. \quad (6.29)$$

This expression reveals the asymmetry in the calculation, one of the 4 Clebsch-Gordon coefficients relevant for the calculation is negative and this produces an F dependent Stark shift.

6.2 The *omg* Protocol

Once the above effect was understood and confirmed to be real, it had immediate implications for how we thought about illuminating an ion with far-detuned laser light. In particular, it became clear that, for a particular mixture of helicities, the differential Stark shift could be engineered to vanish. This offers a practical advantage because it reduces the extent to which fluctuations in laser power are mapped onto qubit detuning errors. This is especially useful in the context of the recently proposed *omg* protocol for trapped-ion quantum computing ([Allcock et al., 2021](#)).

The *omg* protocol was introduced as a way to obtain some of the advantages of using multiple ion species while still working with only a single atomic species. In many trapped ion architectures, different species are used for different tasks. One species may store quantum information, while another is used for sympathetic cooling, state preparation, detection, or auxiliary operations. This separation can be extremely useful, but it also makes the experiment substantially more complicated. Each ion species requires its own photoionization, cooling, repumping, and detection lasers. The different masses of the ions also change the normal mode structure of the crystal, which can complicate gates, transport, and cooling. In a large processor, these complications become increasingly costly.

The basic idea of the *omg* protocol is that different long lived manifolds of the same ion can play the role of different qubit types. The name *omg* refers to optical, metastable, and ground state qubits. A ground state qubit can provide a stable memory, an optical qubit can be useful for fast coherent operations, and a metastable qubit can provide an auxiliary storage space that is spectroscopically separated from the ground state manifold. In this sense, a metastable state can act like a programmable second species, but without changing the ion species in the trap.

This is useful because the role of an ion no longer has to be fixed by its atomic identity. Instead, quantum information can be moved between different internal manifolds

depending on the operation being performed. Information can be stored in a ground state qubit when long coherence is needed, transferred to an optical qubit when fast control is useful, or shelved into a metastable manifold when it should be protected from light resonant with the ground state. This provides a route toward some of the flexibility of mixed species ion chains while avoiding many of the technical costs associated with using two different atomic species.

From this perspective, the ability to engineer differential Stark shifts with polarization is particularly important. Laser beams used for gates, cooling, or other operations can still shift spectator qubits even when they are far detuned. If the polarization can be chosen so that the relevant differential Stark shift vanishes, then the qubit becomes less sensitive to fluctuations in the laser intensity. This provides an additional control knob for making different qubit encodings compatible within the same ion. More broadly, it shows that far detuned light does not simply produce an unavoidable shift, but can be engineered so that its effect on a chosen qubit is strongly suppressed.

6.3 Experimental Measurements of Magic Angle Gates

The *omg* protocol assumes that one subspace of the atomic Hilbert space can be accessed without causing unacceptable errors in another subspace. In this work, we studied this assumption in $^{133}\text{Ba}^+$. We found that high-intensity laser operations necessary for gates in the metastable qubit space can indeed disturb information stored in the ground-state qubit space. The mechanism is a differential AC Stark shift of the ground-state clock qubit caused by the metastable-qubit gate lasers. Intensity noise in the gate beams is then mapped into qubit-frequency noise, which causes decoherence. We also found that this effect can be mitigated by using a magnetic-field-induced vector light shift to cancel the scalar and tensor contributions to the differential Stark shift. With the correct laser polarization, which we call the magic polarization, the differential light shift of the ground-state qubit is nulled. This protects information stored in the ground-state qubit

while high-intensity laser operations are performed in the metastable manifold.

The implementation considered here uses the $\{g, m, g\}$ version of the *omg* architecture. The ordered triplet denotes the subspace used for cooling, gates, and storage (Allcock et al., 2021). Cooling and readout are performed using the ground-state, or g -type, qubit in the $^2S_{1/2}$ manifold. Gate operations are performed by first activating the ion into the metastable, or m -type, qubit in the $^2D_{5/2}$ manifold. After the gate operation, the ion can be returned to the g manifold for storage or readout.

The advantage of this architecture is that the high-power laser beams used for gates can be applied globally if individual ions can be selectively activated into the m manifold. This would greatly simplify the optical control system. Instead of tightly focusing high-power gate beams onto individual ions, the same beams could illuminate the full chain, while only the ions activated into the m manifold participate in the gate. The other ions, which remain in the g manifold, should ideally be protected from the gate operation.

The g -type qubit is defined using the $F = 0$ and $F = 1$ hyperfine clock states of the $^2S_{1/2}$ manifold. Both states are zero-field clock states with $m_F = 0$. The m -type qubit is defined using the $F = 2$ and $F = 3$ hyperfine clock states of the $^2D_{5/2}$ manifold. Laser cooling, state preparation, and state readout are performed through the g manifold (Christensen, 2020). After high-fidelity preparation of the $F = 1, m_F = 0$ state of the g qubit, population is transferred to the $F = 3, m_F = 0$ state of the $^2D_{5/2}$ manifold using a coherent operation on the 1762 nm optical qubit transition.

Readout of the m qubit is performed using a shelving operation. Light at 1762 nm is applied together with light at 493 nm to transfer population from the target m state to the $^2D_{3/2}$ manifold, where it can be detected by resonant fluorescence. This allows the state of the metastable qubit to be mapped onto a bright or dark fluorescence outcome.

The stimulated Raman beams used for m -qubit operations are generated from a continuous-wave 532 nm laser. The laser is passed through acousto-optical modulators to create two beams whose frequency difference is near the m -qubit splitting. An electro-optical modulator in the same beam path allows the same laser system to also drive




Figure 6.1: Implementation of the $\{g, m, g\}$ omg protocol in $^{133}\text{Ba}^+$. Cooling and readout are performed using g qubits in the $^2S_{1/2}$ manifold. Gates are performed using global beams that address the m qubit. The level diagram shows both the g and m manifolds in $^{133}\text{Ba}^+$. Raman spectroscopy of the m -qubit hyperfine splitting gives a splitting of 89.6697(4) MHz at 5.036(1) G.

operations on the g qubit. Although this is not required for the $\{g, m, g\}$ architecture, the ability to use the same optical path for both g - and m -qubit operations adds useful flexibility to the experiment.

Before performing gates in the m manifold, it was necessary to measure the hyperfine splitting of the $5d6s^2D_{5/2}$ state. This splitting was not known with sufficient precision for high-resolution coherent control. We measured it using stimulated Raman spectroscopy on the $F = 3, m_F = 0 \leftrightarrow F = 2, m_F = 0$ transition at a magnetic field of $B = 5.036(1)$ G, as shown in Figure 6.1. From this measurement, we obtained the first high-resolution

measurement of the $^{133}\text{Ba}^+ 5d6s^2D_{5/2}$ hyperfine constant,

$$A_{5/2} = 29.75650(10)_{\text{sys}}(4)_{\text{stat}} \text{ MHz}. \quad (6.30)$$

This value provided the frequency reference needed to drive coherent operations in the m manifold. Using a single m -qubit rotation, we measured an intensity-noise-limited m -qubit coherence time of 2.6(9) ms.

The experimental sequence used to prepare, manipulate, and read out the m qubit is shown in Figure 6.2. The narrow electric quadrupole transition at 1762 nm transfers population between the g and m manifolds. The sequence begins with Doppler cooling for 8 ms using 493 nm and 650 nm light. Both cooling beams use electro-optical modulators to generate the hyperfine tones needed to address the relevant levels in the $S_{1/2}$ and $D_{3/2}$ manifolds.

After cooling, the 493 nm hyperfine tones are turned off and a tone near 1840 MHz, resonant with the $|S_{1/2}, F = 1, m_F = 0\rangle$ to $|P_{1/2}, F = 1\rangle$ transition, optically pumps the ion into the $|S_{1/2}, F = 0, m_F = 0\rangle$ state in approximately 150 μs . A 130 μs microwave π pulse then transfers the population to the $|S_{1/2}, F = 1, m_F = 0\rangle$ state. A 45 μs pulse on the 1762 nm transition transfers approximately 90 to 95 percent of the population into the $|D_{5/2}, F = 3, m_F = 0\rangle$ state.

The transfer to the metastable state is heralded by turning on the cooling lasers for 8 ms and collecting photon counts. If the ion has been successfully transferred into the $D_{5/2}$ manifold, it remains dark during this detection period. The two co-propagating 532 nm Raman beams are then applied to drive coherent transfer between the $|D_{5/2}, F = 3\rangle$ and $|D_{5/2}, F = 2\rangle$ hyperfine clock states. The frequency of one beam is varied with an acousto-optical modulator to scan across the Raman resonance.

To read out the m qubit, the $|D_{5/2}, F = 3\rangle$ state is selectively shelved by applying 1762 nm and 493 nm light together for 7 ms. This transfers the target population to the $D_{3/2}$ manifold with a measured fidelity of 0.9989(2). The cooling lasers are then turned




Figure 6.2: Experimental sequence used for metastable-qubit operations. The ion is cooled and optically pumped in the ground-state manifold, transferred to the $D_{5/2}$ manifold using the 1762 nm transition, manipulated with 532 nm Raman beams, and then read out through a shelving operation.

on for 10 ms to measure fluorescence. Population originally in the $|D_{5/2}, F = 3\rangle$ state is mapped to bright fluorescence, while population that remains in the $|D_{5/2}, F = 2\rangle$ state is dark. Finally, 614 nm light resonant with the $D_{5/2} \rightarrow P_{3/2}$ manifold returns any remaining population in the metastable manifold to the ground-state manifold so that the next experimental cycle can begin.

A Raman spectroscopy scan of the metastable hyperfine splitting is shown in Figure 6.3. The two Raman beams are derived from a single laser using two acousto-optical modulators. For the scan shown here, the Raman illumination time was $268 \mu\text{s}$ and the power in each beam was approximately 35 mW. The beam diameter was chosen to be large, roughly $100 \mu\text{m}$, to reduce sensitivity to beam pointing fluctuations. The measured splitting was $89.6697(4)$ MHz at $5.036(1)$ G. After accounting for the Zeeman shifts of the $D_{5/2}$ levels, the zero-field hyperfine splitting was approximately $89.2695(4)$ MHz.

The measured m -qubit splitting depends on the intensity of the Raman beams because the same beams produce a differential light shift of the $D_{5/2}$ clock states. Figure 6.4 shows the measured frequency splitting as a function of the m -qubit Rabi rate, which provides a proxy for the laser intensity at the ion. The intercept of the linear fit gives the splitting at zero differential light shift. The data point near a 4 kHz Rabi rate corresponds to the intensity used in the spectroscopy scans. The intercept is approximately 400 Hz lower than the splitting measured at that intensity, and this correction was subtracted when extracting the zero-field hyperfine coefficient.




Figure 6.3: Raman spectroscopy measurement of the $D_{5/2}$ hyperfine splitting using high-power 532 nm light to drive a stimulated Raman transition.

The metastable splitting was measured multiple times, with a corresponding magnetic-field measurement for each data set. The magnetic field was determined using a detuned microwave Ramsey sequence on the g qubit in $^{133}\text{Ba}^+$. The standard error of these magnetic-field measurements was approximately 4×10^{-4} G.

A central question for the $\{g, m, g\}$ architecture is whether high-intensity laser operations in the m manifold disturb information stored in the g manifold. To test this, we performed a detuned microwave Ramsey sequence on the g qubit while illuminating the ion with the 532 nm beams used for m -qubit gates. The Ramsey oscillation frequency measures the detuning between the applied microwave field and the actual g -qubit splitting. A differential AC Stark shift from the 532 nm light therefore appears as a shift of the Ramsey frequency. Intensity noise in the laser causes this frequency to fluctuate, which produces decoherence of the Ramsey fringe.

The Ramsey sequence begins by preparing the $F = 0$ state of the g qubit. A detuned microwave $\pi/2$ pulse creates a superposition of the two clock states. The system then evolves for a variable wait time, during which the 532 nm beams may be applied. A second




Figure 6.4: Measured m -qubit frequency splitting as a function of Rabi rate. The measured splitting shifts with Raman-beam intensity due to the differential AC Stark shift of the $D_{5/2}$ clock states. The intercept of the fit gives the splitting extrapolated to zero differential light shift.

detuned microwave $\pi/2$ pulse maps the accumulated phase onto the population, and the final state is measured. The sequence is repeated 300 times for each wait time. The decay of the Ramsey fringe amplitude gives the coherence time of the g qubit under the chosen illumination condition.

Figure 6.5 shows three Ramsey measurements at a magnetic field of 5.00(1) G. With no 532 nm illumination, the Ramsey fringe shows no significant decay over the 30 ms measurement window. The fitted Ramsey frequency corresponds to a microwave detuning of 440.4(1) Hz. When the ion is illuminated by two co-propagating 532 nm Raman beams at a combined intensity of approximately 150 MW/m² and with the polarization set to $A = -1$, the Ramsey contrast decays much more quickly. The fitted coherence time is 46.80(3.8) ms, or approximately 50 ms. This decoherence is consistent with fluctuations in the differential Stark shift of the g -qubit transition caused by intensity noise in the Raman beams. The illumination produces fluctuations on more than one time scale. Fast intensity noise, at frequencies comparable to or above the inverse Ramsey interrogation




Figure 6.5: Detuned microwave Ramsey measurements of the g qubit under different illumination conditions. The Ramsey oscillation frequency gives the detuning of the microwave drive from the shifted g -qubit transition. Measurements are shown with no 532 nm illumination, with 532 nm illumination at a non-magic polarization, and with 532 nm illumination at the magic polarization.

time, directly dephases the Ramsey fringe. Slower drifts, likely due to thermal changes in the optical system, shift the Ramsey frequency over the course of the measurement. In the non-magic data, slow frequency drifts of roughly 20 Hz over 60 minutes were observed. This corresponds to a change in laser intensity at the ion of roughly 10 MW/m^2 during a typical Ramsey scan. The measured coherence decay corresponds to a g -qubit crosstalk error of approximately $2.7(2) \times 10^{-5}$ per m -qubit π pulse. Extrapolating from this error gives an estimated crosstalk error of order 10^{-4} for a two-qubit gate in the m space. Although this level is below typical fault-tolerance thresholds, it is still desirable to reduce the error because doing so lowers the overhead required for error correction.

The crosstalk described above could be reduced with technical improvements to the intensity stability of the 532 nm beams, or by using dynamical decoupling. A more direct solution is to eliminate the differential light shift of the g -qubit transition itself. To see how this can occur, we write the light shift of a hyperfine state due to a laser of frequency

ω_L , polarization $\hat{\epsilon}_L$, and intensity I_L as

$$\Delta E_{F,m_F} \propto \alpha_F(m_F, \omega_L, \hat{\epsilon}_L) I_L, \quad (6.31)$$

where α_F is the dynamic polarizability. The dynamic polarizability can be decomposed into scalar, vector, and tensor parts,

$$\begin{aligned} \alpha_F(m_F, \omega_L, \hat{\epsilon}_L) &= \alpha_F^s(\omega_L) + A \frac{m_F}{2F} \alpha_F^a(\omega_L) \\ &+ \frac{1}{2} (3|\hat{\epsilon}_L \cdot \hat{z}|^2 - 1) \frac{3m_F^2 - F(F+1)}{F(2F-1)} \alpha_F^t(\omega_L), \end{aligned} \quad (6.32)$$

where α_F^s , α_F^a , and α_F^t are the scalar, vector, and tensor dynamic polarizabilities. The helicity projection is

$$A \equiv -i(\hat{\epsilon}_L^* \times \hat{\epsilon}_L) \cdot \hat{z}, \quad (6.33)$$

and $\hat{z}$ is the quantization axis.

At first glance, Eq. (6.32) suggests that the vector light shift should not help for the g -qubit clock states, since both states have $m_F = 0$. In this simple picture, the vector contribution is proportional to m_F and therefore vanishes for both clock states. The tensor contribution remains, but in this system its magnitude is not large enough to match the light shifts of the two qubit states. This conclusion changes in a nonzero magnetic field. The magnetic field breaks rotational invariance, so F is no longer an exact quantum number. Equation (6.32) is then only approximate. The result is a nonzero vector light shift contribution even for nominally $m_F = 0$ clock states (Schleier-Smith et al., 2010; Chicireanu et al., 2011; Kim et al., 2013; Choi and Cho, 2007). This magnetic-field-induced vector light shift can be used to match the dynamic polarizabilities of the two clock states and null the differential light shift.

The effect can be understood intuitively by considering the light shift of the g -qubit states under illumination by σ^+ light. The two states couple primarily to the $F = m_F = 1$ state of the $^2P_{1/2}$ manifold. At zero magnetic field, the $F = 0$ g -qubit state experiences a




Figure 6.6: Measurement of the magic polarization condition. The differential light shift of the g -qubit transition is measured as a function of the helicity projection A for several magnetic fields. The zero crossing gives the magic polarization. The extracted magic polarization varies with magnetic field, and the critical magnetic field B_c is obtained from the condition $A = 1$. The measured value is compared with an atomic-structure calculation, and the calculation is extended to other gate-laser wavelengths.

larger light shift than the $F = 1$ g -qubit state because it is closer in energy to the $^2P_{1/2}$ manifold. Therefore, The g qubit has a negative differential light shift. In a nonzero magnetic field, the two g -qubit states are mixed and their optical couplings are modified. The light shift then depends on the degree of helicity projection. To quantify this effect, we measured the differential light shift of the g qubit using a Ramsey sequence as a function of the helicity projection A and the applied magnetic field $B\hat{z}$. The wave vector of the 532 nm light was aligned nearly antiparallel to the magnetic field, with an angle $\theta_{kz} = 180.0(8)^\circ$. The result is shown in Figure 6.6. The data show a clear magnetic-field-dependent vector light shift. Above a certain magnetic field, this vector contribution is large enough to completely cancel the difference in the scalar light shifts. The polarization at which the differential Stark shift vanishes is the magic polarization.

The magnetic field required for a given magic polarization is found from the zero crossings of the measured differential Stark shift versus A . The minimum magnetic field needed for cancellation occurs for pure σ^+ polarization, corresponding to $A = 1$. We call this field the critical magnetic field, B_c . In our apparatus, direct measurement at the critical field was limited by the need to maintain efficient Doppler cooling at fields above approximately 2 G. At lower fields, coherent population trapping reduced the Doppler cooling efficiency. We therefore extracted B_c by fitting the measured magnetic-field dependence of the magic polarization.

The dependence of the magic polarization on magnetic field is found by diagonalizing the optical Hamiltonian that leads to Eq. (6.32) in the presence of an applied magnetic field. This produces a vector light shift term proportional to B , leading to

$$A \approx -\frac{\hbar\omega_q}{\mu_B B} \sqrt{\frac{2I+1}{2I(2I+2)}} \frac{\alpha_{F'}^s(\omega_L) - \alpha_F^s(\omega_L)}{\alpha_F^a(\omega_L)} \approx \frac{B_c}{B}. \quad (6.34)$$

Here $F' = 0$ and $F = 1$ label the two qubit states, and $\hbar\omega_q$ is the qubit splitting. Fitting this expression to the data gives

$$B_c = 1.15(3) \text{ G} \quad (6.35)$$

at 532 nm for the $^{133}\text{Ba}^+$ g qubit.

Calculating the expected behavior requires a third-order hyperfine-interaction-mediated polarizability treatment. Without this contribution, the scalar polarizabilities of the two clock states cancel. The hyperfine interaction mixes the total electronic angular momentum J into the atomic wave functions, which produces different transition dipole moments from the two hyperfine qubit states to the same fine-structure state. Using empirical data for the low-lying electronic states together with an *ab initio* relativistic RPA+BO many-body calculation, we evaluate the polarizabilities in Eq. (6.34). The calculation predicts

$$B_c = 1.17 \text{ G} \quad (6.36)$$

at 532 nm, in good agreement with the measured value.

The same calculation can be performed as a function of gate-laser wavelength. The predicted critical field is shown in Figure 6.6. At longer wavelengths, the critical field increases because interference from the counter-rotating term reduces the induced vector light shift. This wavelength dependence is important for choosing a practical gate laser. For single-qubit and two-qubit Raman gates, $A = 1$ polarization is preferred because it yields the largest Rabi rate and therefore reduces spontaneous Raman scattering error

(Moore et al., 2023b; Boguslawski et al., 2023). To use the magic condition with pure circular polarization, the applied magnetic field should therefore be near the critical field. At 532 nm, the critical field for $^{133}\text{Ba}^+$ is below the 2 G field required for efficient Doppler cooling in our apparatus. At longer wavelengths the critical field grows. For example, near 1130 nm, $B_c \approx 2$ G and the two-qubit gate infidelity due to Raman scattering is expected to be below 10^{-5} (Boguslawski et al., 2023).

To demonstrate the utility of the magic polarization for *omg* operation, we repeated the Ramsey measurement at $B = 5.00(1)$ G with the polarization set to the measured magic value, $A = 0.212$. Under this condition, the decoherence caused by the differential light shift of the g qubit from the m -qubit gate beams was strongly suppressed. The measured coherence time increased by approximately two orders of magnitude. The fitted coherence time was $\tau = 3.6(2.5)$ s, consistent with the laser-free Ramsey measurement even though the ion was illuminated by a laser field with an intensity of approximately 100 MW/m².

Using the magic polarization, the estimated g -qubit crosstalk error per m -qubit π pulse was reduced from approximately $2.7(2) \times 10^{-5}$ to approximately $3(2) \times 10^{-7}$. A similar 100-fold reduction is expected for two-qubit gates. These infidelities are below state-of-the-art single-qubit gate infidelities (Ballance et al., 2016), so they should not significantly limit the use of the *omg* protocol.

The magic-polarization measurement shows that the differential light shift between zero-field clock states can be nulled using a magnetic-field-induced vector light shift. In the context of the *omg* protocol, this provides a practical way to protect information stored in the g qubit while high-intensity laser operations are performed on the m qubit. More generally, it shows that the polarization of far detuned light can be used as an engineering tool for controlling qubit sensitivity to intensity noise.

This section gives the perturbative description used to understand why a vector light shift can appear for an $m_F = 0$ clock transition in a magnetic field. The energy shift of an

atomic state

$$|i\rangle \equiv |n, F, m_F, (J, I, \ell, S)\rangle, \quad (6.37)$$

in the presence of nonresonant classical electromagnetic radiation can be approximated to second order in the field as (Rosenbusch et al., 2009)

$$\Delta\omega_i = - \sum_k \left(\frac{\langle i | (\vec{d} \cdot \vec{E}_+)^{\dagger} | k \rangle \langle k | \vec{d} \cdot \vec{E}_+ | i \rangle}{4(\omega_{ki} - \omega_L)} + \frac{\langle i | \vec{d} \cdot \vec{E}_+ | k \rangle \langle k | (\vec{d} \cdot \vec{E}_+)^{\dagger} | i \rangle}{4(\omega_{ki} + \omega_L)} \right), \quad (6.38)$$

where $\omega_{ki} = (E_k - E_i)/\hbar$, $\vec{d} = e\vec{r}$ is the electric dipole operator, and the positive-frequency electric field is $\vec{E}_+ = E_+ \hat{e}$.

