

**Achieving a long coherence time for a multiplexed,
ion-based eEDM experiment**

by

Sun Yool Park

B.A., Oberlin College, 2019

M.S., University of Colorado Boulder, 2021

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2026

Committee Members:

Eric A. Cornell, Chair

Jun Ye

James Thompson

John Bohn

Andrew Hamilton

Park, Sun Yool (Ph.D., Physics)

Achieving a long coherence time for a multiplexed, ion-based eEDM experiment

Thesis directed by Prof. Eric A. Cornell

The observed Universe is made almost entirely of matter, with very little remaining anti-matter. The Standard Model of particle physics does not contain enough CP violation to explain the observed matter–antimatter asymmetry of the Universe. This motivates searches for additional sources of CP violation beyond the Standard Model. A nonzero permanent electric dipole moment of the electron (eEDM) would violate time-reversal symmetry and would imply CP violation, assuming the CPT theorem. The eEDM is a powerful low-energy observable for probing physics beyond the Standard Model.

The JILA eEDM experiment searches for the eEDM using trapped molecular ions. Previous generations of JILA eEDM experiments used HfF^+ and recently set a stringent upper bound on the eEDM. The next-generation JILA experiment uses ThF^+ and aims to improve the statistical sensitivity by combining long coherence (interrogation) times with a multiplexed, high-throughput experimental configuration.

This thesis presents work on the prototype apparatus of the next generation JILA eEDM experiment using ThF^+ . The main focus is understanding and extending coherence time of the eEDM-sensitive states within ThF^+ . I characterize several mechanisms that limit the experimental coherence time, including depopulation of internal states, trap loss of ions, motional heating and collisions, and dephasing from field inhomogeneity. I also present proof-of-principle studies toward continuous-mode operation, including Ramsey spectroscopy during ion translation. Together, these studies demonstrate long-coherence Ramsey spectroscopy with ThF^+ and identify several key requirements for the full next-generation experiment.

Acknowledgements

I would like to first thank my advisor Professor Eric Cornell. He has been very patient with me from very beginning. His ability to explain complicated physics in simple terms and his intuition have always amazed me. He has also created a welcoming environment for his students. Whenever I felt lost, I could go talk to him and leave his office feeling better, more grounded, and more motivated. I am deeply grateful for his guidance, patience, and support throughout my time in the group.

I would also like to thank Professor Jun Ye. His knowledge of physics, lasers, and experimental techniques is truly impressive, and I have learned a great deal from him. I am also grateful for the care he shows for students. I still remember when the HVAC in the lab was having problems and the lab became extremely cold. I was working in puffy jackets and gloves just to stay warm, and I clearly remember how strongly he advocated for us and pushed to have the problem fixed as quickly as possible, and that moment meant a lot to me.

I am also very grateful to Kia Boon Ng. When I first joined the group, he was the main person working on the ThF eEDM project, and I learned an enormous amount from him, including lasers, vacuum systems, experimental techniques, and the basic physics of the experiment. The prototype apparatus presented in this thesis would not have existed without his work. I also thank Noah Schlossberger, from whom I learned so much about machining, engineering, and practical experimental problem solving. I would also like to thank the past HfF eEDM members that I overlapped with including Tanya Roussy, Trevor Wright, Luke Caldwell, and Yuval Shagam.

I would like to thank Anzhou Wang, who joined the group only a year after me. He has

been like a little brother to me, and it has been a great pleasure to grow into physicists together. I also had the pleasure of working with the younger graduate students Addison Hartman, Patricia Hector Hernandez, and Tuan Nguyen and also undergraduate students Antonio Vigil and Rohan Kompella. Their energy, curiosity, and eagerness to learn have always been inspiring and motivating. In particular, Patricia and Tuan will be the graduate students carrying the generation-three experiment forward, and I deeply appreciate their enthusiasm and willingness to absorb the overwhelming amount of information I sometimes—well often—dump on them. I am also grateful to our postdoc Michail Athanasakis. He gave me new perspectives on what it means to be a hard-working physicist (while staying somewhat sane). He also invited me to join an experiment at ISOLDE at CERN, where I had a wonderful time learning about new physics and experiencing a very different kind of experiment. I'm very appreciative of that opportunity. I feel very close to all of my lab members. This is a difficult experiment and I think we all bonded over trauma of things breaking. I have no doubt that the experiment is in good hands, and I will miss my advisors and my lab members very much.