For $m_F = 0$ states and σ^+ polarization, the rotating term couples only to intermediate states with $m_F = 1$, while the counter-rotating term couples only to intermediate states with $m_F = -1$. Although the intermediate states are different, the numerators are identical at zero magnetic field for $m_F = 0$. The light shift can then be written as

$$\begin{aligned} \Delta\omega_i &= - \sum_k \frac{\langle i | \vec{d} \cdot \vec{E}^* | k \rangle \langle k | \vec{d} \cdot \vec{E} | i \rangle}{4} \left(\frac{1}{\omega_{ki} - \omega_L} + \frac{1}{\omega_{ki} + \omega_L} \right) \\ &= - \sum_k \frac{|\langle k | \vec{d} \cdot \vec{E}_+ | i \rangle|^2}{2} \frac{\omega_{ki}}{\omega_{ki}^2 - \omega_L^2}. \end{aligned} \quad (6.39)$$

Comparing this with the scalar, vector, and tensor decomposition of the light shift gives the usual result that there is no vector light shift for an $m_F = 0$ state at zero magnetic field (Hu et al., 2018).

For Raman gates, it is beneficial for the gate laser to be red detuned from the lowest excited state in order to reduce Raman scattering errors (Boguslawski et al., 2023; Moore et al., 2023b). For nondegenerate ground-state qubits with $\omega_{ki} > 0$ in zero magnetic field, this means that one cannot generally find a zero-field magic wavelength or polarization that nulls the differential light shift of the qubit. In a magnetic field, states with different

F are mixed according to

$$\hat{H}_B = \mu_B(\hat{L}_z + g_e \hat{S}_z)B_z - g_N \mu_N \hat{I}_z B_z. \quad (6.40)$$

The rotating and counter-rotating terms then no longer have identical numerators. Equation (6.38) can be factored as

$$\Delta\omega_i = - \sum_k a_{ki}^{(+)} \frac{\omega_{ki}}{\omega_{ki}^2 - \omega_L^2} - \sum_k a_{ki}^{(-)} \frac{\omega_L}{\omega_{ki}^2 - \omega_L^2}, \quad (6.41)$$

where

$$a_{ki}^{(+)} = |\langle k | \vec{d} \cdot \vec{E}_+ | i \rangle|^2 + |\langle k | (\vec{d} \cdot \vec{E}_+)^{\dagger} | i \rangle|^2, \quad (6.42)$$

$$a_{ki}^{(-)} = |\langle k | \vec{d} \cdot \vec{E}_+ | i \rangle|^2 - |\langle k | (\vec{d} \cdot \vec{E}_+)^{\dagger} | i \rangle|^2. \quad (6.43)$$

The $a_{ki}^{(+)}$ terms have the form of a scalar light shift, while the $a_{ki}^{(-)}$ terms have the form of a vector light shift. The coefficients can have opposite signs for the two clock states, which allows the differential light shift

$$\delta\omega_q(\lambda, \vec{B}, \hat{\epsilon}) \equiv \Delta\omega_0 - \Delta\omega_1 \quad (6.44)$$

to vanish for certain choices of wavelength, magnetic field, and polarization.

6.4 Calculation of the Critical Magnetic Field

A full calculation of the critical field requires hyperfine-interaction-mediated AC polarizabilities, $\alpha_F^{(k),\text{HFI}}(\omega_L)$, evaluated to third order in the light-matter and hyperfine interactions. A useful estimate can be obtained by considering only the magnetic-field-induced mixing of the hyperfine clock states. We consider a $^{133}\text{Ba}^+$ ion exposed to light of frequency ω with σ^+ polarization at zero magnetic field. The two g -qubit states are

denoted as

$$|6s^2S_{1/2}, F=0\rangle \equiv |S, 0\rangle \quad (6.45)$$

and

$$|6s^2S_{1/2}, F=1, m_F=0\rangle \equiv |S, 1, 0\rangle. \quad (6.46)$$

They couple primarily to the excited states $|6s6p^2P_{1/2}, F=1, m_F=1\rangle$, $|6s6p^2P_{3/2}, F=1, m_F=1\rangle$, and $|6s6p^2P_{3/2}, F=2, m_F=1\rangle$. Ignoring the counter-rotating term, the differential light shift is approximately

$$\delta\omega_q(B=0) \approx -\frac{|E|^2}{4} \sum_k \left(\frac{|\langle k|\hat{d}|S, 0\rangle|^2}{4(\omega_k - \omega)} - \frac{|\langle k|\hat{d}|S, 1, 0\rangle|^2}{4(\omega_k - \omega + \omega_q)} \right), \quad (6.47)$$

where k indexes the excited $|P_J, F, m_F=1\rangle$ states, ω_k is the transition frequency from $|S, 0\rangle$ to state k , and ω_q is the ground-state hyperfine qubit splitting. Defining the detunings Δ_1 and Δ_2 from the $P_{1/2}$ and $P_{3/2}$ manifolds gives

$$\delta\omega_q(B=0) \approx -\frac{\omega_q|E|^2}{24} \left(\frac{d_1^2}{\Delta_1^2} + \frac{d_2^2}{\Delta_2^2} \right). \quad (6.48)$$

At zero field, the $|S, 0\rangle$ state is closer to the excited states than the $|S, 1, 0\rangle$ state. The differential light shift is therefore negative and scales linearly with the qubit splitting ω_q .

When a magnetic field is applied along the quantization axis, the two clock states mix as

$$|S, \psi_+\rangle \approx |S, 0\rangle + \frac{\mu_B B}{\omega_q} |S, 1, 0\rangle, \quad (6.49)$$

$$|S, \psi_-\rangle \approx |S, 1, 0\rangle - \frac{\mu_B B}{\omega_q} |S, 0\rangle. \quad (6.50)$$

The magnetic field also mixes excited states, but this contribution is neglected in the estimate because the laser detuning is large compared with the excited-state hyperfine structure. The first contribution to the field-dependent shift comes from the change in the

transition moments to the $|P_{1/2}, F = 1, m_F = 1\rangle$ state,

$$\langle P_{1/2}, 1, 1 | \hat{d} | S, \psi_+ \rangle \approx \frac{d_1}{\sqrt{6}} \left(1 - \frac{\mu_B B}{\omega_q} \right), \quad (6.51)$$

$$\langle P_{1/2}, 1, 1 | \hat{d} | S, \psi_- \rangle \approx -\frac{d_1}{\sqrt{6}} \left(1 + \frac{\mu_B B}{\omega_q} \right). \quad (6.52)$$

This modifies the differential light shift by

$$\delta\omega_q(B) \approx -\frac{\omega_q |E|^2}{24} \left(\frac{d_1^2}{\Delta_1^2} + \frac{d_2^2}{\Delta_2^2} \right) + \frac{d_1^2 |E|^2}{24\Delta_1} \frac{4\mu_B B}{\omega_q}. \quad (6.53)$$

This term has the opposite sign to the zero-field shift. Physically, the $|S, 0\rangle$ state acquires some $|S, 1, 0\rangle$ character, and the opposite sign of that state's transition dipole moment reduces the light shift of the mixed state. The opposite occurs for the other clock state. The second contribution comes from the unequal transition moments to the $P_{3/2}$ states. In a magnetic field,

$$\langle P_{3/2}, 1, 1 | \hat{d} | S, \psi_+ \rangle \approx \frac{d_2}{\sqrt{6}} \left(1 + \frac{\mu_B B}{2\omega_q} \right), \quad (6.54)$$

$$\langle P_{3/2}, 1, 1 | \hat{d} | S, \psi_- \rangle \approx \frac{d_2}{2\sqrt{6}} \left(1 - \frac{2\mu_B B}{\omega_q} \right), \quad (6.55)$$

$$\langle P_{3/2}, 2, 1 | \hat{d} | S, \psi_+ \rangle \approx \frac{d_2}{2\sqrt{2}} \frac{\mu_B B}{\omega_q}, \quad (6.56)$$

$$\langle P_{3/2}, 2, 1 | \hat{d} | S, \psi_- \rangle \approx \frac{d_2}{2\sqrt{2}}. \quad (6.57)$$

To first order in $\mu_B B/\omega_q$, the contribution from the $|P_{3/2}, F = 2, m_F = 1\rangle$ state does not modify the differential light shift. The relevant correction comes from the $|P_{3/2}, F = 1, m_F = 1\rangle$ state, giving

$$\delta\omega_q(B) \approx -\frac{\omega_q |E|^2}{24} \left(\frac{d_1^2}{\Delta_1^2} + \frac{d_2^2}{\Delta_2^2} \right) + \frac{d_1^2 |E|^2}{24\Delta_1} \frac{4\mu_B B}{\omega_q} - \frac{d_2^2 |E|^2}{24\Delta_2} \left(\frac{2\mu_B B}{\omega_q} \right). \quad (6.58)$$

The critical field is estimated by setting this differential light shift to zero,

$$B_c \approx \frac{\omega_q^2 (d_1^2/\Delta_1^2 + d_2^2/\Delta_2^2)}{2\mu_B (2d_1^2/\Delta_1 - d_2^2/\Delta_2)}. \quad (6.59)$$

Using reduced matrix elements $d_1 = 3.3251(21)$ a.u. and $d_2 = 4.7017(27)$ a.u., this estimate gives $B_c \approx 1.07$ G at 532 nm for $^{133}\text{Ba}^+$. This differs from the full theoretical treatment by about 10 percent, mainly because higher-lying states are neglected in the estimate.

6.5 Experimental Fitting Details

The Ramsey curves in Figure 6.5 were fit with

$$P(t) = A_0 \exp\left(-\frac{t}{\tau}\right) \sin^2(\pi f_{\text{Ramsey}} t + \alpha t^2), \quad (6.60)$$

where A_0 is the initial fringe amplitude, τ is the coherence time, and f_{Ramsey} is the Ramsey oscillation frequency. The additional term αt^2 accounts phenomenologically for slow frequency drifts during the data set. Two relevant time scales of amplitude noise were observed. Noise on short time scales causes decoherence on the millisecond scale, while slow thermal drifts in the laser intensity change the Ramsey frequency over tens of minutes. The magic polarization reduces sensitivity to both types of noise by nulling the differential light shift.

The uncertainty in the angle between the laser wave vector $\vec{k}$ and the magnetic field $\vec{B}$ was estimated from the beam divergence,

$$\theta = \tan^{-1}\left(\frac{w - w_0}{2F}\right), \quad (6.61)$$

where w_0 is the beam waist at the focus, w is the beam radius at the focusing lens, and F is the focal length. The beam was focused with a 400 mm focal length lens and positioned using a high-accuracy three-dimensional micrometer stage. This gives a beam divergence

of approximately 13.5 mrad.

The fidelity of quantum information stored in the g qubit after a given number of m -qubit π pulses was estimated as

$$\mathcal{F} = \left(1 - e^{-\tau_{\text{gate}}/\tau_{\text{coherence}}}\right)^n, \quad (6.62)$$

where τ_{gate} is the gate time, $\tau_{\text{coherence}}$ is the measured coherence time, and n is the number of gate operations. Using the measured decay envelopes, the non-magic condition gives an estimated error of $2.7(2) \times 10^{-5}$ per m -qubit π pulse. The magic condition gives $3(2) \times 10^{-7}$ per π pulse. For two-qubit gates, the corresponding estimates are $2.4(2) \times 10^{-4}$ and $3(2) \times 10^{-6}$ for the non-magic and magic conditions, respectively.

CHAPTER 7

Efficient loading of barium ions using autoionizing resonances

“All science is either physics or stamp collecting.” – Ernest Rutherford

After the construction of the second-generation $^{133}\text{Ba}^+$ trap, ion loading became one of the main experimental challenges. This was primarily due to the much smaller trap volume compared with the first-generation apparatus. In addition, the amount of $^{133}\text{Ba}^+$ available in the loading target was reduced by roughly a factor of 100. The new trap was also designed to be loaded by photoionization, which can further reduce the loading efficiency compared with less isotope-selective methods. Taken together, these factors made the expected loading time impractically long. We therefore set out to identify the most efficient photoionization scheme for reducing the loading time of $^{133}\text{Ba}^+$ in the new apparatus.

This chapter describes a theoretical and experimental study of autoionizing resonances in neutral barium, with the goal of improving ion trap loading efficiency. The work was motivated in part by the efficient loading schemes demonstrated by Greenberg *et al.* (N. Greenberg and Senko, 2024). Here, we focus on autoionizing resonances reached from the neutral barium $5d6p\ ^3\text{D}_1^o$ state. We calculate and measure the photoionization cross sections of several resonances in the region from 450 to 454 nm, and we identify a broad and particularly useful autoionizing resonance at 530.8 nm. The 530.8 nm transition combines a large cross section with a broad spectral width and the practical availability of high-power lasers at this wavelength.




Figure 7.1: Level diagram of the photoionization schemes measured in this work. The scheme using resonant excitation at 413 nm followed by excitation to an autoionizing resonance at 530.8 nm was found to be the most efficient pathway for loading barium ions into the trap.

Figure 7.1 summarizes the photoionization pathways considered in this work. Previous studies have measured photoionization cross sections and ion-trap loading rates using the $6s6p\ ^1P_1^\circ$ state (Leschhorn et al., 2012; N. Greenberg and Senko, 2024; White et al., 2022; Armstrong et al., 1993), the $6s6p\ ^3P_J^\circ$ state (Steele et al., 2007), and the $5d6p\ ^3D_1^\circ$ state (Wood et al., 1993). The measurements presented here allow direct comparison with the most efficient previously demonstrated barium loading schemes. We find that the 530.8 nm autoionizing transition is approximately an order of magnitude more efficient than the best previously known pathway (N. Greenberg and Senko, 2024; White et al., 2022; Steele et al., 2007; Leschhorn et al., 2012).

Several features make the 530.8 nm transition especially useful for experiments. First, the second photon absorption remains linear with laser power over the range measured here, with no observed saturation. Second, high optical power is readily available near this

wavelength. Third, the transition lies in the visible spectrum, which should reduce the anomalous charging effects associated with ultraviolet photoionization light (Harlander et al., 2010; Lucas et al., 2004; Qiao et al., 2021). Finally, the same wavelength region is also useful for high-fidelity stimulated Raman gates in Ba^+ (Boguslawski et al., 2023), making the required laser technology compatible with other parts of a barium quantum information experiment.

The theoretical methods used to identify useful barium photoionization pathways are based on techniques that have been developed for two-valence-electron systems over several decades. In this work, the calculations use variational R -matrix theory combined with multichannel quantum defect theory, following the general approach of Refs. (Aymar et al., 1996; M. Aymar and Himdy, 1983). The goal of the calculation is to predict the positions, widths, and relative cross sections of autoionizing resonances that can be accessed from experimentally convenient intermediate states. The starting point is an effective single-electron model potential $V_{lj}(r)$ that describes the interaction between one valence electron and the closed-shell ionic core. Here l is the electron orbital angular momentum, j is the electron total angular momentum, and r is the distance between the electron and the nucleus of the closed-shell ion. The potential is local in r for each l, j channel and is optimized to reproduce the energy levels of Ba^+ as accurately as possible for $l < 5$. The model potential includes the screening of the valence electron by the closed-shell core electrons, the core polarizability interaction, and the spin-orbit interaction of the valence electron in the field of the doubly charged ion. The dielectronic polarization term was not included in the present calculation. This choice is consistent with earlier variational eigenchannel R -matrix calculations that found good agreement with experiment for photoionization from both the $6s6p\ ^1P_1^o$ and $5d6p\ ^3D_1^o$ states in this spectral region without including that term (Wood et al., 1993; Armstrong et al., 1993; Greene and Aymar, 1991). The resulting calculation provides the expected autoionizing resonance structure and guides the experimental search for transitions with large effective loading rates.

7.1 Experimental Procedure

The relative photoionization cross sections were measured by producing a plume of neutral barium atoms inside the trap volume and counting the number of ions loaded into the trap. Neutral barium was generated by laser ablation of a BaCl_2 target. The ablation light came from a 532 nm Nd:YAG laser focused to a beam waist of approximately $300\text{ }\mu\text{m}$, with a pulse energy of 0.5 mJ. The target was positioned 4.3 cm from the trap center. The first step of the photoionization process used approximately $500\text{ }\mu\text{W}$ of 413 nm light focused to a beam waist of about $50\text{ }\mu\text{m}$. This light resonantly excited neutral barium from the $6s^2\text{}^1\text{S}_0$ ground state to the $5d6p\text{}^3\text{D}_1^o$ state. The 413 nm beam was vertically polarized and overlapped with the 493 nm and 650 nm Doppler-cooling beams used for trapped Ba^+ . These beams passed through the trap center perpendicular to the direction of the ablation plume in order to minimize Doppler shifts on the first excitation step. The second-step spectroscopy laser was propagated approximately antiparallel to the 413 nm beam and was horizontally polarized. In the first set of measurements, this laser was scanned between 450.5 and 454.8 nm to drive transitions from the $5d6p\text{}^3\text{D}_1^o$ state to autoionizing states. In the second set of measurements, the spectroscopy laser was tuned near the broad resonance at 530.8 nm.

Measurements were averaged over multiple ablation shots, typically using 30 shots per data point. To reduce sensitivity to slow drifts in the ablation plume, data at different wavelengths were interleaved when comparing cross sections. The ablation pulse energy was monitored using a Gentec energy meter and remained stable over the course of the measurements. All loaded barium isotopes were recorded. To place the measured cross sections on an absolute scale, the experimental cross section at 413 nm was set equal to the calculated value of 12 Mb, and the other measured cross sections were scaled relative to this value. A related $1 + 1$ photoionization scheme using 413 nm + 413 nm light from the same intermediate state has also been demonstrated previously (Leschhorn et al., 2012).

The ion trap used for these measurements is a four-blade segmented linear Paul trap




Figure 7.2: Spectrum of autoionizing resonances originating from the $5d6p\ ^3D_1^o$ state. The black points show the measured cross section, and the red curves show the theoretical calculation. The calculated peak positions are in reasonable agreement with the measured frequencies, although the peak heights of the narrower resonances differ from the unbroadened theoretical prediction.

made from gold-coated ceramic blades. The ion-to-electrode distance is $200\ \mu\text{m}$, and the trap drive frequency is $\Omega_{\text{RF}} = 2\pi \times 32.0\ \text{MHz}$. The RF voltage amplitude was $200\ \text{V}$, giving a radial secular frequency of $\omega_{\text{radial}} = 2\pi \times 2.0\ \text{MHz}$ and a radial trap depth of approximately $4.3\ \text{eV}$. Axial confinement was produced by applying $20\ \text{V}$ of static bias to the four outermost electrodes, giving an axial secular frequency of $\omega_{\text{axial}} = 2\pi \times 0.55\ \text{MHz}$. The trap was housed in a sealed vacuum chamber with a measured pressure of approximately $2 \times 10^{-10}\ \text{Torr}$.

7.2 Spectroscopy From the $5d6p\ ^3D_1^o$ State

The first photoionization pathway newly studied in this work begins with excitation at $413\ \text{nm}$ from the neutral barium ground state to the $5d6p\ ^3D_1^o$ state. A second laser is then scanned between 450.5 and $454.8\ \text{nm}$ to excite autoionizing states. The resulting spectrum is shown in Fig. 7.2. The calculated spectrum provides a useful guide to the locations and relative strengths of the measured resonances. However, a direct comparison between the calculated peak heights and the measured peak heights requires accounting for Doppler broadening in the ablation plume. This effect is particularly important for




Figure 7.3: Autoionization resonances to the $5d_{3/2}15d_{3/2} J = 2$ state and the $5d_{3/2}15d_{5/2} J = 1$ state. The unbroadened theoretical calculation is shown in red, while the blue curve shows the same calculation convolved with a Gaussian Doppler profile corresponding to the temperature of the ablation plume. Including Doppler broadening significantly improves the agreement between theory and experiment, especially for the narrower resonance.

narrow resonances. Even if the intrinsic peak cross section of a narrow resonance is large, the effective cross section sampled by a hot atomic plume can be much smaller because only a fraction of the atoms are resonant with the laser at any given frequency.

The temperature of the ablation plume was estimated by measuring the broadening of the first-step $6s^2\ ^1S_0 \rightarrow 5d6p\ ^3D_1^o$ transition at 413 nm. The measured line shape was fit to a Voigt profile, and the Gaussian contribution was interpreted as Doppler broadening. This procedure gave a temperature of approximately 350 K. A Gaussian profile with a full width at half maximum corresponding to this temperature was then convolved with the calculated autoionization spectrum. The effect of this convolution is shown in Fig. 7.3 for the two strongest resonances in the scanned region. Including Doppler broadening brings the calculated cross section into much better agreement with the measured peak heights, especially for the narrower feature. This result indicates that the effective cross section observed during trap loading can be strongly limited by the velocity distribution

of the neutral atoms. A colder atomic source, such as loading from a magneto-optical trap, should therefore allow narrower autoionizing resonances to reach substantially larger effective cross sections (Johansen et al., 2022).

7.3 530.8 Transition

After surveying the calculated spectrum and accounting for the practical effects of Doppler broadening, a broad resonance near 530.8 nm was identified as a particularly promising photoionization pathway. This resonance drives from the $5d6p\ ^3D_1^o$ state to the $5d_{3/2}8d_{5/2}\ J = 2$ autoionizing state. It has a large predicted cross section, a broad linewidth, and lies at a wavelength where high-power lasers are readily available. These three properties make it especially attractive for ion trap loading.

Precision spectroscopy of the 530.8 nm transition was performed using a 1.1 W fiber laser system. The measured spectrum is shown in Fig. 7.4. The measured cross section of the resonance is

$$\sigma_{530.8} = (7400 \pm 870) \text{ Mb}, \quad (7.1)$$

which is consistent with the theoretical calculation within the experimental uncertainty. The broad width of this resonance is important. For the narrow resonances between 450 and 454 nm, Doppler broadening of the ablation plume suppresses the experimentally accessible peak cross section. In contrast, the 530.8 nm resonance remains broad enough that the measured line shape and cross section agree well with the unbroadened theoretical prediction. This makes the transition more robust under the conditions normally encountered during ablation loading of an ion trap.

7.4 Comparison with Previous Photoionization Schemes

To compare the newly characterized pathways with existing methods, we also measured the previously studied scheme in which neutral barium is resonantly excited at 554 nm to




Figure 7.4: Spectrum of the autoionizing resonances near 530.8 nm from the $5d6p\ ^3D_1^o$ state. The black points show the measured spectrum, and the red curve shows the theoretical prediction. The broad resonance agrees well with the theoretical calculation without applying additional Doppler broadening.

the $6s6p\ ^3P_1^o$ state and then ionized nonresonantly with 405 nm light. This measurement allows direct comparison with Greenberg *et al.*, who reported both the 405 nm nonresonant scheme and a resonant autoionizing scheme at 389 nm (N. Greenberg and Senko, 2024). In the present work, the cross section at 405 nm was measured to be

$$\sigma_{405} = (72 \pm 3) \text{ Mb}, \quad (7.2)$$

which agrees well with the calculated value of 76 Mb.

The 389 nm autoionizing transition was not measured directly in this work. Instead, the loading results reported in Ref. (N. Greenberg and Senko, 2024) are used as an experimental benchmark. Greenberg *et al.* reported that the 389 nm scheme loaded approximately seven times more efficiently than the 405 nm nonresonant scheme. In our model, the ratio of the corresponding calculated cross sections is approximately 17. No separate theoretical model was presented in Ref. (N. Greenberg and Senko, 2024). The 530.8 nm transition provides a substantial improvement over these previous schemes. Its measured cross section




Figure 7.5: Average number of ions loaded per ablation shot as a function of 530.8 nm laser power. The red dashed line shows a linear fit to the data with $\chi^2/\text{dof} = 1.01$. The linear dependence shows that the loading rate does not saturate over the measured power range.

is about an order of magnitude larger than the most efficient previously demonstrated scheme. In addition, the available optical power at 530.8 nm is much larger than the power typically available for the ultraviolet and near-ultraviolet photoionization wavelengths used in other schemes. The combination of large cross section and high available power directly translates into faster loading.

The power dependence of the 530.8 nm scheme is shown in Fig. 7.5. The loading rate remains linear over the measured range, with no visible saturation. This behavior is expected because the autoionizing state has a lifetime on the order of picoseconds, which gives a very high saturation intensity. The calculated saturation intensity is approximately

$$I_{\text{sat}} = 7.1 \text{ MW/cm}^2. \quad (7.3)$$

The first-step 413 nm power was held fixed during these measurements, while the second-step 530.8 nm power was varied. The observed linearity indicates that additional laser




Figure 7.6: Comparison of relevant photoionization cross sections for barium loading. The blue star shows the measured cross section of the 530.8 nm transition. Black triangles show other cross sections measured in this work. The black circle shows the measured cross section reported in Ref. (N. Greenberg and Senko, 2024) for photoionization from the $6s6p\ ^3P_1^o$ state. Red squares show theoretical calculations from this work.

power at 530.8 nm can be used directly to increase the ion loading rate, at least until other effects such as trap-loading dynamics or RF-induced heating become important.

A summary of the relevant photoionization pathways is shown in Fig. 7.6. The first new pathway, using an autoionizing transition near 454 nm, gives a loading efficiency and available laser power comparable to the previously known 389 nm autoionization scheme. The second pathway, using the resonant 530.8 nm transition, provides an approximately order-of-magnitude increase in cross section compared with the 389 nm scheme. This improvement is compounded by the practical availability of high laser power near 530.8 nm.

Another practical benefit is that 530.8 nm is in the visible spectrum. Photoionization with ultraviolet light can produce anomalous charging of trap surfaces, which can shift ion positions and change trapping potentials (Harlander et al., 2010; Lucas et al., 2004; Qiao et al., 2021). A visible photoionization beam should reduce this effect, making the loading process less disruptive for precision ion-trap experiments and integrated trap technologies.

The improvement in loading efficiency is especially important for experiments with trace or radioactive barium isotopes. For $^{133}\text{Ba}^+$, the available source material is limited, and radiation-safety considerations favor using the smallest possible quantity of radioactive material. An efficient loading scheme makes it possible to produce ions quickly while using less source material and fewer ablation shots. This reduces both experimental downtime and the total amount of radioactive material required in the apparatus.

Efficient loading is also important for scaling trapped-ion quantum computers. Even with naturally abundant isotopes, loading a large processor one ion at a time can become a significant overhead. If loading a moderate-size chain requires minutes, then loading larger systems or replacing lost ions during operation becomes increasingly impractical. A photoionization scheme with a large cross section, high available laser power, and reduced surface charging can make the loading process faster and more reliable. The 530.8 nm scheme therefore provides both an immediate practical advantage and a route toward more scalable barium-based ion trapping. It is compatible with visible-wavelength optical systems, uses a strong and broad autoionizing resonance, and can operate far from saturation with linearly increasing loading rate as optical power is increased. These properties make it useful not only for present $^{133}\text{Ba}^+$ experiments, but also for future trapped-ion processors that require faster and more reliable loading.

In this work, we identified and characterized autoionizing resonances in neutral barium for efficient ion trap loading. The measurements focused on transitions from the $5d6p\ ^3D_1^o$ intermediate state. Several narrow resonances between 450 and 454 nm were measured and compared with theoretical calculations. Although the unbroadened calculations predict large peak cross sections for some of these features, Doppler broadening in the ablation plume reduces the effective cross section observed during trap loading. This effect was confirmed by convolving the theoretical spectrum with a Gaussian Doppler profile corresponding to a plume temperature of approximately 350 K. A broad resonance at 530.8 nm was identified as the most useful transition for efficient loading. This resonance has a measured cross section of (7400 ± 870) Mb and remains well matched to the theoretical

prediction without additional broadening. The loading rate increases linearly with laser power over the measured range, with no evidence of saturation. Compared with previously demonstrated loading pathways, this scheme provides an order-of-magnitude improvement in loading efficiency, together with the practical advantages of high available laser power and visible-wavelength operation. These results provide an efficient path for loading barium ions, especially scarce or radioactive isotopes such as $^{133}\text{Ba}^+$. The 413 nm + 530.8 nm scheme should reduce loading times, improve ion yield, and decrease the amount of source material required for experiments. More generally, the results show that autoionizing resonances can be used as a practical tool for improving ion-trap loading efficiency in systems where source material is limited or where fast loading is important for scalable operation.

7.5 Doppler Broadening

Ablation loading produces a plume of neutral atoms with a finite velocity distribution. This velocity distribution leads to Doppler broadening of the optical transitions used for photoionization. For broad autoionizing resonances, this broadening has little effect on the measured peak cross section. For narrow resonances, however, Doppler broadening can strongly reduce the effective cross section sampled by the laser. The temperature of the neutral barium plume was estimated from spectroscopy of the first-step transition,

$$6s^2\ ^1S_0 \rightarrow 5d6p\ ^3D_1^o, \quad (7.4)$$

at 413 nm. The measured line shape was fit to a Voigt profile, and the Gaussian component was attributed to Doppler broadening. This gave an effective plume temperature of approximately

$$T \approx 350\text{ K}. \quad (7.5)$$

To compare the measured autoionization spectra with theory, the calculated spectra were convolved with a Gaussian profile corresponding to this temperature. This convolution significantly improved the agreement between the theoretical and measured peak heights for the narrow resonances between 450 and 454 nm. The result indicates that the discrepancy between the unbroadened theory and the measured peak heights is not primarily due to an error in the calculated resonance strengths. Rather, it is a practical consequence of performing spectroscopy on a thermal ablation plume. This also suggests that the same narrow resonances could be more useful in experiments with colder neutral sources. If the neutral atoms were prepared in a magneto-optical trap or another colder source, the Doppler width would be reduced, and the experimentally accessible peak cross section of narrow resonances would increase.




Figure 7.7: A frequency scan of the 413 nm transition was fit with a Voigt profile, from which a Gaussian component of $\sigma = 335$ MHz was extracted. This Doppler width corresponds to a temperature of 318 K at the center of the trap. Time-of-flight data were also taken at this ablation fluence and were used to measure an atom velocity of 205 m/s at the trap center, corresponding to a temperature at ablation of 350 K. The center frequency of 725259000 MHz agrees with the previously reported spectroscopy in Ref. (Choi et al., 2024).

7.6 Calculation of Saturation Intensity

We extract the transition dipole moment associated with each autoionizing resonance by first determining the oscillator strength f from the calculated photoionization cross section.

The oscillator strength is a dimensionless quantity that characterizes the strength of an electric dipole transition. The calculated cross section σ is given in units of megabarns. To relate this to the oscillator strength, we convert σ into the differential oscillator strength $df/d\epsilon$, expressed per unit energy. We use Eq. 2.18 of Ref. (Fano and Rau, 1986):

$$\sigma(\epsilon) = 2\pi^2 \frac{e^2 \hbar}{mc} \frac{df}{d\epsilon}, \quad (7.6)$$

where ϵ is the photon energy, e is the electron charge, m is the electron mass, and c is the speed of light. Using Eq. (7.6), we convert the calculated cross section spectrum into $df/d\epsilon$. Near an isolated autoionizing resonance, we approximate the differential oscillator strength by a Lorentzian profile:

$$\frac{df}{d\epsilon} = \frac{f_s}{\pi} \frac{\gamma}{(\epsilon - \epsilon_0)^2 + \gamma^2}, \quad (7.7)$$

where f_s is the integrated oscillator strength of the resonance, ϵ_0 is the resonance energy, and γ is the half-width at half-maximum. This approximation is valid for well-isolated resonances where interference effects are negligible. Fitting the converted spectrum to Eq. (7.7) yields an oscillator strength of $f_s = 0.01051$ and a linewidth of $\gamma = 0.1329$ meV for the 530.8 nm resonance. The oscillator strength is related to the electric dipole matrix element by Eq. 2.23 of Ref. (Fano and Rau, 1986):

$$f_s = \frac{2m\omega_s}{\hbar} |\langle s|x|0 \rangle|^2, \quad (7.8)$$

where ω_s is the transition angular frequency. Solving for the dipole matrix element gives

$$|\langle s|x|0 \rangle| = \sqrt{\frac{\hbar f_s}{2m\omega_s}}. \quad (7.9)$$

Substituting $f_s = 0.01051$ and $\omega_s = 2\pi \times 564.7 \times 10^{12}$ Hz yields a dipole moment

$$|\langle s|ex|0 \rangle| = 2.05 \times 10^{-30} \text{ C} \cdot \text{m}. \quad (7.10)$$

Finally, using the extracted dipole moment and the measured linewidth of 32.1 GHz, we estimate the saturation intensity using

$$I_{sat} = \frac{c\epsilon_0\hbar^2\Gamma^2}{2|\langle e|\mu|g\rangle|^2}. \quad (7.11)$$

This yields a saturation intensity of 7.1 MW/cm².

CHAPTER 8

State Preparation and Measurement in $^{133}\text{Ba}^+$

“There are two possible outcomes: if the result confirms the hypothesis, then you’ve made a measurement. If the result is contrary to the hypothesis, then you’ve made a discovery” – Enrico Fermi

8.1 What is SPAM?

State preparation and measurement, hereafter referred to as SPAM, is the process of preparing a qubit in either of its two states (or one of several states in the case of a qudit) and subsequently measuring that state. More precisely, the SPAM fidelity is defined as

$$F_{\text{SPAM}} = \frac{1}{2} [P(0|0) + P(1|1)]. \quad (8.1)$$

The SPAM error can also be written in terms of the incorrect assignment probabilities as

$$\epsilon_{\text{SPAM}} = \frac{1}{2} [P(1|0) + P(0|1)]. \quad (8.2)$$

Experimentally, $P(0|0)$ is estimated by repeatedly preparing the qubit in $|0\rangle$ and recording the fraction of shots assigned to outcome 0. Similarly, $P(1|1)$ is estimated by repeatedly preparing the qubit in $|1\rangle$ and recording the fraction of shots assigned to outcome 1. The SPAM fidelity is then obtained by averaging these two correct-assignment probabilities.

Now that we have a formal definition of SPAM fidelity, it is important to ask how

SPAM fits into the broader picture of a quantum computer. SPAM errors are often treated as secondary to gate errors and are not always considered a dominant limitation in quantum-computing architectures. This does not mean that large SPAM errors are considered acceptable. Rather, the usual intuition is that errors from single-qubit and two-qubit gates are more important because they occur many more times during a computation. This intuition is easy to understand. State preparation is typically performed once at the beginning of a computation for each qubit, while measurement is performed once at the end. By contrast, single-qubit and two-qubit gates may be applied many, or even millions, of times over the course of an algorithm. From this perspective, SPAM appears to contribute only a small number of error opportunities compared with the large number of gates in a deep circuit.

However, as discussed earlier, the long-term goal is to build a *fault-tolerant* quantum computer. This will almost certainly require quantum error correction, which changes the role of state preparation and measurement. In an unencoded circuit, preparation and measurement appear only at the beginning and end of the computation. In an error-corrected circuit, however, additional syndrome ancillas are repeatedly prepared and measured in order to extract information about errors. These ancillas are entangled with the data qubits and then measured. The measurement outcomes are used to infer which errors have occurred, either leading to active feedback or, more commonly, to an update of the Pauli frame in software. Thus, SPAM is no longer only an initialization and final-readout operation. It becomes part of the repeated error-correction cycle.

To see this more concretely, consider a rotated surface-code patch of distance d . Such a patch contains d^2 data qubits and approximately $d^2 - 1$ syndrome ancillas. During each syndrome-extraction cycle, every syndrome ancilla is prepared, entangled with neighboring data qubits, and measured. The number of ancilla preparations and measurements per cycle is therefore

$$N_{\text{prep}} = d^2 - 1, \quad N_{\text{meas}} = d^2 - 1, \quad (8.3)$$

or $N_{\text{SPAM}} = 2(d^2 - 1)$ SPAM operations per cycle. The number of data-ancilla two-qubit

gates is obtained by summing the weights of the stabilizer checks. For the rotated surface code, the bulk checks are weight four and the boundary checks are weight two, giving

$$N_{2q} = 4d(d - 1) \quad (8.4)$$

two-qubit gates per syndrome-extraction cycle. Thus, the number of SPAM operations and the number of two-qubit gates have the same leading-order scaling with code distance, $N_{\text{SPAM}} \sim 2d^2$ and $N_{2q} \sim 4d^2$. SPAM is therefore smaller than the two-qubit gate count by only a constant factor, not by a different scaling law. Since syndrome-extraction cycles are repeated throughout a fault-tolerant computation, SPAM errors scale with the size and duration of the computation in the same broad sense as gate errors. They should therefore be treated as part of the repeated physical error budget, rather than as a one-time boundary cost.

Another important practical consideration is the time required for SPAM compared with typical gate operations. Time is a crucial resource in a quantum computation because memory errors accumulate throughout the algorithm. As qubits sit idle, the probability of decoherence increases, and coherent errors can continue to build up if they are not sufficiently randomized or corrected. In many trapped-ion architectures, SPAM operations, especially state detection, are much slower than the average gate operation. Therefore, it is important to evaluate state-preparation and measurement methods in terms of both fidelity and duration. This is one of the major strengths of $^{133}\text{Ba}^+$. Its level structure enables state preparation that is both high fidelity and fast, reducing the SPAM contribution to both the error budget and the total computation time.

8.2 State Preparation

One of the key advantages of trapping the radioactive isotope $^{133}\text{Ba}^+$ is the extremely low-error state preparation. Among the ion species considered here, $^{133}\text{Ba}^+$ is unusual in having nuclear spin $I = 1/2$, similar to $^{171}\text{Yb}^+$. This level structure enables a dissipative




Figure 8.1: Level diagram of the typical optical pumping scheme in $^{137}\text{Ba}^+$ using σ^+ polarization.

optical-pumping scheme that does not rely on precise polarization control. As shown in Figure 8.1, state preparation in $^{137}\text{Ba}^+$ is typically performed by applying light with a well-defined polarization to pump population into one of the stretched states. The population can then be transferred to the desired qubit manifold using microwave pulses. A major error source in this approach is the difficulty of maintaining high polarization purity at the ion. Small temperature changes, optical drifts, and changes in the direction of the quantization axis set by the magnetic field can reduce the purity of the optical-pumping light. These effects can limit the state-preparation fidelity to approximately 99%, while also requiring frequent calibration.

By contrast, in an $I = 1/2$ ion, mixed-polarization light can be used for optical pumping. Because the $|F = 0, m_F = 0\rangle$ state is well separated in frequency from the $F = 1$ manifold, population accumulates in this state and becomes trapped there, as shown in Figure 8.2. This makes the preparation protocol less sensitive to polarization imperfections and reduces the experimental overhead required for stable operation. A caveat is that




Figure 8.2: Level diagram of the typical optical pumping scheme in $^{133}\text{Ba}^+$ using mixed polarization.

recent work has shown that high state-preparation fidelity can also be achieved in isotopes with $I \neq 1/2$ by reshuffling population within the ground-state manifold and applying multiple rounds of polarized-light depletion operations (An et al., 2022b). While this approach can greatly improve the preparation fidelity, it comes at a substantial cost in time. In $^{133}\text{Ba}^+$, the optical-pumping steady state can be reached in approximately $1.5 \mu\text{s}$, whereas the preparation sequence in Ref. (An et al., 2022b) takes roughly 3.5 ms . In addition to this slowdown, the method introduces significant practical overhead. The 1762 nm beam requires substantially higher optical power than the lasers used for cooling and preparation, and in many experiments this beam is implemented globally rather than locally. Similarly, microwave operations can introduce cross talk and often require dedicated hardware. In contrast, state preparation in an $I = 1/2$ ion can be implemented by simply turning off a single EOM tone on the Doppler-cooling laser.

Now, we can derive the fundamental limit to the optical pumping fidelity of an $I = 1/2$

ion. Because the ground-state hyperfine splitting is finite, the same light can also off-resonantly scatter from the nominally dark state $|6S_{1/2}, F = 0, m_F = 0\rangle$. Thus we can estimate the ratio between the desired on-resonant scattering rate out of the $F = 1$ manifold and the unwanted off-resonant scattering rate out of the $F = 0$ target state. For a two-level transition with natural linewidth Γ , Rabi frequency Ω , and detuning Δ , the photon scattering rate is

$$R_{\text{sc}}(\Delta) = \frac{\Gamma}{2} \frac{s}{1 + s + (2\Delta/\Gamma)^2}, \quad (8.5)$$

where the saturation parameter is

$$s = \frac{2|\Omega|^2}{\Gamma^2}. \quad (8.6)$$

In the weak-saturation limit, $s \ll 1$, this becomes

$$R_{\text{sc}}(\Delta) \simeq \frac{\Gamma|\Omega|^2}{\Gamma^2 + 4\Delta^2}. \quad (8.7)$$

On resonance, $\Delta = 0$, giving

$$R_{\text{sc}}(0) = \frac{|\Omega|^2}{\Gamma}. \quad (8.8)$$

In what follows, Ω_q denotes the Rabi frequency associated with polarization component q . For the $F = 1 \rightarrow F' = 1$ transition in a $J = 1/2 \rightarrow J' = 1/2$ system with $I = 1/2$, the allowed transition strengths are $|C_{m_F, q}|^2 = \frac{1}{2}$. Using Eq. (8.8), the on-resonant scattering rates out of the three Zeeman states in the $F = 1$ manifold are therefore

$$R_{-1}^{\text{on}} = \frac{1}{2\Gamma} (|\Omega_0|^2 + |\Omega_+|^2), \quad (8.9)$$

$$R_0^{\text{on}} = \frac{1}{2\Gamma} (|\Omega_-|^2 + |\Omega_+|^2), \quad (8.10)$$

$$R_{+1}^{\text{on}} = \frac{1}{2\Gamma} (|\Omega_-|^2 + |\Omega_0|^2). \quad (8.11)$$

For balanced mixed-polarization light,

$$|\Omega_-|^2 = |\Omega_0|^2 = |\Omega_+|^2 \equiv \Omega^2, \quad (8.12)$$

so each state in the $F = 1$ manifold scatters at the same rate

$$R_{F=1}^{\text{on}} = \frac{\Omega^2}{\Gamma}. \quad (8.13)$$

For the $F = 0 \rightarrow F' = 1$ transition, each polarization component has angular strength $|C_{0,q}|^2 = \frac{1}{2}$. Using Eq. (8.7), the total off-resonant scattering rate out of the $F = 0$ target state is

$$R_{F=0}^{\text{off}} = \frac{\Gamma}{\Gamma^2 + 4\Delta_{\text{hf}}^2} \frac{1}{2} (|\Omega_-|^2 + |\Omega_0|^2 + |\Omega_+|^2). \quad (8.14)$$

For balanced mixed-polarization light,

$$R_{F=0}^{\text{off}} = \frac{3}{2} \frac{\Gamma \Omega^2}{\Gamma^2 + 4\Delta_{\text{hf}}^2}. \quad (8.15)$$

Dividing Eq. (8.15) by Eq. (8.13) gives

$$\frac{R_{F=0}^{\text{off}}}{R_{F=1}^{\text{on}}} = \frac{3}{2} \frac{\Gamma^2}{\Gamma^2 + 4\Delta_{\text{hf}}^2}. \quad (8.16)$$

For $^{133}\text{Ba}^+$,

$$\frac{R_{F=0}^{\text{off}}}{R_{F=1}^{\text{on}}} = \frac{3}{2} \frac{(20.1 \times 10^6)^2}{(20.1 \times 10^6)^2 + 4(9.925 \times 10^9)^2}. \quad (8.17)$$

Evaluating this expression gives

$$\frac{R_{F=0}^{\text{off}}}{R_{F=1}^{\text{on}}} = 1.538 \times 10^{-6}. \quad (8.18)$$

This number is not directly the state-preparation infidelity. It is only the ratio of the unwanted off-resonant scattering rate from the nominally dark $F = 0$ state to the desired resonant scattering rate out of the $F = 1$ manifold. For the closed $S_{1/2}$ - $P_{1/2}$ model,




Figure 8.3: 7 level system simulation to model the effective pumping rate and fidelity of $^{133}\text{Ba}^+$ optical pumping with mixed polarization light at a 5 G magnetic bias field.

resonant scattering from $F = 1$ enters the target $F = 0$ state with branching probability $1/3$, while off-resonant scattering from $F = 0$ leaves the target state with probability $2/3$. The steady-state error is therefore approximately

$$\epsilon_{\text{prep}} \simeq \Gamma_{\text{out}}/\Gamma_{\text{in}} = 2R_{F=0}^{\text{off}}/R_{F=1}^{\text{on}} = 2(1.538 \times 10^{-6}) = 3.08 \times 10^{-6} \quad (8.19)$$

It is useful to compare the analytical estimate with a numerical simulation of the $S_{1/2}$ and $P_{1/2}$ manifolds. In the simulation, the resonant $6S_{1/2}, F = 1 \leftrightarrow 6P_{1/2}, F' = 1$ optical-pumping dynamics are treated explicitly, while the far-detuned coupling from the target state $|6S_{1/2}, F = 0, m_F = 0\rangle$ to the $P_{1/2}$ manifold is treated perturbatively. This removes the need to numerically resolve the fast off-resonant optical coherence. The eliminated coupling is included as an effective AC Stark shift of the target state and as effective Lindblad scattering terms that describe off-resonant depumping from $F = 0$. Since the resulting Lindblad master equation is time independent, the density matrix is vectorized and evolved using the Liouvillian propagator, $\rho_{\text{vec}}(t) = \exp(\mathcal{L}t)\rho_{\text{vec}}(0)$. The




Figure 8.4: Comparison of different intensities to ultimate optical pumping fidelity.

resulting ground-state populations demonstrate agreement with the analytical model and are shown in Figure 8.2. While this is the theoretical limit to the optical pumping fidelity, there is a tradeoff between the speed of pumping and the steady-state fidelity. Due to the effect of power broadening if the pumping occurs at a higher intensity (or s value) the pumping will reach steady state faster, but at the cost of power broadening the transition and increasing the rate of the off resonant scattering from the $|F = 0, m_f = 0\rangle$ state. We can model this effect for various values of the intensity (relative to the saturation intensity s). This comparison is shown in Figure 8.4.