I would also like to thank the JILA staff members who made this work possible. In particular, I thank Terry Brown and Felix Vietmeyer from the electronic shop; James Uhrich, Hans Greens, and Calvin Schwardron from the machine shop; and J.R Raith from the computing shop. I would also like to thank Krista Beck, who is always looking after graduate students and helping us as much as possible, even if it means that she becomes buried under work. We would not have been able to carry out and do research without the support, patience, and expertise of the JILA staff. I am especially grateful given the odd hours that graduate students often work, and I know that working with tired and grumpy graduate students trying to fix their experiments is not always fun.

Finally, on a more personal note, I would like to thank my parents and my sister, who have been very supported of me from the very beginning. I am also very thankful for my new family Andy, Mimi, Cora, and Dane. I am very grateful for the awesomest group of friends I found in Boulder. They have always been supportive and encouraging when things were difficult, whether by listening to me complain or by going out and having fun at the Downer.

Many good things came out of my life in Boulder. The best thing is my fur baby Cassi. She helped me through some rough times, especially during COVID. It hurts that I cannot explain to her how much she means to me, but I hope she knows in her own doggy way.

A (close) second best thing is my husband Benton Miller. I cannot fully express my gratitude and love for him. He was always there for me when I complained, when I was stressed, and when I was grumpy, which I hate to admit was often. I am so grateful for his patience, love, and support. I cannot wait for all the adventures that lie ahead of us.

Contents

Chapter

1	Introduction	1
1.1	The Electron’s Electric Dipole Moment	1
1.2	Some Questions Behind the Standard eEDM story	3
1.2.1	T-reversal symmetry and electron’s EDM and MDM	4
1.2.2	Why must the EDM be aligned with the spin?	4
1.2.3	What is special about the EDM of an electron?	6
1.2.4	Where does the EDM figure of merit come from?	7
1.2.5	What would a nonzero eEDM measurement actually imply?	9
1.3	Thesis outline	11
2	How the JILA eEDM experiment works	13
2.1	Overview of the JILA eEDM experiment	13
2.1.1	How to measure eEDM	13
2.1.2	eEDM measurement sensitivity figure of merit	14
2.2	Why the JILA eEDM approach is powerful	15
2.2.1	Feature 1: molecules	15
2.2.2	Brief background on molecular structure	16
2.2.3	Feature 2: ions	17
2.2.4	Feature 3: $^3\Delta_1$ manifold	19

2.2.5	Feature 4: Ω doubling for differential measurement	19
2.2.6	Feature 5: REMPAD	21
2.2.7	Feature 6: experimental switches	24
2.3	Gen. III experiment Concept	24
2.3.1	New molecule ThF^+	26
2.3.2	Multiplexed experiment protocol	26
2.4	Gen. III prototype experiment	28
2.4.1	Electrodes	29
2.4.2	Lasers	33
2.4.3	Molecular beam	36
2.4.4	Ion Trap	36
2.4.5	Rotating electric field	39
2.4.6	Magnetic fields	39
2.4.7	Cryogenic system	41
2.4.8	Control system	41
2.5	Gen. III prototype measurement	42
2.5.1	Experimental Sequence	42
2.5.2	Ramsey fringe	46
3	Overview of Prototype Coherence Limits	48
3.1	Prototype goals	48
3.2	Categorizing Coherence Limits	48
3.2.1	Loss of count	49
3.2.2	Loss of contrast	50
3.3	Mechanisms vs Effects: A Structural Note	51
3.4	Roadmap of the coherence studies	51