An important and subtle experimental discovery was made in connection with this power-broadening issue. Historically, Doppler cooling was performed with the carrier of the 493 nm laser resonant with the $F = 1 \rightarrow F' = 0$ transition. This choice was motivated by two considerations. First, the available sideband power was relatively small, while decay from $F' = 0$ to the $F = 0$ ground state is suppressed. As a result, this transition provided a nearly closed cycling loop. This was especially useful because the EOM used at the time had a bandwidth of only 8 GHz, so the required 11.2 GHz tone had to be generated using a second-order sideband. Second, before the implementation of electron-shelving




Figure 8.5: EOM sideband power distribution at $\beta = 1.84$. This parameter quantifies the RF power input strength.

detection, state readout was performed by attempting to cycle on this transition without decay into the $F = 0$ state. This method was known as hyperfine-selective fluorescence readout. Both of these constraints have been removed in the current experimental setup because a higher-bandwidth EOM is now available. However, as a consequence of the historical operating procedure, the laser carrier had remained on the $F = 1 \rightarrow F' = 0$ transition.

This configuration required an additional EOM tone near 1.84 GHz to drive the $F = 1 \rightarrow F' = 1$ transition for optical pumping into the $F = 0$ ground state. Ideally, this sideband would be generated while completely depleting the carrier, but in practice complete carrier suppression is difficult to achieve. We recently found that this optical-pumping method produced a significantly higher preparation error. The origin of this error can be understood by noting that the higher-order EOM sidebands place additional optical power closer to the off-resonant scattering pathway out of the $|F = 0, m_F = 0\rangle$ dark state. The relevant sideband powers are shown in Figure 8.5. The effect of this EOM-based pumping scheme can be seen by comparing the steady-state optical-pumping




Figure 8.6: Comparison of different intensities to ultimate optical pumping fidelity, while using an EOM to drive the $|F = 1 \rightarrow F = 1\rangle$ transition.

fidelity in Figure 8.4 with the corresponding result in Figure 8.6. This error channel can be eliminated by moving the laser carrier to the $F = 1 \rightarrow F' = 1$ transition. In that configuration, optical pumping is performed directly by the carrier and no EOM tones are required.

8.3 Single Qubit Gates

To prepare the $|F = 1, m_F = 0\rangle$ state, we apply a single microwave π pulse that transfers population from the $|F = 0, m_F = 0\rangle$ state. This step requires extremely high-fidelity single-qubit control, for which we have developed and implemented extensive calibration metrics. At present, this microwave transfer is the main bottleneck to demonstrating higher-fidelity fast SPAM for both qubit states. Although microwave single-qubit gates with errors at the 1×10^{-7} level have been demonstrated in other systems (Smith et al., 2025), the microwave gates in our lab currently have an error rate of approximately 3×10^{-5} . The goal is to deliver the correctly calibrated microwave amplitude and frequency to the ion with high precision and accuracy, repeatedly and stably over the course of a




Figure 8.7: Diagram of the calibration sequence using a Ramsey experiment.

multi-hour SPAM measurement. Calibration begins by determining the qubit transition frequency to approximately 3 Hz resolution using a Ramsey experiment. Figure 8.7 shows the experimental Ramsey sequence. In this sequence, the qubit is first placed into the superposition state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ by applying a microwave $\pi/2$ pulse. The qubit is then allowed to freely evolve for a variable delay time, after which a second microwave $\pi/2$ pulse is applied. In the ideal case where the microwave frequency is exactly resonant with the qubit transition, the qubit phase remains fixed relative to the microwave reference during the free-evolution time. Therefore, the second $\pi/2$ pulse maps the state to $|1\rangle$ independent of the delay time. If the applied microwave frequency is detuned from the true qubit frequency, the qubit phase slips relative to the microwave reference during the idle time. As a result, the effect of the second $\pi/2$ pulse depends on the accumulated phase. This produces an oscillation in the measured qubit-state population as a function of delay time. The frequency of this Ramsey oscillation is extracted to determine the detuning between the applied microwave frequency and the true qubit transition frequency.




Figure 8.8: Single qubit gate benchmarking. A sequence of π pulses with a random phases are applied to extract the single pulse gate error.

The microwave amplitude is calibrated by applying a sequence of repeated π pulses, all with the same phase. If the pulse area is exactly correct, each π pulse rotates the qubit by 180° , so successive pulses simply move the qubit back and forth between $|0\rangle$ and $|1\rangle$. If the microwave amplitude is slightly too large or too small, each pulse produces a small over-rotation or under-rotation. By repeating the pulse many times, this small error accumulates, making the total rotation error easier to measure. The sequence is repeated for different numbers of π pulses to track how the accumulated error grows with pulse number. A final $\pi/2$ pulse is then added as an analysis pulse, which makes the measurement sensitive to the sign of the accumulated error and allows us to distinguish between over-rotation and under-rotation.

After the amplitude calibration is complete, we benchmark the microwave single-qubit gate fidelity by applying sequences of successive π pulses. Before each pulse, the microwave phase is randomized so that coherent amplitude errors do not add constructively over the full sequence. The measured population is then recorded as a function of the number of




Figure 8.9: Dissipative preparation of the $|1\rangle$ state after the application of a microwave π pulse.

applied π pulses and fit to a linear model. The slope of this fit gives the error accumulated per pulse, which is used to extract the single-pulse gate fidelity. This method is not without limitations. In particular, any time-dependent change in microwave power during the multi-pulse sequence can skew the measured slope, causing the extracted error rate to differ from the true single-pulse gate error. The results of this measurement are shown in Figure 8.8. At the present level of approximately 3×10^{-5} error per microwave π pulse, the limiting mechanisms are likely no longer simple static calibration errors. Possible contributions include residual detuning during the gate, microwave amplitude noise, phase noise, pulse-shape transients, timing or DDS synchronization errors, AC Zeeman shifts during the pulse, weak off-resonant coupling to other hyperfine states, and technical drifts in the microwave delivery chain. In addition, the measured error can be biased by the benchmarking protocol itself, since time-dependent power changes, imperfect state preparation, and detection errors can all affect the extracted slope. These effects occur on timescales or through mechanisms that are not fully removed by the Ramsey frequency calibration or the repeated-pulse amplitude calibration.

In order to achieve higher fidelity preparation of the $|1\rangle$ state, population can be pumped dissipatively to this state after the microwave π pulse. Population in the other Zeeman states of the $F = 1$ ground-state manifold is removed by driving electric-quadrupole transitions at 1762 nm to the $5D_{5/2}$ manifold. The addressed $m_F = \pm 1$ states are coupled to the accessible $5D_{5/2}$ states with $m_F = \pm 2$, while the desired $|F = 1, m_F = 0\rangle$ state is left dark to this excitation. Population transferred to $5D_{5/2}$ is then repumped with 614 nm light through the $6P_{3/2}$ manifold, from which it spontaneously decays back to the $6S_{1/2}$ ground-state manifold. Repeating this excitation and repump cycle continually removes population from the unwanted $F = 1$ Zeeman states, while population that decays into $|F = 1, m_F = 0\rangle$ remains there, causing the target state to accumulate dissipatively. This method is shown diagrammatically in Figure 8.9

8.4 Electron Shelving

After the $|1\rangle$ state is prepared, population in this state is shelved to the $D_{5/2}$ manifold. This shelving allows the Doppler cooling lasers to produce fluorescence only when population remains in the $|0\rangle$ state. The $D_{5/2}$ manifold is not addressed by the Doppler cooling light, so population shelved there decays primarily through the natural lifetime of the state. Shelving is achieved in a two-step process. First, four pulses of 1762 nm light drive the quadrupole transition to levels with $\Delta m_F = \pm 2$. In the second stage, incoherent optical pumping with 455 nm, 650 nm, and 585 nm light pumps population through the $|P_{3/2}, F = 2\rangle$ state.

The 1762 nm multi-pulse shelving technique, referred to hereafter as cabinet shelving, was first demonstrated in Ref. (An et al., 2022b). The shelving levels are chosen to reduce the probability of off-resonant deshelling back to one of the $m_F = \pm 1$ Zeeman sublevels in the ground state. This off-resonant coupling can occur because the relevant Zeeman splittings differ only through a second-order perturbation. Each shelving pulse has a fidelity of approximately 0.9. In the ideal case, four sequential pulses would therefore




Figure 8.10: The temperature dependence of the cabinet shelving method is compared for two different starting thermal distributions.

suppress the shelving error to roughly $(0.1)^4 = 1 \times 10^{-4}$. In practice, however, the fidelity of later pulses in the sequence can be reduced because the 1762 nm transition is sensitive to the motional state of the ion. The ion begins in a thermal mixture of motional states, and the shelving pulse does not couple all motional states with equal strength. Lower- n states are generally closer to the conditions for which the pulse was optimized, so they are shelved more efficiently by the first few pulses. As the sequence proceeds, this colder population is preferentially removed from the unshelved ensemble. The population remaining for later pulses is consequently biased toward hotter motional states, which have different effective Rabi frequencies and poorer overlap with the calibrated pulse. The later pulses therefore exhibit lower apparent shelving fidelity, not necessarily because the pulse quality itself has degraded, but because they act on a progressively more difficult-to-shelve subset of the original thermal distribution. This effect is modeled in Figure 8.10, for a starting average Fock occupational number of 500, and 50. Although this is much larger than the doppler temperature, the trap does exhibit a high heating rate, and depending on the idle time




Figure 8.11: Simulation of the rate of population shelving from the $F = 1$ state versus the rate population accidentally falls into the $F = 0$ causing an error.

this higher temperature could be reached.

After the cabinet shelving pulses, three lasers are applied simultaneously. The 455 nm laser is resonant with $|S_{1/2}, F = 1\rangle \rightarrow |P_{3/2}, F = 2\rangle$, the 650 nm laser is resonant with $|D_{3/2}, F = 1\rangle \rightarrow |P_{1/2}, F = 0\rangle$, and the 585 nm laser is resonant with $|D_{3/2}, F = 2\rangle \rightarrow |P_{3/2}, F = 2\rangle$. These excited states are chosen so that spontaneous decay back to $|S_{1/2}, F = 0\rangle$ is forbidden. Such a decay would constitute an error during shelving, since population in $|S_{1/2}, F = 0\rangle$ would be incorrectly identified as unshelved. The dominant mechanism by which population can leak into $|S_{1/2}, F = 0\rangle$ during this sequence is therefore off-resonant excitation to $|P_{3/2}, F = 1\rangle$, which can subsequently decay to $|S_{1/2}, F = 0\rangle$. A second layer of protection is provided by using only π -polarized light. For pure π polarization, the transition $|S_{1/2}, F = 1, m_F = 0\rangle \rightarrow |P_{3/2}, F = 1, m_F = 0\rangle$ is forbidden by the relevant selection rule. As a result, leakage into $|S_{1/2}, F = 0\rangle$ can occur only through a higher-order process. For example, population must first decay into one of the $|S_{1/2}, F = 1, m_F = \pm 1\rangle$ states. From there, the off-resonant transition $|S_{1/2}, F = 1, m_F = \pm 1\rangle \rightarrow |P_{3/2}, F = 1, m_F = \pm 1\rangle$ can be weakly driven, followed by decay into $|S_{1/2}, F = 0\rangle$. This combination of hyperfine selection and polarization selection strongly suppresses leakage into the other qubit state. This error source is enhanced when the 455 nm light is power broadened, since power broadening increases the rate of off-resonant excitation. We model this effect using a simple




Figure 8.12: Simulation of the rate of population shelving from the $F = 1$ state versus the rate population accidentally shelved that began in the $F = 0$ causing an error. The starting population is assumed to have been cabinet shelved with 1×10^{-3} error if it began in the $F = 1$.

rate-equation simulation, shown in Figure 8.11. The simulation captures the observed scale of the error and gives a shelving error of approximately 3×10^{-4} .

A related error source occurs during the 455 nm optical-pumping shelving step if the ion is prepared in the $|0\rangle$ qubit state. In this case, off-resonant excitation out of $|0\rangle$ can transfer population into the shelved manifold, producing a bright-to-dark error. Given the measured rate for this process, together with the residual population left in $|1\rangle$ after cabinet shelving, we can predict the optimal 455 nm intensity and shelving time that minimizes the total SPAM error. Figure 8.12 shows the result of this rate-equation simulation, assuming that 1×10^{-3} of the population remains in $|1\rangle$ after cabinet shelving. The simulation indicates that a dark-state shelving error of 2×10^{-6} and a bright-state shelving error of 1×10^{-5} are achievable.

8.5 State Detection

To determine whether the $|0\rangle$ or $|1\rangle$ state was prepared, the Doppler cooling lasers are turned on after shelving is complete, and fluorescence photons are collected on two PMTs above and below the trap. The counts from the two PMTs are summed during




Figure 8.13: Theoretical photon-count distributions for a $300 \mu\text{s}$ detection window, including the high-count tail of the dark distribution caused by decay out of the shelved D state.

a fixed detection time t_{det} . A threshold n_{th} is chosen a priori. If the total photon count is greater than n_{th} , the ion is declared bright and measured as $|0\rangle$. If the count is less than or equal to n_{th} , the ion is declared dark and measured as $|1\rangle$. For a bright ion, the Doppler cooling lasers drive a fluorescence cycle with detected count rate $R_{\text{B}} = \eta R_{\text{sc}} + R_{\text{bg}}$, where η includes the photon collection efficiency, optical transmission, and PMT quantum efficiency, and R_{bg} includes background light and PMT dark counts. The mean bright count is $\bar{n}_{\text{B}} = R_{\text{B}} t_{\text{det}}$, and the photon counts are modeled by a Poisson distribution,

$$P_{\text{B}}(n) = \frac{\bar{n}_{\text{B}}^n e^{-\bar{n}_{\text{B}}}}{n!}. \quad (8.20)$$

For a shelved ion, the count rate is ideally set only by background counts, giving a mean count $\bar{n}_{\text{bg}} = R_{\text{bg}} t_{\text{det}}$. However, the shelved D state has a finite lifetime. If the ion decays during the detection window, it can return to the Doppler cooling cycle and begin fluorescing. The dark-state distribution is therefore not a single Poisson distribution.

Instead, it is a mixture of Poisson distributions weighted by the probability that the ion decays at each time during the detection window.

Let $\Gamma_D = 1/\tau_D$ be the decay rate of the shelved state. If the ion decays at time t_d , it fluoresces for the remaining time $t_{\text{det}} - t_d$, so the mean number of detected photons is

$$\bar{n}(t_d) = \bar{n}_{\text{bg}} + R_B(t_{\text{det}} - t_d). \quad (8.21)$$

The resulting dark-state count distribution is

$$P_D(n) = e^{-\Gamma_D t_{\text{det}}} \frac{\bar{n}_{\text{bg}}^n e^{-\bar{n}_{\text{bg}}}}{n!} + \int_0^{t_{\text{det}}} \Gamma_D e^{-\Gamma_D t_d} \frac{\bar{n}(t_d)^n e^{-\bar{n}(t_d)}}{n!} dt_d. \quad (8.22)$$

The first term describes ions that remain shelved for the full detection time and produce only background counts. The second term describes ions that decay during detection and then fluoresce for the remaining part of the detection window. This decay process adds a high-count tail to the dark-state distribution. For the threshold rule used here, the probability of incorrectly identifying a bright ion as dark is the probability that the bright-state photon count falls at or below threshold,

$$\epsilon_{B \rightarrow D} = \sum_{n=0}^{n_{\text{th}}} P_B(n). \quad (8.23)$$

Similarly, the probability of incorrectly identifying a dark ion as bright is the probability that the dark-state photon count falls above threshold,

$$\epsilon_{D \rightarrow B} = \sum_{n=n_{\text{th}}+1}^{\infty} P_D(n). \quad (8.24)$$

The total threshold-detection error is set by the overlap of the bright-state Poisson distribution with the decay-modified dark-state distribution. The optimal threshold is




Figure 8.14: Detection-time optimization for fluorescence state discrimination. The background and bright count rates are calibrated from the measured photon counts at a reference detection time, then rescaled for each detection time in the scan. At each detection time, the discriminator threshold is chosen to minimize the equal-prior average detection error. The optimum balances improved separation between the bright and dark photon-count distributions at longer detection times against the increased probability of decay out of the shelved $D_{5/2}$ state.

chosen to minimize the average detection error,

$$\epsilon_{\text{det}}(n_{\text{th}}) = \frac{1}{2} [\epsilon_{\text{B} \rightarrow \text{D}}(n_{\text{th}}) + \epsilon_{\text{D} \rightarrow \text{B}}(n_{\text{th}})]. \quad (8.25)$$

Figure 8.13 shows the expected photon-count distributions for a 300 μs detection window using the experimentally observed mean counts. The optimal threshold is chosen to minimize the average of the bright-to-dark and dark-to-bright detection errors.

Using the known background and bright fluorescence counts we can determine the optimal collection time base on the theoretical distributions. Figure 8.14 shows the average detection error for an optimized detection threshold given different detection times. Using this simulation an optimal time of 180 μs is determined to produce an average detection

State Detection Counts




Figure 8.15: Experimental SPAM data currently demonstrating the highest fidelity measured SPAM of any qubit platform.

error of $\approx 2 \times 10^{-6}$.

Figure 8.15 shows the experimentally measured bright- and dark-state photon-count distributions. At the beginning of each experimental shot, the ion is Doppler cooled for 2 ms. The photon counts collected during this cooling period are used as a preselection check. If the counts do not exceed a chosen threshold, the shot is rejected. This protects the data set against rare events such as background-gas collisions or excess motional heating. During Doppler cooling, the frequency of the 493 nm light is rapidly shifted using a double-pass AOM, allowing the cooling light to be placed farther red of resonance for more efficient cooling. This preselection step is justified because it occurs before any state preparation or quantum operation.

After Doppler cooling and preselection, either $|0\rangle$ or $|1\rangle$ is chosen at random and prepared. The same shelving and detection sequence is then applied to both input states. Shots intended to prepare $|1\rangle$ are shown in red, while shots intended to prepare $|0\rangle$ are shown in blue. The full sequence from state preparation through detection takes

approximately $700\ \mu\text{s}$.

The present experimental SPAM fidelity is limited by the microwave single-qubit gate fidelity used in state preparation. As discussed above, this limitation can be removed by replacing the microwave preparation step with a dissipative preparation scheme, which is currently being implemented.

CHAPTER 9

Future Experiments: The Next Generation

“Dreams are maps.” – Carl Sagan

9.1 Where is the Bottleneck to Building a Fault-Tolerant Quantum Computer Now?

A practical roadmap toward fault-tolerant trapped-ion quantum computing can be divided into three stages. These stages are closely connected and will likely be developed in parallel, but each provides a clear experimental milestone.

9.1.1 Ion Loading, Shuttling, and Recooling

A large-scale trapped-ion quantum computer will require a reliable method for supplying ions, transporting them through the processor, and restoring their motional state after transport. These operations are central to the quantum charge-coupled device architecture. As the system grows, loading and transport can no longer be treated as occasional calibration procedures. They must instead become routine operations that are automated, repeatable, and compatible with continuous quantum error correction.

The loading process begins with the production of a neutral atomic beam. Conventional resistively heated ovens can generate sufficient atomic flux, although they also introduce heat, outgassing, and material deposition inside the vacuum system. These effects are particularly undesirable in cryogenic processors and densely integrated traps. A promising

alternative is the miniature atomic source developed by Quantum Fabrix. This source uses optical heating and integrated collimation to produce a tightly directed atomic beam with a comparatively small thermal load. A well-collimated beam is useful because it directs the barium flux toward the loading region while reducing the deposition of barium on trap electrodes, optical surfaces, and nearby dielectric materials. This becomes increasingly important when the same trap must operate reliably for months or years.

The loading region itself should be designed to maximize the probability that a newly ionized atom is captured. An ideal loading trap would therefore combine a large loading volume with a large trap depth. The loading volume is the region in which the neutral atomic beam, the photoionization beams, and the confining potential overlap. A larger volume relaxes alignment requirements and allows more neutral atoms to pass through a region in which they can be ionized and captured. A larger trap depth increases the range of initial positions and kinetic energies from which a newly created ion can remain confined. Together, these properties can increase the loading rate and improve its robustness against variations in atomic flux, beam alignment, and laser intensity.

These requirements may favor a dedicated loading zone whose geometry differs from that of the computation zones. Computation regions are normally optimized for low motional heating, high optical access, and strong gate interactions. The loading region can instead be optimized for deep confinement, high capture probability, and resistance to contamination. Once an ion has been loaded and Doppler cooled, it can be transported into a cleaner region of the processor. Separating loading from computation also confines barium deposition and loading-light exposure to a replaceable or less sensitive part of the device.

Loading efficiency can be further improved by using a resonantly enhanced autoionization pathway for barium. The recently demonstrated barium autoionization scheme provides a substantially more efficient route from a neutral excited state into the ionization continuum than nonresonant photoionization. Increasing the ionization probability reduces the neutral atomic flux, laser intensity, and exposure time required to load a single ion.

This is valuable for a fault-tolerant processor because lower atomic flux reduces electrode contamination, while shorter optical exposure reduces the opportunity for loading light to charge dielectric surfaces. Efficient autoionization may also improve isotope-selective loading, which is important when different barium isotopes are assigned to computational, cooling, or auxiliary roles.

The exposure of dielectrics to ultraviolet and near-ultraviolet loading light is itself an important engineering concern. Photoemission and charge trapping can produce slowly varying electric fields that displace the ions from the radio-frequency null. These fields increase micromotion and require repeated compensation. Future loading systems should prevent loading beams from directly illuminating exposed dielectric surfaces. This can be achieved through dedicated loading zones, tightly focused and collimated optical beams, metallic shielding, integrated apertures, and carefully positioned beam dumps. Trap designs that place exposed dielectric material outside the direct line of sight of the loading beams would further improve long-term stability.

Three-dimensional and multilayer trap fabrication may provide an important route toward traps with larger loading volumes, greater trap depths, and improved protection from dielectric charging. Conventional surface-electrode traps are attractive because they are compatible with planar microfabrication. Their open geometry, however, generally produces shallower confinement than electrode structures that surround the ion more completely. Surface traps also contain dielectric gaps near the electrodes that can be illuminated by scattered laser light. Recent three-dimensional printed ion traps have demonstrated that complex electrode structures can produce substantially greater trap depths than comparable planar geometries. Multilayer fabrication could provide similar advantages by allowing electrodes, shielding layers, and electrical routing to be placed at different heights.

In these geometries, elevated or surrounding metal electrodes can screen the ion from exposed dielectric surfaces. They can also increase the separation between the active trapping region and dielectric material within the substrate. This arrangement should

reduce the sensitivity of the ion to charge deposited on the dielectric. It could therefore support both deeper confinement and greater stability under the ultraviolet illumination required for loading. Three-dimensional fabrication also permits trap junctions and transport channels to be shaped more freely than is possible in a single planar metal layer.

Increasing the speed of ion transport is another important objective. In a large processor, ions may need to undergo repeated translations, separations, combinations, and junction crossings during every error-correction cycle. Even when each operation has high fidelity, slow transport can increase the total cycle time and expose the qubits to additional memory errors. Faster shuttling would allow ions to be rearranged with less latency and could increase the rate at which stabilizer measurements are repeated.

The practical transport speed is limited by the available confinement, the smoothness of the trapping potential, the bandwidth of the control electronics, and the amount of motional excitation that can be tolerated. Deep traps provide greater confinement during rapid waveform changes and reduce the likelihood of ion loss at junctions or during separation operations. Three-dimensional printed and multilayer traps may therefore enable faster transport by providing deeper and more symmetric potentials, particularly in junction regions where planar traps often have weak confinement or pseudopotential barriers. These geometries are not sufficient by themselves. They must be combined with low-noise voltage sources, high-bandwidth waveform generators, accurate models of the electrode potentials, and calibrated transport waveforms.

Rapid transport will generally leave some residual motional excitation. The processor must either suppress this excitation through optimized waveforms or remove it efficiently after transport. Waveform design can minimize excitation by controlling the acceleration of the potential minimum and by avoiding frequencies that resonantly drive the ion motion. For more demanding operations, sympathetic recooling can be performed using a second ion species or an auxiliary ion that does not store computational information. The cooling ion can remove shared motional energy while allowing the internal state of the data ion to remain coherent.

The relevant benchmark is therefore not the fidelity of a single shuttling operation. A useful assessment must measure the complete repeated sequence of transport, separation, recombination, junction crossing, recoupling, and gate execution. It should include ion-loss probability, motional excitation, qubit dephasing, waveform reproducibility, and the time required for the full sequence. These measurements should be performed over the number of repetitions expected during an error-correction experiment.

The long-term objective is an autonomous ion-supply and transport system. A processor should be able to detect the loss of an ion, load a replacement from a collimated atomic source, verify the isotope, cool the ion, transport it to the required location, and return the affected region to operation. Achieving this objective will require progress in atomic sources, efficient barium autoionization, deep loading traps, three-dimensional or multilayer electrode fabrication, ultraviolet-light management, fast control electronics, and sympathetic cooling. These capabilities would transform loading and shuttling from laboratory procedures into dependable components of a fault-tolerant trapped-ion computer.

9.1.2 Parallel Operations and Integrated Laser Delivery

At present, one of the most immediate limitations to scaling trapped-ion systems in academic laboratories is laser delivery. A small experiment can rely on free-space beams that are aligned through vacuum windows and directed toward one or two trapping regions. This approach becomes increasingly difficult as the number of ions and operating zones grows. Each zone may require wavelengths for cooling, repumping, state preparation, coherent control, shelving, and measurement. A fault-tolerant processor must therefore deliver many optical wavelengths to many locations without requiring a separate free-space beam path for every combination of wavelength and trapping zone.

The first requirement is a multi-wavelength optical distribution network that can supply light from the near-ultraviolet to the near-infrared. This large wavelength range makes it difficult to use a single waveguide material, coupler design, or focusing element

with uniformly high efficiency. Future systems must extend current integrated delivery capabilities from a few static outputs to a network that can deliver all required wavelengths to tens or hundreds of independently controlled regions.

Optical delivery must also support the tight spatial pitch needed for dense ion registers. Ions within a chain are commonly separated by only a few micrometers. Individual addressing therefore requires beam waists that are substantially smaller than the ion spacing and have sufficiently weak sidelobes to suppress crosstalk. A small beam waist does not reduce the intensity required at the ion. Instead, it produces the required intensity using less total optical power. This is particularly important for quadrupole transitions and stimulated-Raman gates, which can otherwise require substantial laser power. Tight focusing also limits the area illuminated by scattered light and reduces unintended interactions with neighboring ions.

Hybrid two-dimensional and three-dimensional photonic architectures provide a promising method for achieving broadband, tightly focused delivery. In such an architecture, lithographically patterned waveguides and planar lenses route and shape the light before a three-dimensionally printed reflective micromirror directs it toward the ion. Reflective freeform optics, including parabolic and biconic-aspheric micromirrors, are attractive because their operation can be much less wavelength dependent than that of grating couplers. A single reflective element can therefore deliver several wavelengths to the same ion or register. Three-dimensional printing also allows the optical surface to be shaped for micron-scale focusing at ion heights that are difficult to address using planar structures alone.

These focusing elements must be combined with fast optical routing. Passive splitters can distribute one laser to many zones, although passive distribution alone sends light to locations where it may not be needed and divides the available optical power. A large processor will require active on-chip multiplexing so that a limited number of laser sources can be connected to different trapping zones at different times. Integrated Mach-Zehnder interferometers, electro-optic phase shifters, tunable couplers, and high-extinction

modulators could provide this functionality. The switching time must be short compared with an error-correction cycle so that cooling, gates, state preparation, and measurement can be scheduled independently and in parallel across the processor.

The required switching network must provide more than simple on and off control. It should control the destination, amplitude, frequency, and phase of each beam. High extinction is necessary to prevent residual light from producing Stark shifts, spontaneous scattering, or unwanted qubit rotations in idle zones. Stable splitting ratios are necessary when nominally identical operations are applied in parallel. The system must also avoid converting laser-frequency drift, temperature drift, or electrical control noise into changes in the optical power delivered to the ions.

Optical loss is a central constraint on this architecture. Every fiber interface, waveguide crossing, splitter, switch, bend, and output element contributes some insertion loss. As the number of components grows, even modest losses at each element can produce a large reduction in the power reaching the ions. The immediate consequence is that more laser power must be supplied at the input. More importantly, absorption within the chip converts part of this optical power into heat. Short-wavelength propagation is particularly challenging because scattering and absorption generally increase toward the ultraviolet. Low-loss waveguides, efficient fiber interfaces, reflective output optics, and routing architectures with a limited number of crossings are therefore essential.

Heat dissipation is especially important for cryogenic ion traps. The available cooling power at the lowest temperature stage is limited. Heating the trap can increase electric-field noise, change mechanical alignment, and degrade cryogenic pumping. Excess heat may also increase local outgassing and reduce the quality of the vacuum environment. Similar considerations apply to microwave conductors, radio-frequency resonators, electronic switches, and on-chip modulators. A scalable processor must minimize both optical absorption and electrical dissipation rather than assuming that additional laser or microwave power can be supplied without consequence.

A second major concern is photoinduced charging. Integrated photonics places dielectric

waveguides, cladding layers, lenses, and support structures close to the ion. Light scattered from these components can charge exposed dielectric surfaces and generate electric fields at the ion. These fields shift the ion position, increase micromotion, modify the axial confinement, rotate the motional axes, and complicate transport between zones. The problem is particularly severe at ultraviolet and visible wavelengths.

The silicon substrate can introduce a separate charging mechanism. Light reaching the silicon can generate mobile charge carriers and change the electrical response of the trap. This may alter the radio-frequency impedance, modify the trapping pseudopotential, or produce time-dependent fields that cannot be removed by a single static compensation. These effects demonstrate that an integrated optical component cannot be evaluated independently of the trapping fields.

Future devices should therefore shield both dielectric and semiconductor regions from the ion. Possible approaches include burying the photonic layer beneath continuous metal electrodes, adding a grounded plane above the silicon substrate, coating exposed optical dielectrics with a transparent conductor, and using metallic baffles to intercept scattered light. Three-dimensional fabrication is particularly attractive because it permits metallized overhangs and shielding structures that would be difficult to fabricate in a planar process. A three-dimensionally printed trap or optical structure could place grounded metal between the ion and the exposed dielectric while preserving a clear path for the intended beam. Three-dimensional fabrication does not remove charging automatically, but it provides additional geometric freedom for controlling which surfaces are visible to the ion.

Integrated photonics may also provide capabilities that are difficult to achieve with conventional free-space optics. Waveguides can divide a common laser into two paths with a stable relative geometry. The outputs can then be designed to form a standing wave at the ion. Control of the node, antinode, and field gradient can be used to select whether an ion couples primarily to the optical field or to its spatial gradient. This provides another method for suppressing unwanted carrier excitation, controlling motional sidebands, and engineering spatially selective interactions.

The same architecture may improve the phase stability of two-qubit gates. In free-space Raman systems, beam paths that travel through separate optical components accumulate differential phase noise from vibration, air currents, and thermal expansion. Routing both beams through a common substrate or through closely matched waveguides can reject much of this common-mode motion. Integrated phase shifters could then provide controlled adjustments of the phase at each trapping zone. This would allow the optical phase required for an entangling gate to be set electronically and reproduced across repeated error-correction cycles.

The long-term goal is an optical control layer that distributes multiple wavelengths across the processor, switches them between zones on the timescale of a quantum circuit, and produces tightly focused beams at the ions with low crosstalk. It must perform these functions while maintaining low optical loss, low heat dissipation, stable phase, and minimal charging. Achieving all of these requirements simultaneously remains a major engineering challenge. Integrated waveguides, active photonic switches, reflective micromirrors, three-dimensional optical structures, grounded shielding layers, and phase-stable interferometric circuits provide a plausible route toward this objective. Success would allow many trapping zones to operate in parallel and would remove one of the largest remaining barriers to scaling trapped-ion experiments beyond the capabilities of conventional academic laboratories.

9.1.3 QEC Syndrome Extraction

The most important next step toward fault-tolerant trapped-ion quantum computing is the demonstration of a complete, repeated syndrome-extraction cycle. Individual components of this process have been demonstrated, including high-fidelity gates, ion transport, mid-circuit measurement, and real-time classical control. However, these capabilities have not yet been combined into a system that can rapidly bring an ancilla into contact with data qubits, extract a syndrome, separate and measure the ancilla, and return the result to the quantum controller without introducing substantial errors into the

neighboring data qubits. Demonstrating this full sequence repeatedly and at a speed that suppresses logical errors remains a central challenge.

In a QCCD architecture, an ancilla ion may be prepared in a dedicated region and transported to the data qubits. It must then interact with several data ions through entangling gates so that the value of a stabilizer is mapped onto the ancilla. The ancilla must subsequently be separated from the data and either measured locally or transported to a dedicated detection region. After measurement, it must be reset, recooled if necessary, and made available for the next error-correction cycle. Every part of this sequence contributes to the duration and error probability of syndrome extraction.

The readout operation is particularly difficult because trapped-ion detection is normally dissipative. A bright ion scatters many resonant photons during fluorescence detection. These photons can be absorbed or scattered by nearby ions, producing depumping, dephasing, Stark shifts, or leakage from the computational basis. The detection beams themselves may also illuminate neighboring ions because of imperfect focusing, beam sidelobes, and optical scattering from the trap. Photon recoil heats the measured ancilla and can heat other ions through their shared motion. If the ions remain in the same crystal during detection, this heating must be removed before another high-fidelity entangling gate can be performed.

One approach is to physically separate the ancilla from the data before detection. The ancilla could be shuttled to a measurement region that is optically isolated from the computation zones. This reduces direct optical crosstalk, although it adds the time and motional excitation associated with separation, transport, and recombination. The measurement region must therefore be close enough to support rapid operation while providing sufficient isolation from resonant detection light. Optical baffles, shielded detection zones, tightly focused beams, integrated photon collection, and spectral isolation between the ancilla and data qubits may all be required.

The speed of this sequence is as important as its fidelity. Data qubits continue to accumulate memory errors while an ancilla is transported, measured, reset, and processed

by the classical controller. A highly accurate syndrome measurement is not useful if it takes so long that the additional memory errors exceed the errors that the measurement is intended to detect. Fault tolerance therefore requires optimization of the complete error-correction cycle rather than independent optimization of the gate and measurement fidelities.

Fast readout requires photons to be collected efficiently. Conventional free-space fluorescence detection often collects only a small fraction of the emitted light. The ion must then scatter many photons before the bright and dark distributions can be distinguished with high confidence. Increasing the numerical aperture of the collection optics and integrating detectors near the trap can reduce the required detection time. Time-resolved decision rules can provide an additional improvement by ending the measurement as soon as the photon record identifies the state with sufficient confidence. This avoids using the same fixed detection time for every measurement outcome.

Optical cavities may provide a more powerful route to fast ancilla readout. A cavity can increase the probability that photons are emitted into a collected optical mode or can convert the internal state of the ion into a change in the phase or amplitude of a reflected probe beam. Either approach could reduce the number of spontaneously scattered photons required to determine the state. Cavity-assisted readout may therefore provide faster detection with less recoil heating and less scattered light reaching the data qubits. The difficulty is integrating a stable, high-finesse cavity with a trap while maintaining optical access, low charging, low motional heating, and compatibility with ion transport.

Nonclassical light may also improve readout. In a dispersive or cavity-assisted measurement, squeezed probe light could reduce the optical noise in the measured quadrature and allow the two ion states to be distinguished using fewer probe photons. This could reduce both the measurement time and the optical disturbance. The advantage will depend strongly on optical loss because loss rapidly destroys squeezing. Squeezed light is therefore most likely to be useful when it is combined with a high-efficiency cavity or collection system rather than conventional low-efficiency fluorescence detection.

A useful QEC readout system does not necessarily need equal performance for both ancilla states. Below threshold, the most common syndrome outcome is normally the outcome corresponding to no detected error. The physical states of the ancilla can therefore be assigned so that the no-error outcome corresponds to the state that is easier to prepare or identify with lower error. This creates a deliberately asymmetric SPAM protocol optimized for the statistics of syndrome extraction.

The optimal assignment depends on the dominant experimental limitation. Mapping the no-error syndrome to a bright state may allow very rapid identification because the first few detected photons provide strong evidence that the ion is bright. An adaptive detector can stop early after these photons arrive. This choice may minimize the average measurement time, although it causes resonant photons to be scattered during nearly every syndrome measurement. If disturbance of neighboring data qubits is the dominant problem, it may instead be advantageous to map the no-error syndrome to a dark state. Most measurements would then produce little or no fluorescence, reducing recoil and optical crosstalk. The cost is that a dark result generally requires waiting long enough to establish that photons are truly absent.

The assignment should therefore be selected using the logical consequences of the two error types. Mistaking an error syndrome for no error can allow a physical error to remain undetected. Mistaking a no-error syndrome for an error can cause an unnecessary correction or create a false detection event. These errors need not have the same effect on logical performance. The decoder should be supplied with the measured asymmetry of the readout rather than assuming that the two outcomes have equal error probabilities. It may also use the full photon-arrival record or an analog confidence value instead of reducing every measurement to an unweighted binary result.

The classical processing of the syndrome is the next part of the bottleneck. A laboratory demonstration can record the ancilla measurements and decode them after the experiment. A fault-tolerant computer cannot rely on this approach. Syndrome results must be transmitted to a decoder, combined with earlier measurements, and converted into a

correction or an updated Pauli frame while the computation is still running. The decoder must sustain the average syndrome production rate without developing an increasing backlog.

For circuits containing only Clifford operations, many corrections do not need to be applied immediately. They can be tracked in software as a Pauli frame and interpreted when later measurements are performed. This flexibility allows some decoder latency to be hidden behind subsequent quantum operations. The situation becomes more demanding when the computation reaches a non-Clifford operation. A Pauli correction cannot always be propagated through a non-Clifford gate without changing the required operation or measurement basis. Magic-state injection and gate teleportation also require measurement-dependent feedforward. The decoder must therefore provide the relevant logical information before the processor reaches one of these decision points.

This requires close integration between the detector electronics, decoder, waveform system, and quantum control hardware. The measurement result cannot pass through a slow sequence of independent laboratory computers and software interfaces. Photon counting, state discrimination, syndrome assembly, decoding, and conditional pulse selection must be implemented as a real-time pipeline. Field-programmable gate arrays, dedicated processors, or application-specific decoding hardware may be required to achieve predictable latency.

The relevant experimental benchmark is not simply the fidelity of an isolated ancilla measurement. A complete benchmark must include ancilla preparation, stabilizer mapping, ion separation, transport, readout, classical decoding, conditional control, reset, and recooling. It must quantify the errors induced on neighboring data qubits while these operations are performed. The sequence must then be repeated for many rounds to determine whether the logical error rate decreases as the code distance is increased.

A convincing demonstration of fault tolerance will therefore require more than the observation of a syndrome. It must show that syndromes can be extracted rapidly without destroying the stored quantum information, that the measurement outcomes can be decoded in real time, and that the resulting information can control later operations when

required. Fast cavity-assisted detection, improved photon collection, squeezed probe fields, isolated measurement zones, biased ancilla SPAM, and hardware-integrated decoding are promising directions. Combining these capabilities into a repeated QEC cycle represents the next major step from high-fidelity trapped-ion components toward an operational fault-tolerant quantum computer.

9.2 SPAM in the Age of Quantum Error Correction

State preparation and measurement (SPAM) errors are often treated as errors that occur only at the beginning and end of a quantum experiment. This picture changes substantially in a quantum error-corrected processor. Error correction requires auxiliary qubits to be repeatedly prepared, entangled with the data qubits, measured, and reset. SPAM operations therefore become part of every error-correction cycle rather than simply determining the accuracy of the final experimental result.

Consider the measurement of a stabilizer operator

$$S_Z = Z_{i_1} Z_{i_2} \cdots Z_{i_w}, \quad (9.1)$$

where w is the weight of the stabilizer. The corresponding syndrome can be measured by preparing an auxiliary qubit, coupling it to each of the w data qubits, and then measuring the auxiliary qubit. Ideally, the measurement returns the eigenvalue of S_Z , indicating whether an error has changed the parity of the participating data qubits.

A measurement error does not necessarily introduce an error into the encoded state. Instead, it may produce an incorrect syndrome and cause the decoder to infer an error that did not occur. For this reason, stabilizer measurements are normally repeated. The decoder examines changes in the measured syndrome,

$$d_j(r) = s_j(r) \oplus s_j(r-1), \quad (9.2)$$

where $s_j(r)$ is the measured value of stabilizer j during round r . A syndrome that changes for only one round is more likely to have resulted from a measurement fault, while a persistent change is more consistent with an error on a data qubit. Repeated syndrome extraction therefore makes imperfect measurement compatible with fault-tolerant operation.

This redundancy does not make SPAM errors unimportant. A larger measurement error produces more false syndrome events and makes the decoding problem more difficult. It may also require additional rounds of syndrome extraction before the decoder can reliably distinguish data errors from measurement errors. More importantly, the data qubits continue to accumulate errors while the auxiliary qubits are prepared, measured, and reset. The duration of the complete error-correction cycle is approximately

$$t_{\text{cycle}} = t_{\text{prep}} + t_{\text{gate}} + t_{\text{meas}} + t_{\text{reset}}, \quad (9.3)$$

where t_{gate} includes the operations required to map the stabilizer parity onto the auxiliary qubit. If the data qubits experience errors at an effective rate Γ_{data} , the memory error accumulated during one cycle is approximately

$$p_{\text{memory}} \simeq \Gamma_{\text{data}} t_{\text{cycle}}. \quad (9.4)$$

Consequently, SPAM speed can be as important as SPAM fidelity. A very accurate but slow measurement may produce fewer readout errors while allowing more errors to accumulate on the data qubits. The relevant system-level quantity is therefore the logical error rate per unit time, rather than the isolated fidelity of a single measurement.

Parallel operation is also essential. A large quantum code may require hundreds or thousands of stabilizers to be measured during every cycle. These measurements must be performed simultaneously whenever possible. The practical performance of a SPAM system must therefore include the effects of parallel readout, such as optical crosstalk, detector crosstalk, correlated state-preparation errors, and disturbance of nearby data qubits. A

high single-qubit measurement fidelity does not guarantee equally good performance when many qubits are measured at once.

The consequences of SPAM errors also depend strongly on their structure. An average SPAM fidelity combines the probabilities $P(1|0)$ and $P(0|1)$, even though these errors need not be equal. Such asymmetry can bias the syndrome record unless it is incorporated into the decoder. Errors that are correlated across multiple qubits are generally more damaging than independent errors with the same average probability. Leakage from the computational basis can be particularly problematic because it may persist across several error-correction cycles and may propagate errors through subsequent two-qubit gates. Calibration drift is another concern because the decoder may perform poorly if the true measurement statistics differ from those assumed in its noise model.

An important distinction is also made between an unflagged error and an erasure. In an unflagged error, the recorded state is incorrect and the location of the error is unknown. In an erasure, the apparatus indicates that a particular result is unreliable. Knowledge of the error location gives the decoder additional information. For a code of distance d , a commonly used condition for correcting t unknown errors and e known erasures is

$$2t + e < d. \tag{9.5}$$

A detected erasure therefore consumes approximately half as much of the code-distance budget as an error at an unknown location. This makes state-resolved loss detection and reliable error flags potentially valuable. However, the usefulness of an erasure flag depends on both its false-negative probability, in which an error is incorrectly accepted, and its false-positive probability, in which valid data are unnecessarily discarded.

The demands placed on SPAM also depend on the chosen code. In the surface code, stabilizer checks are geometrically local and are repeatedly measured over many cycles. Measurement errors are explicitly included in the decoding problem, but slow measurement increases the time available for memory errors to accumulate. Quantum low-density

parity-check (qLDPC) codes can encode many logical qubits in a single block and may greatly reduce the physical-qubit overhead. Their stabilizers remain low weight, but their connectivity is generally nonlocal. A large qLDPC block may therefore require many auxiliary qubits, parallel measurements, and efficient routing or long-range interactions. In this setting, SPAM throughput, crosstalk, and correlated errors can be as important as the fidelity of an individual operation.

Logical initialization and final logical readout can themselves be protected by the code. For example, a logical Z measurement can be performed by measuring the physical data qubits and decoding the resulting bit string. A small number of incorrect physical measurements need not change the inferred logical result. Similarly, an imperfectly prepared logical state can be followed by syndrome extraction to detect and correct local preparation faults. Error correction cannot, however, repair arbitrary preparation errors. A correlated fault that acts as a logical operator, an undetected leakage event, or a systematic measurement bias can remain invisible to the stabilizers and directly produce a logical error.

For these reasons, the performance of a SPAM system intended for quantum error correction cannot be summarized by a single average fidelity. Relevant quantities include the two state-dependent error probabilities, preparation and measurement durations, reset fidelity and duration, leakage probability, erasure-detection fidelity, crosstalk, temporal stability, and correlations between qubits. These quantities must ultimately be evaluated through their effect on the logical error rate of the complete error-correction cycle.

Quantum error correction therefore does not eliminate the importance of high-fidelity SPAM. Instead, it changes the engineering objective. The goal is a preparation and measurement system that is fast, repeatable, parallel, minimally disturbing to stored data, and stable over a very large number of error-correction cycles. A platform that combines high SPAM fidelity with short measurement and reset times can reduce both erroneous syndrome information and the memory errors accumulated during syndrome extraction. This combination is especially important for large, high-rate codes, for which the advantages

of reduced qubit overhead can only be realized if the full syndrome-extraction cycle remains sufficiently fast and reliable.