4	Towards Long Coherence Time: Loss of Signal from Quantum Manifold	54
4.1	m_F manifold lifetime	54
4.1.1	Ion-ion collisions	55
4.1.2	Ion-neutral collisions	60
4.2	$J = 1$ state lifetime	62
4.2.1	Blackbody radiation excitation	62
4.2.2	Ion-neutral collisions	64
4.2.3	Ion-ion collisions	74
4.3	Conclusion	74
5	Towards Long Coherence Time: Loss of Signal from Trap	76
5.1	Trap lifetime of ions	76
5.2	Estimating ion temperature	77
5.2.1	Transfer-function thermometry	79
5.2.2	Laser-based thermometry	88
5.3	Observed heating in the prototype	93
5.4	Mechanisms contributing to trap loss	96
5.4.1	Energy associated with $\mathcal{E}_{\text{rot}}$	96
5.4.2	rf heating	103
5.4.3	Ion-neutral collisions	114
5.4.4	Noise on electrodes	118
5.4.5	Other possible contributions	122
5.5	Mitigation and lessons from the prototype	124
6	Towards Long Coherence Time: Alternative Trapping Configurations	126
6.1	Effective trap depth	126
6.2	Octupole trap	129
6.3	Rotating rf (rrf) ion trap	132

6.3.1	Motivation and background	132
6.3.2	Theory and Experimental Realization of the rrf trap	135
6.3.3	Implications for Precision Measurement in Non-ideal Geometries	141
6.3.4	Discussion	147
7	Towards Long coherence time: Loss of coherence	149
7.1	Temporal inhomogeneity	150
7.1.1	Temporal drift of the rotating electric field	150
7.1.2	Temporal drift of the magnetic field	153
7.2	Spatial inhomogeneity of the rotating electric field	155
7.3	Spatial inhomogeneity of the magnetic field	156
7.3.1	Magnetic multipoles	157
7.3.2	Residual magnetization and magnetic-field homogeneity	162
7.4	Collisional narrowing	163
7.4.1	Inhomogeneous dephasing in noninteracting gas	164
7.4.2	Thermalization timescales from LAMMPS	166
7.4.3	LGT rate, CDT rate, and actual “narrowing” rate	169
7.4.4	Numerical Simulations of Collisional Narrowing	172
7.5	Dephasing with applied B_{40}	177
7.5.1	B_{40} gradient coils	178
7.5.2	Experimental data	182
8	Towards continuous-mode operation	185
8.1	Ramsey fringes during translation	187
8.2	Catching and ejecting ions with $\mathcal{E}_{\text{rot}}$ continuously on	189
8.2.1	Catching ions	189
8.2.2	Ejecting ions	195
8.3	“Batch” mode operation in the prototype experiment	203

8.3.1	Timings	204
9	Generation III ThF⁺ eEDM signal	212
9.1	Experimental procedure	212
9.1.1	Experimental block and switch states	212
9.1.2	Introduction to switch-state and parity-channel notation	214
9.1.3	Switch-parity expansion	219
9.1.4	Sign definitions for chops	227
9.1.5	Quick note about the sign convention for the measured frequency	227
9.2	Brief notes on systematics	228
9.2.1	Systematic consequences of enhanced off-diagonal mixing	229
9.2.2	One subtle consequence of large off-diagonal mixing: Berry phase	231
9.2.3	Brot, Erot?	233
9.3	Sensitivity in Gen. III prototype	237
9.3.1	Long coherence time and scatter	237
9.3.2	eEDM sensitivity	240
10	Conclusion	243
	Bibliography	245
	Appendix	
A	Calibration of IG and RGA	254
A.1	Calibration of ion-gauge and RGA readings	254
A.2	Inferring the pressure at the octagon	258
A.2.1	Summary of geometric factors	264
A.3	Procedure for obtaining the H ₂ pressures in Fig. 4.8	264