9.3 Metastable Same-Species Sympathetic Cooling

Motional heating remains an important limitation on the fidelity of trapped-ion entangling gates. Electric-field noise continuously adds energy to the ion crystal, while transport, crystal separation and recombination, fluorescence detection, and imperfect control waveforms can produce additional excitation. Because most trapped-ion entangling gates use the collective motion as a quantum bus, this excitation changes the coupling strength across the motional distribution and can reduce the gate fidelity. A large-scale processor will therefore require a method for periodically restoring the relevant modes near the ground state without destroying quantum information stored in the data ions.

In the optical-metastable-ground architecture currently pursued at IonQ, quantum information is stored in the metastable $D_{5/2}$ manifold while ions in the $S_{1/2}$ ground-state manifold serve as coolant ions. Resonant 493 nm light can then be used to Doppler cool a ground-state ion while the metastable data qubits remain largely dark to the cooling cycle. This provides much of the functionality of a dual-species crystal without requiring a second atomic species. The data and coolant ions retain the same mass and share the same optical hardware, while their electronic manifolds provide the spectral distinction needed to apply dissipative operations selectively.

This assignment is particularly useful when frequent Doppler cooling is required. It is not necessarily optimal for a trap with sufficiently low motional heating. In such a device, the modes may remain cold throughout many gates and transport operations. Occasional resolved-sideband cooling may then be sufficient to remove the relatively small amount of accumulated excitation. This possibility motivates reversing the conventional roles of the two manifolds. Quantum information could be stored in the $S_{1/2}$ clock states while a second ion in the $D_{5/2}$ manifold serves as the coolant.

Storing the data qubits in the ground-state manifold has several potential advantages. The data would not be limited by the finite lifetime of the $D_{5/2}$ state. Spontaneous decay of a metastable data qubit is particularly damaging because it produces leakage and eventually limits the duration of a fault-tolerant computation even when all coherent operations are ideal. Ground-state storage would also reduce the need to transfer data repeatedly between the $S_{1/2}$ and $D_{5/2}$ manifolds. Those transfer operations require additional 1762 nm pulses and can introduce coherent transfer errors, off-resonant excitation, and leakage into unwanted hyperfine or Zeeman states.

The principal disadvantage is the loss of optical isolation from resonant ground-state cooling and detection light. A data ion stored in the $S_{1/2}$ manifold is strongly affected by 493 nm light. Conventional Doppler cooling of a nearby ion would therefore scatter photons from the data ion and destroy its quantum state. The proposed protocol avoids repeated ground-manifold Doppler cooling by performing resolved-sideband cooling within the metastable manifold. Its central requirement is a coolant-reset sequence that returns the metastable coolant to a known internal state while producing negligible scattering, differential light shifts, and off-resonant excitation on the neighboring ground-state data ion.

9.3.0.1 Proposed Cooling Cycle

The coolant ion is first prepared in $|D_{5/2}, F = 2, m_F = 2\rangle$. Two optical fields near 532 nm then drive a stimulated-Raman transition within the $D_{5/2}$ manifold. Their frequency difference is chosen to address the red motional sideband of the hyperfine transition. If ω_{32} is the frequency separating the $F = 2$ and $F = 3$ coolant states and ω_m is the selected motional-mode frequency, the required Raman difference frequency is

$$\omega_1 - \omega_2 = \omega_{32} - \omega_m. \quad (9.6)$$

The pulse transfers population from $|D_{5/2}, F = 2, m_F = 2, n\rangle$ to $|D_{5/2}, F = 3, m_F = 3, n - 1\rangle$.

This is a stretched $\Delta m_F = +1$ transition and has the largest rank-one angular matrix element available for this pair of hyperfine manifolds. The large angular factor increases the achievable two-photon Rabi frequency for a given optical power and detuning. It also helps make the sideband pulse shorter, which reduces the time during which the data ion experiences any residual ac Stark shift from the Raman beams.

In the Lamb-Dicke regime, the red-sideband Rabi frequency for motional level n is approximately

$$\Omega_{n,n-1} \simeq \eta \Omega_c \sqrt{n}, \quad (9.7)$$

where η is the Lamb-Dicke parameter and Ω_c is the corresponding carrier Raman Rabi frequency. The factor of $\sqrt{n}$ causes the red-sideband coupling to vanish when the mode reaches $n = 0$. The motional ground state is therefore dark to an ideal red-sideband pulse. Repeating the cooling cycle preferentially removes population from the excited motional states and accumulates it in the ground state.




Figure 9.1: Proposed metastable same-species sympathetic-cooling cycle and the present optical-reset simulation. The 532 nm Raman red-sideband pulse removes one phonon from a shared motional mode while transferring the coolant within the $D_{5/2}$ manifold. A 614 nm quench, spontaneous emission, recycling of the D manifolds, and a selective 1762 nm pulse return the coolant to its initial state. The neighboring data ion remains in the $S_{1/2}$ clock manifold throughout the cycle. The right panel shows the simulated final ground-state distribution for equal σ^+ , π , and σ^- components in the quench field.

After the red-sideband pulse, the phonon has been mapped onto the internal state of the coolant, but the cooling cycle is not complete until this internal entropy is removed. The proposed reset begins with 614 nm, σ^- -polarized light. This field couples $|D_{5/2}, F = 3, m_F = 3\rangle$ to $|P_{3/2}, F = 2, m_F = 2\rangle$. Spontaneous emission from the $P_{3/2}$ manifold then dissipates the entropy carried by the coolant state. Ideally, the angular-momentum selection rules direct the population into $|S_{1/2}, F = 1, m_F = 1\rangle$.

The excited $P_{3/2}$ state can also decay into either metastable D manifold. Population that returns to $D_{5/2}$ is excited again with 614 nm light. Population that enters $D_{3/2}$ is recycled using 585 nm light. Continued application of these fields removes population

from both metastable manifolds and accumulates it in the $S_{1/2}$ manifold. A spectrally selective 1762 nm pulse then transfers the desired $|S_{1/2}, F = 1, m_F = 1\rangle$ population back to the original $|D_{5/2}, F = 2, m_F = 2\rangle$ coolant state. The coolant is then ready for another red-sideband pulse.

Each ideal cycle removes one phonon from the selected shared mode. The coherent Raman pulse performs the state-selective phonon subtraction, while spontaneous emission during the quench makes the process irreversible. Recoil from the emitted photons adds some motional energy, although this contribution is suppressed in the Lamb-Dicke regime. Net cooling requires the average energy removed by the red-sideband step to exceed the recoil and technical heating added during the reset. The final temperature will be determined by the competition between these processes, imperfect red-sideband transfer, off-resonant carrier excitation, and background motional heating.

9.3.0.2 Present Reset Simulation

An initial simulation was performed to study the quench and repumping portion of the sequence. The coolant was initialized in $|D_{5/2}, F = 3, m_F = 3\rangle$. The 614 nm field was assumed to contain equal σ^+ , π , and σ^- polarization components. The simulation included spontaneous decay and continued optical recycling until all population had left the $D_{5/2}$ and $D_{3/2}$ manifolds.

Under these conditions, the desired state contained 90.69% of the final population. This target state was $|S_{1/2}, F = 1, m_F = 1\rangle$. Its population exceeded 90% after approximately $0.4 \mu\text{s}$. This result shows that the optical part of the reset can be intrinsically fast compared with typical sideband-pulse and transport times. The present limitation is therefore not the rate at which the metastable population is quenched.

The remaining population was distributed among three unwanted states in the $S_{1/2}$ manifold. The final simulated populations were

Final state	Population
$ S_{1/2}, F = 1, m_F = 1\rangle$	90.69%
$ S_{1/2}, F = 1, m_F = 0\rangle$	4.97%
$ S_{1/2}, F = 0, m_F = 0\rangle$	3.39%
$ S_{1/2}, F = 1, m_F = -1\rangle$	0.95%

Table 9.1: Simulated final population after quenching and recycling an initially prepared $|D_{5/2}, F = 3, m_F = 3\rangle$ coolant ion with equal σ^+ , π , and σ^- polarization components.

No population remained in either D manifold. This is an encouraging result because it indicates that the combined 614 nm and 585 nm fields can efficiently empty the metastable manifolds. At the same time, the result identifies the central reset problem. Once population has leaked into an unwanted $S_{1/2}$ hyperfine or Zeeman state, continued application of the quench and repump fields cannot repair it. These fields do not resonantly address the ground-state manifold. A 1762 nm pulse designed only for $|S_{1/2}, F = 1, m_F = 1\rangle$ would therefore return at most 90.69% of the coolant population to the desired initial state.

This reset error would compound over repeated cooling cycles. Population left in the wrong ground-state sublevel would fail to participate in the next intended Raman red-sideband pulse. It could also be transferred unpredictably by later cleanup fields. A single-cycle target population near 90% may be sufficient for an initial proof-of-principle cooling demonstration, but it is unlikely to be sufficient for a long sequence intended to reach a low steady-state motional occupation. Improving the internal-state closure of the reset is therefore the most immediate theoretical task.

9.3.0.3 Improving the Coolant Reset

The first improvement is to repeat the simulation using pure σ^- polarization during the initial 614 nm quench. Beginning from the stretched $|D_{5/2}, F = 3, m_F = 3\rangle$ state, this polarization preferentially selects the intended $|P_{3/2}, F = 2, m_F = 2\rangle$ excited state. Restricting the initially excited state should reduce the number of spontaneous-emission

pathways that populate unwanted ground-state sublevels.

The complete recycling sequence must nevertheless be simulated rather than only the first excitation. Spontaneous decay from $P_{3/2}$ can populate multiple states within the $D_{5/2}$ and $D_{3/2}$ manifolds. Once those states are populated, a single fixed polarization may excite them through different intermediate states and reopen leakage pathways into undesired $S_{1/2}$ levels. The polarization that is optimal for the first quench may therefore not be optimal for clearing all subsequent D -manifold population.

A staged protocol may provide better performance. A pure σ^- pulse could first transfer the stretched $D_{5/2}$ state through the desired excited state. Separate polarization components could then be applied only after the initial population has been removed, allowing residual D -manifold states to be repumped without continuously reopening every transition. The pulse durations, detunings, intensities, and order of the 614 nm and 585 nm fields should be optimized using the full multilevel rate or optical-Bloch model.

Additional cleanup operations may be necessary to reach reset errors compatible with repeated sideband cooling. Microwave pulses could transfer unwanted $S_{1/2}$ population into states that can be selectively shelved or optically pumped. A sequence of state-selective 1762 nm pulses could instead transfer several populated ground states into a common metastable state. These methods would improve the reset fidelity at the cost of additional time and coherent-control errors. The relevant figure of merit is not the fidelity of an individual cleanup pulse, but the probability of returning the coolant to the correct state per unit time while producing acceptably small disturbance of the data ion.

The state-preparation protocol should also be optimized for the statistical structure of cooling. An imperfect reset does not necessarily require a fully general high-fidelity SPAM sequence. It requires a rapid map from the small set of populated leakage states back to one specific coolant state. A purpose-built sequence may therefore be shorter and more robust than a generic state-preparation protocol. If residual leakage can be detected nondestructively, the cooling sequence could also be made heralded, although the added measurement and recooling overhead would need to be included in the total performance.

9.3.0.4 Protection of the Ground-State Data Qubit

The defining requirement of same-species sympathetic cooling is that the data ion remain coherent while the coolant undergoes a dissipative reset. This condition must be tested separately for every optical field. The 614 nm and 585 nm fields do not resonantly address an ion in the $S_{1/2}$ manifold, which provides useful spectral isolation. They can still produce off-resonant scattering and ac Stark shifts at sufficiently high intensity. The differential shift between the two clock states is likely to be much smaller than the common shift, but even a small differential component can accumulate over many cooling cycles.

The 532 nm Raman beams present a different error mechanism. Although these beams are designed to drive a transition within the coolant's $D_{5/2}$ manifold, they can off-resonantly couple the ground-state data ion through dipole-allowed excited states. This produces both spontaneous Raman scattering and coherent ac Stark shifts. The scattering probability decreases rapidly with detuning, while the Raman coupling required for cooling can be maintained by increasing the optical power. This creates a tradeoff between available laser power, sideband speed, differential light shift, and unwanted scattering.

The 1762 nm reset pulse is potentially the most serious source of crosstalk because it directly couples the $S_{1/2}$ and $D_{5/2}$ manifolds. The data and coolant ions are the same isotope and may be separated by only a few micrometers. Frequency selectivity alone may therefore be insufficient unless the two ions experience a controlled differential Zeeman or Stark shift. Tight spatial addressing, a small beam waist, low optical sidelobes, polarization selection, and shaped pulses will be important for ensuring that the reset acts only on the coolant ion.

The disturbance should be characterized after many complete cooling cycles rather than after a single optical pulse. Population measurements can detect unintended excitation or leakage of the data qubit. Ramsey measurements can reveal differential light shifts and dephasing. Repeating the cooling sequence during a spin-echo experiment can separate coherent, reproducible phase shifts from irreversible decoherence. Randomized

benchmarking with interleaved cooling cycles would provide a more operational measure of the error added to a quantum circuit.

Motional coupling creates an additional connection between the coolant and data ions. The two ions must share the mode being cooled, which requires appreciable participation of both ions in that normal mode. At the same time, spontaneous-emission recoil acts locally on the coolant and is redistributed through the shared normal modes. The choice of axial or radial mode, trap frequency, ion spacing, and Raman-beam geometry will therefore influence both the cooling rate and the recoil limit. A complete model should include all relevant modes rather than treating the crystal as a single harmonic oscillator.

9.3.0.5 Experimental Roadmap

The first experimental milestone is a demonstration of the coherent 532 nm Raman transition within the $D_{5/2}$ manifold. Carrier spectroscopy should determine the two-photon Rabi frequency, differential ac Stark shift, polarization dependence, and off-resonant coupling to other hyperfine states. Red and blue motional sidebands can then be resolved to calibrate the Lamb-Dicke parameter and determine whether the available optical power is sufficient for rapid phonon subtraction.

The second milestone is an experimental measurement of the coolant-reset map. The ion should be initialized in $|D_{5/2}, F = 3, m_F = 3\rangle$, exposed to the proposed 614 nm and 585 nm sequence, and analyzed state by state in the $S_{1/2}$ manifold. This measurement would test the predicted 90.69% transfer into $|S_{1/2}, F = 1, m_F = 1\rangle$ and determine whether imperfect polarization, laser detuning, coherent dark states, or unmodeled branching processes produce additional errors. The reset should then be optimized using pure σ^- excitation, staged polarization, and targeted cleanup of the unwanted ground states.

The third milestone is repeated cooling of a single motional mode. The mode can be deliberately heated to a known thermal occupation, followed by a variable number of cooling cycles. Red-to-blue sideband asymmetry can then be used to extract the final

mean occupation $\bar{n}$ and the number of cycles required to approach the steady state. These measurements should be compared with a model including imperfect red-sideband transfer, reset leakage, spontaneous-emission recoil, off-resonant carrier excitation, and background heating.

The fourth milestone is the demonstration of sympathetic cooling in a two-ion crystal. One ion would remain in the $S_{1/2}$ clock manifold as the data qubit while the second ion executes the metastable cooling cycle. The shared motion should cool while the internal state of the data ion remains unchanged. The experiment should measure the cooling rate and final $\bar{n}$ together with the population error, dephasing, leakage, and coherent phase shift induced on the data qubit.

The final milestone is integration with the operations that motivate the protocol. The ion crystal should be transported, separated, recombined, or subjected to a sequence that produces a controlled amount of motional excitation. Metastable sympathetic cooling should then restore the required modes before an entangling gate is performed. The decisive benchmark is whether inserting the cooling sequence improves the total circuit fidelity after accounting for the errors and time introduced by the cooling itself.

If successful, this scheme would provide same-species sympathetic cooling while allowing quantum information to remain in the ground-state clock manifold. It would avoid the finite lifetime of metastable data qubits, reduce the number of coherent manifold-transfer operations applied to the data, and remove the hardware complexity associated with introducing a second atomic species. The present simulation already indicates that the optical reset can transfer more than 90% of the coolant population into the desired ground state on a sub-microsecond timescale. The central unresolved questions are whether this reset can be made sufficiently closed, whether repeated cycles can reach the motional ground state, and whether the entire process can be performed without measurably disturbing a neighboring ground-state data qubit. Answering these questions would determine whether metastable same-species sympathetic cooling can become a practical component of fault-tolerant $^{133}\text{Ba}^+$ processors.

9.4 Autler–Townes Shielding of Readout Light

Fast mid-circuit measurement is a central requirement for fault-tolerant trapped-ion quantum computing. During a stabilizer measurement, an ancilla ion must be read out while neighboring data ions preserve their populations and coherence. Standard laser-induced-fluorescence readout is intrinsically dissipative. The measured ion repeatedly absorbs resonant 493 nm photons and spontaneously emits fluorescence. Some of this light propagates toward nearby ions and can be reabsorbed, producing population transfer, dephasing, leakage, and motional heating. Even if the detection beam is tightly focused on the ancilla, fluorescence is emitted over a broad range of directions and cannot be eliminated through beam focusing alone.

This form of crosstalk can force a QCCD processor to move an ancilla to a distant measurement region before readout. Greater separation reduces the probability that a fluorescence photon reaches a data ion, while physical barriers and dedicated optical collection can provide further protection. The cost is additional shuttling, separation, recombination, and recoiling. These operations lengthen the error-correction cycle and allow the data qubits to accumulate memory errors. A method that protects nearby ions spectroscopically could permit readout at much smaller separations and reduce this transport overhead.

9.4.0.1 Scale of Fluorescence Reabsorption

The possible size of the problem can be estimated using a two-level resonant-scattering model. For light with wavelength $\lambda = 493$ nm, the ideal resonant scattering cross section is

$$\sigma_0 = \frac{3\lambda^2}{2\pi} \simeq 1.16 \times 10^{-13} \text{ m}^2. \quad (9.8)$$

Consider two barium ions separated by $r = 5 \mu\text{m}$. If the fluorescence is treated as isotropic, the optical power from one emitted photon is distributed over a sphere with area

$4\pi r^2$ at the position of the neighboring ion. The probability that this photon is absorbed by an ideal resonant spectator is then approximately

$$p_\gamma = \frac{\sigma_0}{4\pi r^2} \simeq 3.7 \times 10^{-4}. \quad (9.9)$$

For a scattering rate of $R_{\text{sc}} = 10^7 \text{ s}^{-1}$ and a readout duration of $t_r = 100 \mu\text{s}$, the measured ion emits approximately

$$N_\gamma = R_{\text{sc}} t_r = 10^3 \quad (9.10)$$

photons. Treating the absorption events as independent gives a probability for at least one reabsorption event of

$$P_{\text{cross}} = 1 - \exp(-N_\gamma p_\gamma) = 1 - \exp(-0.37) \simeq 31\%. \quad (9.11)$$

This estimate should be interpreted as a high-end scale calculation rather than a quantitative prediction for the full $^{133}\text{Ba}^+$ system. The expression for σ_0 assumes an ideal closed two-level transition driven by weak resonant light. A real spectator occupies a particular hyperfine and Zeeman state, and its absorption strength depends on the relevant Clebsch–Gordan coefficient and photon polarization. Dipole radiation is not perfectly isotropic. The emitted spectrum is broadened by the natural linewidth, the driving fields, and atomic motion. The spectator absorption spectrum is also power broadened and may be shifted by magnetic and optical fields. Branching into other internal states further changes the probability that one absorbed photon produces a lasting qubit error.

These corrections may reduce the numerical value substantially. The estimate nevertheless shows why fluorescence reabsorption cannot be dismissed based only on the small geometric cross section. A single emitted photon has a small probability of reaching and exciting the neighboring ion, but fluorescence readout requires hundreds or thousands of scattering events. The total error can therefore become appreciable at the few-micrometer

spacings used in dense ion registers.

The same estimate also shows the basic geometric scaling. In the limit of small total absorption probability, P_{cross} scales approximately as N_γ/r^2 . Increasing the separation by a factor of ten suppresses the probability by roughly two orders of magnitude, which explains the attraction of distant measurement zones. Spectral suppression offers an alternative because it reduces the effective absorption cross section without requiring a corresponding increase in distance.

9.4.0.2 Autler–Townes Shielding

A possible solution is to apply a strong 650 nm control field to the spectator ion. The control field couples the $5D_{3/2}$ and $6P_{1/2}$ manifolds. The unwanted 493 nm fluorescence couples the $6S_{1/2}$ manifold to the same $6P_{1/2}$ excited state. These transitions form a three-level system in which the strong control field modifies the spectator’s response to the weak, unwanted fluorescence.

When the 650 nm control field is sufficiently strong, it dresses the $6P_{1/2}$ excited state together with the coupled $5D_{3/2}$ state. The original 493 nm absorption resonance is replaced by two dressed resonances separated approximately by the control-field Rabi frequency. This is the Autler–Townes regime. A 493 nm photon centered at the original bare resonance then lies between the dressed absorption peaks and has a much smaller probability of being absorbed.

Depending on the coherence, detuning, and decay rates of the three-level system, destructive quantum interference can produce an additional transparency feature near two-photon resonance. The shielding can therefore contain features of both Autler–Townes splitting and electromagnetically induced transparency. For the present proposal, the important operational quantity is not the name assigned to the line shape. It is the integrated probability that the broadband fluorescence spectrum of the measured ion overlaps the dressed absorption spectrum of the spectator.




Figure 9.2: Concept for spectroscopic shielding of a spectator ion during fluorescence readout. A measured ancilla emits approximately 10^3 photons during a $100\ \mu\text{s}$ readout, giving a high-end reabsorption estimate of approximately 31% for a spectator $5\ \mu\text{m}$ away. A strong 650 nm control field splits the spectator’s 493 nm absorption resonance and reduces its overlap with the fluorescence spectrum. For a $^{133}\text{Ba}^+$ ground-state clock qubit at 5 G, the calculated differential Stark shift can be cancelled near a circular-polarization parameter of $|\mathcal{A}| \simeq 0.29$. The spectra in the center panel are schematic and are not plotted from the numerical optical-Bloch calculation.

The spectral-overlap calculation can be written in a form that separates the emission and absorption processes. Let $S_{\text{em}}(\omega)$ be the normalized spectrum of the 493 nm photons emitted by the measured ion, and let $A_{\text{sp}}(\omega; s_{650})$ be the spectator absorption profile in the presence of a 650 nm field with saturation parameter s_{650} . The relevant overlap is

$$\mathcal{O}(s_{650}) = \int S_{\text{em}}(\omega) A_{\text{sp}}(\omega; s_{650}) d\omega. \quad (9.12)$$

The shielding factor can be defined by normalizing this overlap to the undressed value,

$$\eta_{\text{shield}}(s_{650}) = \frac{\mathcal{O}(s_{650})}{\mathcal{O}(0)}. \quad (9.13)$$

The crosstalk estimate in Eq. (9.11) can then be modified by replacing p_γ with $p_\gamma \eta_{\text{shield}}$. This emphasizes that the dressing beam does not prevent photons from reaching the spectator. Instead, it reduces the probability that an arriving photon is absorbed.