B	Anti-Helmholtz configuration	266
C	Timing code	269

Tables

Table

2.1	Spectroscopic constants for the eEDM-sensitive $^3\Delta_1$ states of HfF^+ and ThF^+ . . .	27
2.2	Lasers used in the Gen. III prototype apparatus	34
4.1	Quadratic fit parameters for the secular-frequency scan for Ramsey decoherence rate	60
4.2	Predicated BBR excitation lifetime out of $X^3\Delta_1(v=0)$	63
4.3	Langevin rate constants k_L for selected gases	70
5.1	Numerical simulation of the q dependence on the fractional heating rate	108
5.2	Numerical simulation of the ion number dependence on the fractional heating rate .	109
5.3	Numerical simulation of $T/T_{\text{crossover}}$ dependence on the fractional heating rate	110
7.1	Leading rotating-frame magnetic multipoles	159
7.2	Representative LAMMPS datasets used in collisional-narrowing comparison	173
9.1	Table of old convention and new convention for parity notation	222
9.2	Sign conventions for the main chops	227
9.3	Table of old channel labels and new channel labels	229
A.1	Measured response of the RGA and ion gauge relative to the Baratron-based reference	258
B.1	On-axis magnetic field derivatives from anti-Helmholtz pair	267

Figures

Figure

1.1	Schematic of eEDM violating T reversal symmetry	2
2.1	Basic principle of an eEDM measurement	14
2.2	Rotating electric field in trapped-ion eEDM experiment	18
2.3	The $^3\Delta_1, J = 1$ manifold with Ω -doublet in zero electric field	20
2.4	Electronic structure of $^3\Delta_1$ with Ω =doublets in the presence of an electric field . . .	22
2.5	eEDM-sensitive level shifts in the two oriented Stark doublets	22
2.6	Geometry of the two-state REMPAD readout	23
2.7	Example of two-state REMPAD on MCP	25
2.8	Schematic of the multiplexed Gen. III experiment	28
2.9	Gen. III prototype experiment schematic	30
2.10	Gen. III prototype electrodes	31
2.11	Electrode assembly	32
2.12	Electrode assembly 2	33
2.13	Gen. III ThF zoo	35
2.14	Experimental sequence for Gen. III prototype	43
2.15	Example Ramsey fringe measurement for the upper and lower Stark doublets	47
4.1	Schematic of ion–ion collision tilting the local quantization axis	56
4.2	$\mathcal{E}_{\text{rot}}$ dependence on decoherence rate from ion–ion collisions	59

4.3	Secular frequency dependence on Ramsey decoherence rate from ion–ion collisions	61
4.4	CAD drawing and measured temperature along the prototype apparatus	65
4.5	Measured $J = 1$ lifetime at several positions along the prototype apparatus	66
4.6	Schematic picture of rotationally inelastic ion–neutral collisions	69
4.7	Calculated Langevin rate Γ_L for selected gases	71
4.8	Pressure-dependence $J = 1$ decay rate with added N_2 and H_2	72
5.1	Schematic illustration of the relation between confinement and effective trap depth.	78
5.2	Example of transfer function, slosh-kick method	82
5.3	Example of transfer function, shim-and-kick method	83
5.4	Relation between measured detector width and ion number 1	85
5.5	Relation between measured detector width and ion number 2	86
5.6	Photograph of the old MCP	87
5.7	Example of imperfections of old MCP	87
5.8	Photograph of the new mcp	89
5.9	Doppler spectroscopy using Hen to study number dependence	91
5.10	Example of width-based heating measurement	94
5.11	Doppler spectroscopy using Hen to study time dependence	95
5.12	COMSOL simulation of ponderomotive energy from uneven $\mathcal{E}_{\text{rot}}$	100
5.13	Measured ion survival with amplitude offset for $\mathcal{E}_{\text{rot}}$	102
5.14	Example of heating rate extraction from LAMMPS simulation	108
5.15	The ratio $A = \text{FHR}_{\text{sim}}/\text{FHR}_{\text{model}}$ across the simulated parameter space	112
5.16	Width-based heating measurement at different rf frequencies	113
5.17	Estimated ion–neutral heating rate	116
5.18	Measured heating with added neutral gas	117
5.19	Differential- and common-mode noise on the rod pairs	120
5.20	Ion heating in the vicinity of common-mode noise present in the rods	121

6.1	Different components of the effective trap depth in the prototype trap	128
6.2	Illustrative octupole-trap configurations	130
6.3	Octupole wall improvement for lowest-confinement achievable	131
6.4	Landscape of instantaneous rf field of the lrf and rrf traps	133
6.5	Transverse mode frequency dependence on q in rrf trap	137
6.6	Transverse ion dynamics comparison in lrf vs rrf trap	140
6.7	Detected transverse ion trajectory of rrf trap	142
6.8	Pseudopotential field curvature of the lrf and rrf configurations in the prototype apparatus	144
6.9	Ion survival as a function of adiabatic tilt for lrf and rrf trap	146
7.1	Illustration of Ramsey-time correction	152
7.2	Ramsey fringes with the corrected Ramsey times based on the rotating field magni- tude fluctuation	154
7.3	Example extraction of the longitudinal thermalization time τ_{LGT}	167
7.4	Example extraction of the cross-dimensional thermalization time τ_{CDT}	168
7.5	Extracted thermalization rates versus naive collision rate	170
7.6	Simulated Ramsey coherence decay for Dataset 1	174
7.7	Simulated Ramsey coherence decay for Dataset 2	175
7.8	Simulated Ramsey coherence decay for Dataset 3	176
7.9	B_{axgrad} and B_{40} coils mounted on the vacuum chamber near the octagon	179
7.10	Measured quadratic frequency curvature f_2 from B_{40} coils	181
7.11	Observed Ramsey decay time versus inferred noninteracting dephasing time	183
8.1	Conceptual illustration of throughput vs batch mode operation	186
8.2	Ramsey fringe measured during translation	188
8.3	Number of trapped ions as a function of starting phase of $\mathcal{E}_{\text{rot}}$	190
8.4	Number of trapped ions as a function of magnitude of applied anti-kick voltage . . .	191

8.5	Observation of ion sloshing in the y =direction after catching ions with $\mathcal{E}_{\text{rot}}$	192
8.6	Comparison of ion dynamics after catching ions with and without $\mathcal{E}_{\text{rot}}$	194
8.7	SIMION simulation of the exit $\mathcal{E}_{\text{rot}}$	196
8.8	Numerical simulation of ion dynamics in the presence of the exit $\mathcal{E}_{\text{rot}}$	200
8.9	Experimental result of ejecting ions with $\mathcal{E}_{\text{rot}}$	202
9.1	Ramsey asymmetry with long coherence time	238
A.1	Calibration vacuum chamber used to calibrate the ion gauge and RGA against a Baratron reference	255
A.2	Baratron measurement used to infer the N_2 pressure	257
A.3	Linearity of the N_2 calibration for the RGA and ion gauge	259
A.4	Geometry used to relate the RGA pressure reading to the octagon pressure	260
A.5	Empirical geometric calibration of N_2	262
A.6	Empirical geometric calibration for H_2	263

Chapter 1

Introduction

The search for the electron electric dipole moment (eEDM) at JILA has developed over many years, across multiple generations of experiments, molecules, and measurement techniques.

In this chapter, I begin with the “classic” motivation for eEDM experiments. After that, I take a more personal approach—I will address several questions on the eEDM experiment, some quite simple, others more subtle, that I encountered from my first year as a Ph.D student up to now. Then, I finish the chapter with the outline of the rest of the thesis.

1.1 The Electron’s Electric Dipole Moment

Since the pioneering work of Purcell and Ramsey and the early EDM searches that followed, the search for permanent eEDM became a wide experimental efforts around the world. There are many efforts on measuring the eEDM using different experimental techniques and platforms [1–8]. In 2023, our group set a new upper bound limit on the eEDM of [8]

$$|d_e| < 4.1 \times 10^{-30} \text{ e} \cdot \text{cm}.$$

What is the eEDM and why are so many groups interested in measuring this quantity? A permanent electric dipole moment of a fundamental particle violates time T symmetry. This can be seen in Fig. 1.1; under time reversal, the spin changes sign whereas the EDM is unchanged. Time reversal therefore maps a state with the EDM aligned with the spin to one with the EDM anti-aligned with

the spin. These two states are physically distinct, so a nonzero eEDM violates T symmetry.¹ By the CPT theorem, T violation implies CP violation. For this reason, a nonzero eEDM would be a clear signal of CP -violating physics.