For the present calculation, the measured ion was assumed to scatter with $s_{493} = 1$ and $s_{650} = 1$. These fields determine the emitted 493 nm spectrum. The spectator absorption was then calculated while varying the strength of its additional 650 nm dressing field. Under these assumptions, modest optical power in the dressing beam produces greater than 99% suppression of the spectrally allowed readout crosstalk. This conclusion refers to the reduction in the spectral-overlap contribution. The total spectator error must also include imperfect polarization, residual dressing-beam crosstalk, off-resonant scattering, and any population outside the protected internal-state manifold.

These two figures are retained conditionally so that they will appear automatically when the original calculation files are placed in the thesis directory. The new overview figure summarizes the physical mechanism independently of those numerical plots.

9.4.0.3 Dressing Power and Raman Scattering

The dressing field must be strong enough to move the spectator absorption away from the emitted fluorescence spectrum. A large saturation parameter does not necessarily require a large total optical power because the 650 nm transition is strong and the beam can be tightly focused. For the proposed operating point, a saturation parameter of $s_{650} = 640$ was estimated to require approximately $23 \mu\text{W}$ of optical power for a nominal $10 \mu\text{m}$ beam size.

The beam-size convention must be stated carefully when quoting this number. For a Gaussian beam with $1/e^2$ intensity radius w_0 , the peak intensity is

$$I_0 = \frac{2P}{\pi w_0^2}. \quad (9.14)$$

An intensity of approximately 50 W/cm^2 requires about $20 \mu\text{W}$ when the quoted $10 \mu\text{m}$ beam size is the full $1/e^2$ diameter, corresponding to $w_0 \simeq 5 \mu\text{m}$. If $10 \mu\text{m}$ instead denotes the standard Gaussian waist radius, the required power is approximately $79 \mu\text{W}$. The stated $23 \mu\text{W}$ estimate is therefore consistent with a beam diameter near $10 \mu\text{m}$ rather than a $10 \mu\text{m}$ radius. This convention should be made explicit in the final calculation and figure labels.

One possible error from the dressing beam is off-resonant Raman scattering on the ground-state spectator. For the calculated $s_{650} = 640$ operating point, the estimated Raman scattering rate is

$$\Gamma_{\text{Raman}} \simeq 9 \times 10^{-6} \text{ Hz}. \quad (9.15)$$

During a $100 \mu\text{s}$ readout, the corresponding error is

$$\epsilon_{\text{Raman}} \simeq \Gamma_{\text{Raman}} t_{\text{r}} \simeq 9 \times 10^{-10}. \quad (9.16)$$

Within this model, Raman scattering from the dressing beam is negligible compared with the unshielded fluorescence-reabsorption estimate. This very small result should still be tested against a complete multilevel calculation that includes both fine-structure manifolds, all relevant hyperfine states, the actual polarization at the ion, and the spectral content of the dressing laser. Coherent ac Stark shifts and their fluctuations may become important well before the Raman-scattering probability does.

The dressing beam can also affect the measured ancilla if it is not sufficiently localized. The measured ion already requires 650 nm repumping during ordinary fluorescence detection, but a much stronger or differently polarized field can alter its steady-state fluorescence rate and emission spectrum. Spatial addressing must therefore ensure that the

additional shielding field reaches the spectator with the intended intensity while producing a known and acceptable change at the measured ion. This is an important systems-level constraint because the same optical field influences both the absorber and the source spectrum if it overlaps both ions.

9.4.0.4 Magic-Helicity Protection of the Ground-State Qubit

Even when spontaneous Raman scattering is negligible, a strong dressing field produces an ac Stark shift. For a spectator $^{133}\text{Ba}^+$ ground-state clock qubit at $B = 5\text{ G}$ and $s_{650} = 640$, corresponding to an estimated intensity of approximately

$$I_{650} \simeq 50\text{ W/cm}^2, \quad (9.17)$$

the calculated differential clock shift is

$$\delta\nu_g \simeq -0.0325 + 0.1135\mathcal{A} \text{ Hz}. \quad (9.18)$$

The parameter $\mathcal{A}$ describes the circular polarization of the field,

$$\mathcal{A} = \frac{I_{\sigma^+} - I_{\sigma^-}}{I_{\sigma^+} + I_{\sigma^-}}. \quad (9.19)$$

The constant term in Eq. (9.18) represents the polarization-independent part of the differential shift within the model. The term proportional to $\mathcal{A}$ is controlled by the helicity of the dressing field relative to the quantization axis. Because these terms have opposite signs for an appropriate polarization, the differential clock shift can be cancelled without turning off the shielding interaction.

At $B = 5\text{ G}$, the cancellation occurs near

$$|\mathcal{A}_{\text{magic}}| \simeq 0.29. \quad (9.20)$$

For the sign convention in Eq. (9.19), $\mathcal{A} = 0.29$ corresponds to intensity fractions

$$f_{\sigma^+} = \frac{1 + \mathcal{A}}{2} \simeq 0.65, \quad f_{\sigma^-} = \frac{1 - \mathcal{A}}{2} \simeq 0.35. \quad (9.21)$$

The resulting field is therefore approximately 65% σ^+ and 35% σ^- . Although this condition can be set experimentally by rotating a wave plate, it is more accurately described as a magic-helicity or magic-polarization condition than as a literal magic angle. The relevant quantity is the circular Stokes component at the ion rather than the mechanical angle of a particular optical element.

Even without cancellation, Eq. (9.18) predicts a shift with magnitude below approximately 0.2 Hz for all allowed values of $\mathcal{A}$. During one 100 μ s readout, a 0.2 Hz differential shift produces a phase below approximately 1.3×10^{-4} rad. This is small for one measurement, but a coherent phase can accumulate over many repeated QEC cycles. Operating near the magic helicity reduces this systematic accumulation and makes the remaining error depend primarily on polarization and intensity fluctuations.

The cancellation must be calibrated at the ion. Vacuum windows, fibers, waveguides, and integrated optics can all change the polarization between the external diagnostic point and the trapped ion. Small birefringence drifts may change $\mathcal{A}$ even when the wave-plate angle is fixed. A practical implementation should measure the clock shift as the polarization is scanned, identify the zero crossing directly, and monitor the cancellation over the duration of the experiment.

9.4.0.5 Experimental Tests and Remaining Error Channels

The first experimental step is to characterize the dressed spectator spectrum. A weak 493 nm probe can be scanned across the $6S_{1/2}$ to $6P_{1/2}$ transition while the 650 nm control intensity and detuning are varied. This measurement should resolve the transition from a single broadened resonance to a split or transparent spectrum and determine the suppression at the frequency occupied by the measured ion's fluorescence.

The second step is a direct two-ion crosstalk measurement. One ion should undergo a normal fluorescence-detection pulse while the internal state of the neighboring ion is analyzed before and after the pulse. Repeating the measurement as a function of ion separation, readout duration, photon-scattering rate, spectator state, and photon polarization would replace the isotropic two-level estimate with an empirical error probability. Applying the 650 nm dressing field to the spectator would then provide a direct measurement of the shielding factor.

The fluorescence spectrum of the measured ion and the absorption spectrum of the spectator should be measured under the same conditions used in the crosstalk experiment. At $s_{493} = 1$, the emitted light is not perfectly described by a single natural-linewidth Lorentzian. Power broadening and driven-atom fluorescence can place spectral weight away from line center, where it may overlap one of the dressed spectator resonances. The final shielding prediction must therefore integrate the measured or calculated emitter spectrum against the full dressed absorption profile.

The third step is to measure the effect of the dressing field on the spectator clock qubit. Ramsey spectroscopy can determine the differential Stark shift, while repeated dressing pulses inserted into a spin-echo sequence can distinguish a stable coherent shift from irreversible dephasing. Scanning the circular polarization should reveal the predicted zero crossing near $|\mathcal{A}| = 0.29$. The sensitivity to polarization drift, control-field detuning, and intensity noise should then be measured around this operating point.

Additional errors may arise from imperfect state preparation, leakage of the spectator into the $D_{3/2}$ manifold, coherent population trapping, dressing-beam scattering from the trap, and motional excitation. The shielding may also be state dependent across the two ground-state clock levels. A useful protocol must protect an arbitrary superposition, not only prevent population loss from one basis state. Both the diagonal population error and the off-diagonal coherence must therefore be characterized.

The final test is integration into a mid-circuit measurement sequence. An ancilla should be entangled with one or more data ions, measured with fluorescence while the spectators

are dressed, and then reset for another round. The data-qubit process fidelity should be compared with and without shielding. The comparison should include the optical error, the additional time required to turn the dressing field on and off, and any reduction in ancilla photon collection or detection fidelity. It should also be compared with the alternative strategy of shuttling the ancilla to a distant readout zone.

If successful, Autler–Townes shielding would provide a new way to trade optical control power for reduced spatial separation during measurement. It would not eliminate the need for careful beam delivery or integrated measurement zones, but it could substantially relax the distance required between an ancilla and stored data. The current calculations suggest that greater than 99% suppression of spectrally allowed reabsorption may be possible with modest total 650 nm power, while Raman scattering and differential Stark shifts remain extremely small. Demonstrating these predictions in a two-ion experiment would establish whether coherent spectral dressing can become a practical tool for fast, low-crosstalk syndrome extraction in a fault-tolerant $^{133}\text{Ba}^+$ processor.

9.5 Microwave-Dressed Magnetic-Field-Insensitive Qubits in Barium

Magnetic-field noise remains an important source of dephasing in trapped-ion quantum processors. A nominal clock transition can eliminate the leading linear Zeeman shift at one magnetic field, but the transition generally retains curvature with respect to the field. This residual curvature converts slow magnetic-field drift, spatial field variation between trapping zones, and technical current noise into qubit-frequency fluctuations. The problem becomes more important as a processor grows because a single calibration must remain valid across many locations and throughout long error-correction sequences. A broad operating region with a nearly constant transition frequency would therefore be more useful than a mathematically exact cancellation at only one field value.

Continuous microwave dressing offers a way to engineer this magnetic response. An off-

resonant microwave field shifts the two qubit states through their couplings to neighboring hyperfine-Zeeman levels. Because the detuning of each auxiliary transition depends on magnetic field, the associated ac Zeeman shift is also field dependent. A suitable combination of dressing frequency, microwave amplitude, and polarization can therefore add controlled slope and curvature to the bare qubit transition. The goal is to cancel the unwanted magnetic dependence while preserving a qubit that can still be prepared, controlled, and measured with high fidelity.

The present calculation is motivated by microwave dressing of the states

$$|0\rangle = |6S_{1/2}, F = 1, m_F = -1\rangle, \quad |1\rangle = |6S_{1/2}, F = 2, m_F = +1\rangle \quad (9.22)$$

in $^{137}\text{Ba}^+$. These two states have nearly equal linear Zeeman slopes. Their difference frequency consequently has a shallow magnetic-field dependence and a nonzero-field minimum. They are not the conventional $m_F = 0$ clock states. Instead, their usefulness comes from the near cancellation of the two individual linear Zeeman shifts. This state pair also has the same hyperfine topology as the atomic system used in the original microwave-dressing proposal, making it a natural starting point for translating the method to barium.

An isotope distinction is essential. The exploratory notes and numerical plots supplied for this proposal describe $^{137}\text{Ba}^+$, which has nuclear spin $I = 3/2$ and ground-state manifolds with $F = 1$ and $F = 2$. The surrounding processor architecture in this thesis uses $^{133}\text{Ba}^+$, which has a different nuclear spin and a different ground-state hyperfine structure. The specific state pair in Eq. (9.22), the three auxiliary detunings, and their angular factors cannot be copied directly into a $^{133}\text{Ba}^+$ calculation. The general strategy remains applicable, but the $^{133}\text{Ba}^+$ Hamiltonian and available dressing pathways must be constructed independently. The $^{137}\text{Ba}^+$ calculation should therefore be viewed as a barium case study and as a test of the numerical method, rather than as a completed calculation for the $^{133}\text{Ba}^+$ processor.

9.5.0.1 Bare Magnetic-Field Dependence

For a hyperfine state $|F, m_F\rangle$ in the low-field regime, the first-order Zeeman shift is determined by the hyperfine Landé factor. Including both electronic and nuclear contributions, this factor is

$$g_F = g_J \frac{F(F+1) + J(J+1) - I(I+1)}{2F(F+1)} + g_I \frac{\mu_N}{\mu_B} \frac{F(F+1) + I(I+1) - J(J+1)}{2F(F+1)}. \quad (9.23)$$

For the two states in Eq. (9.22), it is convenient to collect the first-order coefficients into μ_1 and μ_2 and the leading quadratic coefficient into β . The bare state energies can then be written locally as

$$E_0(B) = \hbar (\mu_1 B - 3\beta B^2), \quad (9.24)$$

$$E_1(B) = \hbar (\omega_{\text{hfs}} + \mu_2 B + 3\beta B^2), \quad (9.25)$$

where ω_{hfs} is the zero-field hyperfine splitting in angular-frequency units. The corresponding bare qubit frequency is

$$\omega_q^{(0)}(B) = \omega_{\text{hfs}} + (\mu_2 - \mu_1)B + 6\beta B^2. \quad (9.26)$$

The first derivative vanishes at

$$B_0 = \frac{\mu_1 - \mu_2}{12\beta}. \quad (9.27)$$

The exploratory $^{137}\text{Ba}^+$ calculation places this bare minimum near $B_0 \simeq 2.35$ G. At this point the transition is insensitive to small field fluctuations to first order, but it still

has the quadratic sensitivity 12β . Moving away from B_0 reintroduces a linear sensitivity. Microwave dressing is intended to reduce both effects over a wider interval of magnetic field.

The numerical notes reveal an important units issue that should be corrected before further optimization. Values near 4.4×10^6 were labeled as hertz per gauss, but their magnitude is consistent with angular frequency per gauss. Dividing by 2π gives slopes near 7.0×10^5 Hz/G. Mixing hertz and radians per second changes the inferred detunings, Rabi frequencies, and derivatives by factors of 2π , and it can move or even create apparent optimization points. A clean implementation should use angular-frequency units throughout the Hamiltonian and convert to hertz only when producing final plots and tables.

9.5.0.2 Off-Resonant Microwave Dressing

A microwave magnetic field polarized perpendicular to the static quantization field contains σ^+ and σ^- components. These components couple each qubit state to nearby states in the other hyperfine manifold, as shown in Fig. 9.3. The transition strengths are not equal because the relevant Clebsch–Gordan coefficients differ. The microwave polarization at the ion also need not contain equal σ^+ and σ^- amplitudes. Both effects must be included explicitly.




Figure 9.3: Microwave dressing of magnetic-field-sensitive clock states in barium. The left panel shows the $^{137}\text{Ba}^+$ ground-state topology used in the exploratory calculation. An off-resonant microwave field couples the qubit states to neighboring Zeeman levels through several detunings. The center panel illustrates the desired result. The undressed qubit has a first-order-insensitive point near 2.35 G but retains quadratic curvature. Proper dressing can, in principle, create a broader plateau and make the operating frequency insensitive to both magnetic-field and microwave-amplitude fluctuations. The plotted frequency curves are schematic and are not fitted to a validated barium solution. The right panel distinguishes the results already supported by the exploratory notes from the preliminary candidate range and the still-unresolved double-magic condition.

If the dressing field has angular frequency ω_d , its detuning from the zero-field hyperfine splitting can be defined as

$$\Delta_d = \omega_d - \omega_{\text{hfs}}. \quad (9.28)$$

Using the same low-field expansion as Eqs. (9.24) and (9.25), three relevant signed detunings can be written as

$$\Delta_1(B) = \Delta_d + (\mu_1 + 2\mu_2)B - 3\beta B^2, \quad (9.29)$$

$$\Delta_2(B) = \Delta_d + \mu_1 B - 7\beta B^2, \quad (9.30)$$

$$\Delta_3(B) = \Delta_d - \mu_2 B - 7\beta B^2. \quad (9.31)$$

In the far-detuned limit, the shift produced by one auxiliary two-level coupling has the familiar form

$$\delta\omega_{\text{ac}} \simeq \frac{|\Omega_{\text{eff}}|^2}{4\Delta}, \quad (9.32)$$

up to the sign determined by the signed-detuning and Hamiltonian convention. The exact dressed eigenvalues of an isolated two-level block are

$$\lambda_{\pm} = \frac{1}{2} \left(\Delta \pm \sqrt{\Delta^2 + |\Omega_{\text{eff}}|^2} \right). \quad (9.33)$$

For the multilevel $^{137}\text{Ba}^+$ problem, the differential dressing shift is a weighted sum of contributions from the three pathways. Under one commonly used definition of a state-independent dressing Rabi frequency, the perturbative result takes the form

$$\begin{aligned} \omega_q(B, \Omega_d) = & \omega_{\text{hfs}} + (\mu_2 - \mu_1)B + 6\beta B^2 \\ & - \Omega_d^2 \left[\frac{3}{\Delta_1(B)} + \frac{1}{2\Delta_2(B)} + \frac{3}{2\Delta_3(B)} \right]. \end{aligned} \quad (9.34)$$

Equation (9.34) is useful for building intuition, but it is also where implementation errors can enter. The overall sign depends on the rotating-frame convention, and each denominator must retain its signed detuning. Reversing the definition of one detuning

changes the apparent sign of that term without changing the physical Hamiltonian. The exploratory work initially failed to reproduce the reference atomic calculation and was later brought into qualitative agreement after correcting the treatment of the signed detunings and the relative σ^+ and σ^- coupling strengths. This history makes direct matrix diagonalization the safer basis for the final barium calculation.

The perturbative expression also becomes unreliable near any field at which a denominator approaches zero. The poles visible in several exploratory scans are not useful high-sensitivity operating points. They indicate that the far-detuned expansion has left its domain of validity. Near these regions, the system must be treated as a coupled multilevel Hamiltonian, and a candidate operating point must be rejected if the dressed qubit contains excessive admixture of an auxiliary state or becomes vulnerable to leakage.

9.5.0.3 Meaning of a Magic and Double-Magic Operating Point

Several different robustness conditions appear in the exploratory notes, and they should be distinguished carefully. A first-order magnetic sweet spot satisfies

$$\left. \frac{\partial \omega_q}{\partial B} \right|_{B_*, \Omega_*} = 0. \quad (9.35)$$

Suppressing the residual curvature additionally requires

$$\left. \frac{\partial^2 \omega_q}{\partial B^2} \right|_{B_*, \Omega_*} \simeq 0. \quad (9.36)$$

Together, these two conditions create a broad magnetic-field plateau. They are sufficient to suppress the leading first- and second-order Zeeman sensitivity, but they do not guarantee immunity to noise in the microwave source. A dressing-amplitude-insensitive point also satisfies

$$\left. \frac{\partial \omega_q}{\partial \Omega_d} \right|_{B_*, \Omega_*} = 0. \quad (9.37)$$

The phrase *double magic* is most useful when it refers to simultaneous first-order insensitivity to two independent noise coordinates, here magnetic field and dressing amplitude. In this usage, Eqs. (9.35) and (9.37) are the defining conditions, while Eq. (9.36) is an additional and highly desirable flattening condition. A thesis discussion should state explicitly which derivatives vanish rather than relying on the phrase alone.

One microwave tone does not necessarily provide enough independent controls to satisfy all three conditions at a conveniently chosen field. The available control parameters include B , the common detuning Δ_d , the total Rabi frequency, and the amplitude ratio and relative phase of the σ^+ and σ^- components. A second dressing tone would add further degrees of freedom if the single-tone solution is either absent or impractical. The polarization imbalance is particularly important because the three pathways have different angular matrix elements. Treating the field as perfectly linear can hide a useful control parameter or produce a predicted cancellation that is not present for the measured microwave polarization.

9.5.0.4 Interpretation of the Preliminary Barium Scans

The exploratory calculation successfully reproduced the qualitative bare $^{137}\text{Ba}^+$ frequency minimum near 2.35 G and demonstrated that microwave dressing can strongly modify both the slope and curvature of the transition. It did not yet establish a validated double-magic operating point for barium.

One early trial used values reported as

$$\Omega_d = 4.51497 \times 10^5 \text{ rad/s}, \quad \Delta_d = 2.82467 \times 10^7 \text{ rad/s}, \quad (9.38)$$

which correspond to approximately $\Omega_d/2\pi = 71.9 \text{ kHz}$ and $\Delta_d/2\pi = 4.50 \text{ MHz}$. These parameters altered the magnetic response but did not yield the desired simultaneous cancellation near the initially considered field around 3.5 G.

Later coarse searches returned candidate regions rather than exact solutions. One

tolerance-based search selected $\Omega_d/2\pi \simeq 27.5$ kHz and $B \simeq 2.712$ G. The reported first and second derivatives at that point were still nonzero, so the result should be interpreted as the best grid point under loose numerical thresholds rather than as a magic solution. A separate scan at a reported detuning of approximately 0.941 MHz suggested an operating region near $\Omega_d/2\pi \simeq 37.6$ kHz and $B \simeq 2.57$ G. Several plot annotations labeled these magnetic fields in tesla even though the axes and numerical scale clearly corresponded to gauss. That labeling error should be removed in the reconstructed analysis.

Taken together, the scans suggest a preliminary search window of approximately

$$2.57 \text{ G} \leq B \leq 2.71 \text{ G}, \quad 27 \text{ kHz} \leq \frac{\Omega_d}{2\pi} \leq 38 \text{ kHz}. \quad (9.39)$$

This window is useful for directing the next calculation, but it is not a predicted operating point. The attached notebook contains only one empty code cell and no stored outputs, so the code that generated the exported plots is not available for inspection or reproducibility testing. The calculation must be rebuilt before any parameter from the scans is used experimentally.

9.5.0.5 Reconstructing and Validating the Calculation

The first priority is to create one explicit Hamiltonian with a documented basis. For $^{137}\text{Ba}^+$, the most reliable model begins with the complete $6S_{1/2}$ hyperfine-Zeeman Hamiltonian rather than a quadratic expansion. The static Hamiltonian can be written as

$$H_0 = A \mathbf{I} \cdot \mathbf{J} + \mu_B g_J J_z B - \mu_N g_I I_z B, \quad (9.40)$$

where A is the ground-state magnetic-dipole hyperfine constant. Diagonalizing Eq. (9.40) automatically includes the correct nonlinear Zeeman response and avoids ambiguity in the low-field coefficients.

The time-dependent magnetic interaction should then be expressed in terms of the

measured complex microwave field at the ion,

$$H_d(t) = -\boldsymbol{\mu} \cdot \text{Re}[\mathbf{B}_d e^{-i\omega_d t}] . \quad (9.41)$$

Writing $\mathbf{B}_d$ in the spherical basis exposes the σ^+ , π , and σ^- components and their phases. The rotating-wave Hamiltonian can then be diagonalized at each value of B , ω_d , and microwave amplitude. The dressed eigenstates should be followed continuously across parameter space by their overlap with the bare qubit states. This prevents an eigenvalue-sorting routine from jumping between branches near avoided crossings.

Before searching barium parameters, the reconstructed code should reproduce the known rubidium example using the same conventions. This is not merely a qualitative comparison of curve shapes. The benchmark should verify the bare magic field, the dressed plateau, the dependence on microwave amplitude, and the published relative σ^+ and σ^- coupling strengths. Unit tests should compare direct diagonalization with second-order perturbation theory far from resonance and confirm that the disagreement grows only where the perturbative assumptions fail.