Figure 1.1: Schematic illustration of time reversal acting on the electron with a permanent electric dipole moment. Under T , the spin reverses while the EDM is unchanged. In other words, upon T , the physical configurations of the electron are different. Therefore, a permanent EDM of an electron violates time-reversal symmetry. Figure adapted from Ref [9].

This is what makes the eEDM so interesting. The observed matter–antimatter asymmetry of the Universe suggests that additional sources of CP violation may exist. This is because CP violation is present in the Standard Model, but the known contribution from the CKM (Cabibbo–Kobayashi–Maskawa) phase is far too small to account for the observed baryon asymmetry in conventional Standard Model scenarios [10–12]. In this context, the eEDM provides one of the most sensitive probes of additional CP -violating interactions beyond the Standard Model. Within the Standard Model itself, the expected eEDM is extremely small, many orders of magnitude below present experimental reach [12, 13]. For this reason, any observation at currently accessible scale would be compelling evidence for new physics.

Experiments searching for the eEDM have steadily improved over many decades. Particu-

¹ It is also not hard to see that it also violates P symmetry. In that case, the eEDM flips sign while the spin does not.

larly, at JILA, there have been ten theses so far over the last two decades, spanning from early experimental realizations to detailed spectroscopic studies of HfF^+ and ThF^+ and to precision eEDM measurements. Although the experimental platforms have evolved, all these efforts share the same central motivation, so I keep this overview very brief in this thesis. From the experimental side, the eEDM measurements are best understood as precision phase measurements. The rest of this thesis is, in large part, about how such a measurement is made in practice.

In addition to the previous theses from our group, there are many excellent reviews of these efforts. For a broad overview of EDM searches across many physical systems, see Ref. [12]. For a review focused on T -reversal symmetry violation in atoms and molecules, see Ref. [9]. For a more general discussion of precision AMO searches for physics beyond the Standard Model, Ref. [14] is also very useful. I personally very much enjoyed Tanya Roussy’s section on how Eric got to start the eEDM search with trapped ions in her thesis [15].

Finally, I want to emphasize that a single EDM measurement does not by itself identify the underlying source of CP violation. Different systems—electrons, atoms, molecules, and nuclei—are sensitive to different combinations of CP -violating operators. As a result, a full particle-physics interpretation of EDM results requires combining many EDM measurements in a *global analysis*. This broader perspective has become increasingly important as experimental limits continue to improve; for more discussion, see Ref. [16].

1.2 Some Questions Behind the Standard eEDM story

The standard motivation for the eEDM search in the previous section is concise and compelling: a nonzero eEDM violates discrete symmetries, provides a sensitive probe for new CP -violating physics, and can be experimentally accessed through precise phase measurements. However, learning this “polished” story is not quite the same as understanding it in detail. In practice, this standard narrative hides many smaller conceptual questions. For example, even trying to understand the most basic motivation for the EDM experiment in detail turned out to be built on many layers of subtle points.

In this section, I present some non-exhaustive list of questions on these “hidden” assumptions or layers I had throughout my years working on the eEDM experiment. These are often not the kinds of questions that arise in everyday discussion. I hope that these could be useful to future students encountering the same questions.

1.2.1 T-reversal symmetry and electron’s EDM and MDM

At some point, I wondered why the electron’s electric dipole moment (EDM) $\mathbf{d}_e$ violates T reversal symmetry (or P symmetry) but its magnetic dipole moment (MDM) $\boldsymbol{\mu}_e$ does not. After all, both are tied to the angular momentum of the electron. This is a rather simple question to answer: the electric and magnetic dipole moments transform differently under P and T . The Hamiltonian can be written as

$$H_{\text{int}} \propto -\boldsymbol{\mu}_e \cdot \mathbf{B} - \mathbf{d}_e \cdot \mathbf{E}. \quad (1.1)$$