The barium search should then minimize a dimensionless cost function that includes all desired sensitivities and physical penalties. A useful form is

$$\begin{aligned} \mathcal{C} = & w_1 \left(\frac{\partial_B \omega_q}{S_B} \right)^2 + w_2 \left(\frac{\partial_B^2 \omega_q}{C_B} \right)^2 + w_3 \left(\frac{\partial_\Omega \omega_q}{S_\Omega} \right)^2 \\ & + w_4 P_{\text{leak}} + w_5 P_{\text{near resonance}} , \end{aligned} \quad (9.42)$$

where S_B , C_B , and S_Ω are scaling constants that place unlike derivatives on comparable numerical scales. The last two terms penalize auxiliary-state admixture and operation too close to a dressing resonance. A candidate should be refined with a root-finding or constrained-optimization method after the coarse global scan. Its derivatives should be recomputed with several finite-difference step sizes or automatic differentiation to ensure

that a numerical artifact is not mistaken for a cancellation.

For $^{133}\text{Ba}^+$, the same workflow should be repeated from the full isotope-specific Hamiltonian. The smaller hyperfine manifold changes the number and strength of the available microwave pathways. If one tone cannot provide enough differential curvature while remaining far from resonance, two microwave tones or a combination of microwave and radio-frequency dressing could be considered. The expanded control space is useful only if the additional amplitude and phase noise does not exceed the magnetic noise being removed.

9.5.0.6 Experimental Demonstration

The first experimental measurement should be a dressed Ramsey map. The bare qubit frequency is measured as a function of magnetic field to locate the undressed minimum and calibrate the field axis. The dressing field is then applied continuously, and the measurement is repeated while its frequency and amplitude are scanned. This directly produces $\omega_q(B, \Omega_d)$ and allows the first and second magnetic derivatives to be inferred from local fits.

The microwave-amplitude sensitivity should be tested independently. At each candidate field, the dressing power can be modulated by a known small amount while the qubit phase is measured. A true amplitude sweet spot produces no first-order phase modulation. The measurement should be repeated for the actual microwave polarization at the ion because cable phase, impedance mismatch, and the electrode geometry can change the ratio of the two circular components.

Coherence measurements must accompany spectroscopy. A flat frequency curve is useful only if the continuous drive does not introduce larger errors through phase noise, amplitude noise, or leakage. Ramsey and spin-echo measurements with the dressing field on should be compared with the undressed coherence time. Controlled noise can then be injected into the magnetic coils and microwave amplitude to confirm that the observed

improvement agrees with the predicted filter response.

Continuous dressing also changes how quantum gates are implemented. Turning the field on and off abruptly can project a dressed eigenstate onto several bare states. The dressing amplitude should therefore be ramped adiabatically or left on during computation. In the latter case, single- and two-qubit gates must be calibrated in the dressed basis. Their drive frequencies, phases, and selection rules may differ from those of the bare qubit. The relevant performance metric is the complete algorithmic error with dressing present, not the memory coherence time in isolation.

9.5.0.7 Role in a Fault-Tolerant Processor

If the technique can be transferred to $^{133}\text{Ba}^+$, microwave dressing could reduce spatial variation of the qubit frequency across a segmented trap and suppress slow magnetic drift during syndrome extraction. This would relax the requirement that every trapping zone experience an identical static magnetic field. It could also reduce the frequency recalibration required after ion transport and make idle qubits less sensitive to current noise while other zones perform gates or measurement.

The present work establishes the physical mechanism and identifies a plausible $^{137}\text{Ba}^+$ search region, but it does not yet establish a deployable operating point. The immediate milestone is a reproducible exact-diagonalization model that resolves the unit, sign, and polarization conventions and reproduces the reference atomic calculation. The next milestone is an experimentally accessible barium solution that is simultaneously insensitive to magnetic field and dressing amplitude while remaining well separated from all resonances. The final milestone is operation of high-fidelity gates and long QEC sequences in the continuously dressed basis. Reaching all three milestones would turn microwave dressing from a spectroscopy technique into a practical method for suppressing memory errors in a large trapped-ion quantum computer.

REFERENCES

- Aliferis, P., Gottesman, D., and Preskill, J. (2006). Quantum accuracy threshold for concatenated distance-three codes. *Quantum Information and Computation*, 6:97–165. [19](#)
- Allcock, D. T. C., Campbell, W. C., Chiaverini, J., Chuang, I. L., Hudson, E. R., Moore, I. D., Ransford, A., Roman, C., Sage, J. M., and Wineland, D. J. (2021). *omg* blueprint for trapped ion quantum computing with metastable states. *Applied Physics Letters*, 119:214002. [159](#), [161](#)
- An, F. A., Ransford, A., Schaffer, A., Sletten, L. R., Gaebler, J., Hostetter, J., and Vittorini, G. (2022a). High fidelity state preparation and measurement of ion hyperfine qubits with $I > 1/2$. *Physical Review Letters*, 129:130501. [20](#)
- An, F. A., Ransford, A., Schaffer, A., Sletten, L. R., Gaebler, J., Hostetter, J., and Vittorini, G. (2022b). High fidelity state preparation and measurement of ion hyperfine qubits with $I > \frac{1}{2}$. *Phys. Rev. Lett.*, 129:130501. [197](#), [207](#)
- Armstrong, D. J., Wood, R. P., and Greene, C. H. (1993). Photoionization of the $5d6p\ ^3d_1$ state of barium. *Phys. Rev. A*, 47:1981–1988. [179](#), [180](#)
- Aspect, A., Dalibard, J., and Roger, G. (1982). Experimental test of bell’s inequalities using time-varying analyzers. *Physical Review Letters*, 49(25):1804–1807. [14](#)
- Aspuru-Guzik, A., Dutoi, A. D., Love, P. J., and Head-Gordon, M. (2005). Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707. [12](#)
- Aymar, M., Greene, C. H., and Luc-Koenig, E. (1996). Multichannel Rydberg spectroscopy of complex atoms. *Reviews of Modern Physics*, 68(4):1015–1123. [180](#)
- Ballance, C. J., Harty, T. P., Linke, N. M., Sepiol, M. A., and Lucas, D. M. (2016). High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117:060504. [171](#)
- Berkeland, D. J. and Boshier, M. G. (2002). Destabilization of dark states and optical spectroscopy in zeeman-degenerate atomic systems. *Physical Review A*, 65(3):033413. [67](#)
- Berkeland, D. J., Miller, J. D., Bergquist, J. C., Itano, W. M., and Wineland, D. J. (1998). Minimization of ion micromotion in a paul trap. *Journal of Applied Physics*, 83(10):5025–5033. Accessed: Feb. 4, 2026. [41](#)
- Boguslawski, M. J., Wall, Z. J., Vizvary, S. R., Moore, I. D., Bareian, M., Allcock, D. T. C., Wineland, D. J., Hudson, E. R., and Campbell, W. C. (2023). Raman scattering errors in stimulated-raman-induced logic gates in $^{133}\text{ba}^+$. *Phys. Rev. Lett.*, 131:063001. [171](#), [172](#), [180](#)

- Bouwmeester, D., Pan, J.-W., Mattle, K., Eibl, M., Weinfurter, H., and Zeilinger, A. (1997). Experimental quantum teleportation. *Nature*, 390:575–579. [14](#)
- Budker, D., Kimball, D. F., and DeMille, D. P. (2008). *Atomic Physics: An Exploration through Problems and Solutions*. Oxford University Press, Oxford, 2 edition. [75](#)
- Chicireanu, R., Nelson, K. D., Olmschenk, S., Lundblad, N., Derevianko, A., and Porto, J. V. (2011). Differential light-shift cancellation in a magnetic-field-insensitive transition of ^{87}Rb . *Phys. Rev. Lett.*, 106:063002. [168](#)
- Choi, J., Lee, E., Yum, D., An, K., and Kim, J. (2024). Direct measurement of isotope shifts in the barium $6s^2\ ^1s_0$ – $5d6p\ ^3d_1^\circ$ 413-nm electric-quadrupole transition. *Phys. Rev. A*, 110(3):032812. [190](#)
- Choi, J. M. and Cho, D. (2007). Elimination of inhomogeneous broadening for a ground-state hyperfine transition in an optical trap. *Phys. Rev. Lett.*, 80(1):012037. [168](#)
- Christensen, J. E. (2020). *High-fidelity operation of a radioactive trapped-ion qubit, $^{133}\text{Ba}^+$* . PhD thesis, University of California, Los Angeles. [82](#), [161](#)
- Deutsch, D. (1985). Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117. [10](#)
- Evans, R., Peterson, L., and Cafiero, M. (2018). Evaluation of hybrid and pure dft methods for the binding of novel ligands in the tyrosine hydroxylase enzyme. *Computational and Theoretical Chemistry*, 1140:145–151. [1](#)
- Fano, U. and Rau, A. R. P. (1986). *Atomic Collisions and Spectra*. Academic Press, Orlando. [191](#)
- Freedman, S. J. and Clauser, J. F. (1972). Experimental test of local hidden-variable theories. *Physical Review Letters*, 28(14):938–941. [14](#)
- Gidney, C. and Ekerå, M. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5:433. [16](#)
- Gottesman, D. (1998). The heisenberg representation of quantum computers. *arXiv preprint quant-ph/9807006*. [16](#)
- Greene, C. H. and Aymar, M. (1991). Spin-orbit effects in the heavy alkaline-earth atoms. *Physical Review A*, 44(3):1773–1790. [180](#)
- Hänsch, T. W. and Schawlow, A. L. (1975). Cooling of gases by laser radiation. *Optics Communications*, 13:68–69. [54](#)
- Harlander, M., Brownnutt, M., Hänsel, W., and Blatt, R. (2010). Trapped-ion probing of light-induced charging effects on dielectrics. *New Journal of Physics*, 12(9):093035. [180](#), [187](#)

- Harty, T. P., Allcock, D. T. C., Ballance, C. J., Guidoni, L., Janacek, H. A., Linke, N. M., Stacey, D. N., and Lucas, D. M. (2014). High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Physical Review Letters*, 113:220501. 66
- Haze, S., Ohno, T., Toyoda, K., and Urabe, S. (2011). Measurement of the coherence time of the ground-state zeeman states in 40Ca^+ . *Applied Physics B*, 105:761–765. 66
- Hendricks, R. J., Grant, E., Muir, B., Webster, S. C., Nevado, M., Camp, J., Warring, U., Hornekær, L., Drewsen, M., Millen, J., Hinds, E. A., and Tarbutt, M. R. (2007). An all-optical ion-loading technique for scalable microtrap architectures. *Applied Physics B*, 88(4):507–513. 82
- Hogle, C. W., Burch, A. D., Sterk, J. D., Chow, M. N. H., Ivory, M., Lobser, D. S., Maunz, P., Wall, J. V. D., Yale, C. G., Clark, S. M., Stick, D., and Revelle, M. C. (2024). Precise micromotion compensation of a tilted ion chain. *Frontiers in Quantum Science and Technology*, 3:1352800. Brief Research Report, Basic Science for Quantum Technologies. 45
- Hu, Q.-Q., Freier, C., Sun, Y., Leykauf, B., Schkolnik, V., Yang, J., Krutzik, M., and Peters, A. (2018). Observation of vector and tensor light shifts in ^{87}Rb using near-resonant, stimulated raman spectroscopy. *Phys. Rev. A*, 97:013424. 172
- Hucul, D. (2015). *Modular Entanglement of Atomic Qubits Using Photons and Phonons*. PhD thesis, University of Maryland. 85
- Hughes, A. C., Srinivas, R., Löschnauer, C. M., Knaack, H. M., Matt, R., Ballance, C. J., Malinowski, M., Harty, T. P., and Sutherland, R. T. (2025). Trapped-ion two-qubit gates with $> 99.99\%$ fidelity without ground-state cooling. 20
- Johansen, J., Estey, B., Rowe, M., and Ransford, A. (2022). Fast loading of a trapped ion quantum computer using a 2d magneto-optical trap. In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 299–303. 184
- Johnson, K. (2020). *Trapped Barium Ion Quantum Information Processing*. PhD thesis, University of California, Los Angeles. 85
- Johnson, K. G. (2016). *Experiments with Trapped Ions and Ultrafast Laser Pulses*. Dissertation, University of Maryland, College Park. 41
- Kim, H., Han, H. S., and Cho, D. (2013). Magic polarization for optical trapping of atoms without stark-induced dephasing. *Phys. Rev. Lett.*, 111:243004. 168
- Kim, M., Kim, K., Hong, J., Lee, H., Moon, Y., Lee, W. C., Kim, S., Ha, T., Sim, J.-Y., and Lee, M. (2022). Ion trap with gold-plated alumina: Substrate and surface characterization. *AIP Advances*, 12(11):115006. 85
- Kovacic, I., Rand, R. H., and Sah, S. M. (2018). Mathieu’s equation and its generalizations: Overview of stability charts and their features. *Applied Mechanics Reviews*, 70(2):020802. 26

- Landau, L. D. and Lifshitz, E. M. (1976). *Mechanics*, volume 1 of *Course of Theoretical Physics*. Butterworth-Heinemann, Oxford, 3 edition. 23
- Leibfried, D., Blatt, R., Monroe, C., and Wineland, D. (2003). Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75(1):281–324. 25
- Leschhorn, G., Hasegawa, T., and Schaetz, T. (2012). "efficient photo-ionization for barium ion trapping using a dipole-allowed resonant two-photon transition". *Appl Phys B*, 108:159–165. 179, 181
- Low, P. J., White, B., and Senko, C. (2025). Control and readout of a 13-level trapped ion qudit. *npj Quantum Information*, 11:85. 64
- Lucas, D. M., Ramos, A., Home, J. P., McDonnell, M. J., Nakayama, S., Stacey, J.-P., Webster, S. C., Stacey, D. N., and Steane, A. M. (2004). Isotope-selective photoionization for calcium ion trapping. *Phys. Rev. A*, 69(1):012711. 180, 187
- M. Aymar, P. C. and Himdy, A. E. (1983). Multichannel Quantum Defect Analysis of the Autoionized Even-parity Spectrum of Neutral Barium below the 5d2D3/2 Limit . *Phys. Scr.*, 193(27). 180
- Menon, A., Strauss, M., Tomaras, G., Jeanette, L., Sheffield, A. X., Valdez, D., Xie, Y., So, V., De Luo, H., Duraisamy Suganthi, M., Dugan, M., Bado, P., Linke, N. M., Pagano, G., and Zhuravel, R. (2026). Monolithic segmented 3d ion trap for quantum technology applications. *arXiv preprint arXiv:2603.16048*. 85
- Monroe, C., Meekhof, D. M., King, B. E., Itano, W. M., Wineland, D. J., and Gould, P. (1995a). Demonstration of a fundamental quantum logic gate. *Physical Review Letters*, 75(25):4714–4717. 116
- Monroe, C., Meekhof, D. M., King, B. E., Jefferts, S. R., Itano, W. M., Wineland, D. J., and Gould, P. (1995b). Resolved-sideband raman cooling of a bound atom to the 3d zero-point energy. *Physical Review Letters*, 75(22):4011–4014. 116
- Moore, I. D. (2023). *Easy on the Ions: Photon Scattering Errors from Far-Detuned Raman Beams in Trapped-Ion Qubits*. PhD thesis, University of Oregon. 130
- Moore, I. D., Campbell, W. C., Hudson, E. R., Boguslawski, M. J., Wineland, D. J., and Allcock, D. T. C. (2023a). Photon scattering errors during stimulated raman transitions in trapped-ion qubits. *Physical Review A*, 107:032413. 130, 139
- Moore, I. D., Campbell, W. C., Hudson, E. R., Boguslawski, M. J., Wineland, D. J., and Allcock, D. T. C. (2023b). Photon scattering errors during stimulated raman transitions in trapped-ion qubits. *Phys. Rev. A*, 107:032413. 171, 172
- N. Greenberg, B. M. White, P. J. L. and Senko, C. (2024). "trapping ba^+ with seven – fold enhanced efficiency utilizing an autoionizing resonance". *Quantum Sci. Technol*, 9 : 035023. 178, 179, 185, 187

- Nadlinger, D. P., Drmota, P., Main, D., Nichol, B. C., Araneda, G., Srinivas, R., Stephenson, L. J., Ballance, C. J., and Lucas, D. M. (2021). Micromotion minimisation by synchronous detection of parametrically excited motion. *arXiv preprint*. 11 pages, 6 figures. [46](#)
- Nafe, J. E., Nelson, E. B., and Rabi, I. I. (1947). The hyperfine structure of atomic hydrogen and deuterium. *Physical Review*, 71:914–915. [54](#)
- Ozeri, R., Langer, C., Jost, J. D., Knill, E., Brown, K. R., Britton, J., Leibfried, D., Ozeri, R., and Wineland, D. J. (2007). Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Physical Review A*, 75(4):042329. [134](#), [135](#), [149](#)
- Paul, W. (1990). Electromagnetic traps for charged and neutral particles. *Rev. Mod. Phys.*, 62:531–540. [22](#)
- Preskill, J. (1998). Lecture notes for physics 219/computer science 219: Quantum computation. Available from the Caltech course website. [14](#)
- Qiao, M., Wang, Y., Cai, Z., Du, B., Wang, P., Luan, C., Chen, W., Noh, H.-R., and Kim, K. (2021). Double-electromagnetically-induced-transparency ground-state cooling of stationary two-dimensional ion crystals. *Phys. Rev. Lett.*, 126:023604. [180](#), [187](#)
- Ransford, A., Allman, M. S., Arkinstall, J., III, J. P. C., Cooper, S. F., et al. (2025). Helios: A 98-qubit trapped-ion quantum computer. *arXiv preprint arXiv:2511.05465*. [16](#)
- Rosenbusch, P., Ghezali, S., Dzuba, V. A., Flambaum, V. V., Beloy, K., and Derevianko, A. (2009). ac stark shift of the cs microwave atomic clock transitions. *Phys. Rev. A*, 79:013404. [172](#)
- Ruster, T., Schmiegelow, C. T., Kaufmann, H., Warschburger, C., Schmidt-Kaler, F., and Poschinger, U. G. (2016a). A long-lived zeeman trapped-ion qubit. *Applied Physics B*, 122:254. [65](#)
- Ruster, T., Schmiegelow, C. T., Kaufmann, H., Warschburger, C., Schmidt-Kaler, F., and Poschinger, U. G. (2016b). A long-lived zeeman trapped-ion qubit. *Applied Physics B*, 122:254. [66](#)
- Sárkány, L., Weiss, P., Hattermann, H., and Fortágh, J. (2014). Controlling the magnetic-field sensitivity of atomic-clock states by microwave dressing. *Physical Review A*, 90:053416. [66](#)
- Schleier-Smith, M. H., Leroux, I. D., and Vuletić, V. (2010). States of an ensemble of two-level atoms with reduced quantum uncertainty. *Phys. Rev. Lett.*, 104:073604. [168](#)
- Schwinger, J. (1948). On quantum-electrodynamics and the magnetic moment of the electron. *Physical Review*, 73:416–417. [54](#)
- Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):R2493–R2496. [17](#)

- Siverns, J. D., Simkins, L. R., Weidt, S., and Hensinger, W. K. (2012). On the application of radio frequency voltages to ion traps via helical resonators. *Applied Physics B*, 107(4):921–934. [91](#)
- Smith, M. C., Leu, A. D., Gely, M. F., and Lucas, D. M. (2024). Individually addressed quantum gate interactions using dynamical decoupling. *PRX Quantum*, 5(3):030321. [115](#)
- Smith, M. C., Leu, A. D., Miyanishi, K., Gely, M. F., and Lucas, D. M. (2025). Single-qubit gates with errors at the 10^{-7} level. *Physical Review Letters*, 134:230601. [203](#)
- Steele, A. V., Churchill, L. R., Griffin, P. F., and Chapman, M. S. (2007). Photoionization and photoelectric loading of barium ion traps. *Phys. Rev. A*, 75:053404. [179](#)
- Tarlton, J. E., Thompson, R. C., and Lucas, D. M. (2023). Surface-electrode ion trap design for near-field microwave quantum gates. *Applied Physics B*, 129:89. [115](#)
- Turchette, Q. A., Myatt, C. J., King, B. E., Sackett, C. A., Kielpinski, D., Itano, W. M., Monroe, C., and Wineland, D. J. (2000). Heating of trapped ions from the quantum ground state. *Physical Review A*, 61(6):063418. [49](#)
- Uhlenbeck, G. E. and Goudsmit, S. A. (1925). Ersetzung der hypothese vom unmechanischen zwang durch eine forderung bezüglich des inneren verhaltens jedes einzelnen elektrons. *Die Naturwissenschaften*, 13:953–954. [54](#)
- Uhlenbeck, G. E. and Goudsmit, S. A. (1926). Spinning electrons and the structure of spectra. *Nature*, 117:264–265. [54](#)
- Wang, P., Luan, C.-Y., Qiao, M., Um, M., Zhang, J., Wang, Y., Yuan, X., Gu, M., Zhang, J., and Kim, K. (2021). Single ion qubit with estimated coherence time exceeding one hour. *Nature Communications*, 12:233. [66](#)
- Wang, Y., Um, M., Zhang, J., An, S., Lyu, M., Zhang, J.-N., Duan, L.-M., Yum, D., and Kim, K. (2017). Single-qubit quantum memory exceeding ten-minute coherence time. *Nature Photonics*, 11:646–650. [66](#)
- Warring, U., Ospelkaus, C., Colombe, Y., Jördens, R., Leibfried, D., and Wineland, D. J. (2013). Individual-ion addressing with microwave field gradients. *Physical Review Letters*, 110(17):173002. [115](#)
- White, A. G. (2001). A raman approach to quantum logic in calcium-like ions. *arXiv e-prints*. 20 pages, 12 figures. [119](#)
- White, B. M., Low, P. J., de Sereville, Y., Day, M. L., Greenberg, N., Rademacher, R., and Senko, C. (2022). Isotope-selective laser ablation ion-trap loading of $^{137}\text{Ba}^+$ using a bacl_2 target. *Phys. Rev. A*, 105:033102. [179](#)
- Wilson, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212. [143](#)

Wineland, D. J. and Dehmelt, H. G. (1975). Proposed $10^{14} \delta\nu < \nu$ laser fluorescence spectroscopy on tl^+ mono-ion oscillator iii. *Bulletin of the American Physical Society*, 20:637. [54](#)

Wineland, D. J. and Itano, W. M. (1979). Laser cooling of atoms. *Physical Review A*, 20:1521–1540. [54](#)

Wood, R. P., Greene, C. H., and Armstrong, D. (1993). Photoionization of the barium $6s6p \ ^1\text{p}^\circ_1$ state: Comparison of theory and experiment including hyperfine-depolarization effects. *Phys. Rev. A*, 47:229–235. [179](#), [180](#)

Yang, D., Wang, C., Zhao, Y., et al. (2025). Ultrahigh-precision analog computing using memory-switching geometric ratio of transistors. *Science Advances*, 11(37):eady4798. [5](#)