For an electron, $\boldsymbol{\mu}_e \propto \mathbf{S}$ and $\mathbf{d}_e \propto \mathbf{S}$, therefore

$$H_{\text{MDM}} \propto \mathbf{S} \cdot \mathbf{B}, \quad H_{\text{EDM}} \propto \mathbf{S} \cdot \mathbf{E}.$$

Then, these vectors transform under time reversal

$$T : \mathbf{S} \rightarrow -\mathbf{S}, \quad \mathbf{B} \rightarrow -\mathbf{B}, \quad \mathbf{E} \rightarrow \mathbf{E}, \quad (1.2)$$

or under P ,

$$P : \mathbf{S} \rightarrow \mathbf{S}, \quad \mathbf{B} \rightarrow \mathbf{B}, \quad \mathbf{E} \rightarrow -\mathbf{E}, \quad (1.3)$$

Therefore, the term $\mathbf{S} \cdot \mathbf{B}$ does not violate T - or P -symmetry, while the term $\mathbf{S} \cdot \mathbf{E}$ does. That is why the electron’s EDM violates T - and P -symmetry and its MDM does not.

1.2.2 Why must the EDM be aligned with the spin?

In the cartoon picture of the electron in Fig. 1.1, the eEDM is drawn in to be pointing in the same direction as the electron spin. In fact, the eEDM can point in the opposite direction as the spin. However, if a permanent eEDM exists, it must be collinear with the spin [12].

The basic reason comes down to rotational symmetry. For an isolated electron, the spin is the only available internal direction available. A permanent EDM is a vector quantity, but a free electron has no preferred axis—it does not have any internal spatial direction, other than its spin angular momentum. Therefore, the expectation value of any permanent vector observable must be along the electron’s angular momentum.

A more rigorous and mathematical way to explain this is with the Wigner-Eckhart theorem. For a vector operator $\mathbf{V}$ (or equivalently, a rank-1 spherical tensor V_q with $q = -1, 0, +1$), the Wigner-Eckhart theorem implies that² [17]

$$\langle \alpha, j m' | V_q | \alpha, j m \rangle = \frac{\langle \alpha j || \mathbf{V} || \alpha j \rangle}{\langle \alpha j || \mathbf{J} || \alpha j \rangle} \langle \alpha, j m' | J_q | \alpha, j m \rangle. \quad (1.4)$$

Here, α denotes any additional quantum numbers that are held fixed. This relation says that within a fixed angular-momentum manifold $\mathbf{J}$, the matrix elements of *any* vector operator are proportional to those of $\mathbf{J}$. In particular, the expectation value of $\mathbf{V}$ follows the relation of

$$\langle \mathbf{V} \rangle \propto \langle \mathbf{J} \rangle. \quad (1.5)$$

Applying this to the eEDM operator gives that result that a permanent eEDM must lie along the electron angular momentum, i.e. the spin.

Another more intuitive way to think about this is that an electron does not come with any other extra internal “direction” besides its spin. There are only two spin states of the electron: spin up and spin down. This simple two-state structure of electrons is a basic principle of atoms and the whole periodic table. If the electron had an additional internal spatial degree of freedom that could define the direction of an eEDM independently of the spin, the electron would carry more internal structure. Since no such additional direction is available, a permanent EDM must be tied to the spin direction.

² Technically, this form is Wigner-Eckhart with projection theorem

1.2.3 What is special about the EDM of an electron?

One thing that I was always curious about is why the electric dipole moment of an electron is so special. Molecules often have ordinary electric dipole moments. Then, why do people not say that because a molecule has a nonzero electric dipole moment, we have observed T -reversal symmetry breaking and therefore discovered new physics?

One explanation that is often given is that molecules are composite objects, whereas the electron is a fundamental particle. However, this explanation is not quite complete, since the EDMs of protons and neutrons are also of great interest—protons and neutrons are *also* composite objects!

The important difference is the kind of internal degrees of freedom that the dipole moment is tied to. In a molecule, there is a natural direction: the molecular axis. For instance, in a polar diatomic molecule, the molecular dipole moment is tied to the separation of positive and negative charge along the internuclear axis. Then, the molecular dipole moment becomes a geometric property of the charge distribution within the molecule. It does not need to be tied to the direction of the spin or angular momentum. That's why its existence does not imply T -violation.

One small thing worth noting here: a freely rotating molecule in a parity eigenstate does not have a permanent *lab-frame* dipole expectation value. Usually, what is meant is that a polar molecule has a permanent dipole moment in the molecular body frame. Only when the molecule is oriented and polarized, for instance with an external electric field, does a lab-frame dipole appear. But, this is again not new symmetry-violating physics.

The situation is different for an intrinsic EDM of an electron (or proton and neutron). In the electron, there is no ordinary axis like the molecular axis or any other preferred axis for the dipole to point along. As we just saw, rotational symmetry requires a permanent EDM to be collinear with the total angular momentum $\mathbf{J}$. It is this relation that the dipole is *tied* to the intrinsic angular momentum, that violates time-reversal symmetry. This is why a nonzero EDM of the electron, proton, or neutron would be evidence for new physics, whereas an ordinary molecular

dipole moment is not.

1.2.4 Where does the EDM figure of merit come from?

The statistical sensitivity of the eEDM is often written as [16]:

$$\delta d_e \propto \frac{\hbar}{2CE_{\text{eff}}\tau\sqrt{N}}, \quad (1.6)$$

where C is the contrast, E_{eff} is the effective electric field³, τ is the coherent precession time, and N is the total number of detected counts. Where did this figure-of-merit equation come from?

To tackle this question, we remind ourselves how the eEDM measurement is performed first. An eEDM measurement is usually a Ramsey measurement, in which an energy splitting is inferred from the phase accumulated between two coherent pulses [18–21]. In an eEDM experiment, the quantity of interest is the the EDM-induced energy splitting between two spin states.

The EDM interaction energy is

$$E_{\text{EDM}} = -d_e E_{\text{eff}},$$

where E_{eff} is the electric field experienced by the electron. The energy splitting between opposite spin orientations corresponds to a frequency

$$\omega_{\text{EDM}} = \frac{\Delta E_{\text{EDM}}}{\hbar} = \frac{2d_e E_{\text{eff}}}{\hbar}.$$

During a Ramsey free evolution time τ , this frequency shift accumulates a phase

$$\phi_{\text{EDM}} = \omega_{\text{EDM}}\tau = \frac{2d_e E_{\text{eff}}\tau}{\hbar}. \quad (1.7)$$

Then, we just need to determine the phase uncertainty σ_ϕ in order to to translate it into an uncertainty on d_e .

Here, it's important to remember that we never measure this phase directly. We only measure population. The Ramsey sequence converts this accumulated phase into a population difference

³ Here, E_{eff} denotes the effective internal electric field experienced by the electron in an atom or molecule. This quantity is discussed in more detail in the next chapter.

between two measured states. For N uncorrelated two-level systems, the ensemble can be represented by a collective Bloch vector [22]. The length of this vector then is approximately

$$|\mathbf{S}| \simeq \frac{CN}{2},$$

where C is the Ramsey contrast. After the first $\pi/2$ pulse, this collective spin lies in the equatorial plane of the Bloch sphere. During the evolution time, it accumulates a phase ϕ . The second $\pi/2$ pulse maps this phase onto the population difference that is then measured as spin projection.

For the most sensitive point of a Ramsey fringe (near zero crossing where the slope is highest), the measured spin projection changes as

$$\delta S_z \approx \frac{CN}{2} \delta\phi. \quad (1.8)$$

The measured spin projection also has quantum projection noise ΔS_z . For N uncorrelated spin-1/2 system in a coherent state pointing perpendicular to the measured projection axis,

$$\Delta S_z \approx \frac{\sqrt{N}}{2}. \quad (1.9)$$

This is the standard quantum noise limit for a Ramsey measurement [23, 24].⁴

Then, the phase uncertainty is simply what phase shift would produce a signal comparable to this projection noise

$$\frac{CN}{2} \sigma_\phi \simeq \frac{\sqrt{N}}{2}.$$

Finally, we get the corresponding statistical uncertainty in the EDM measurement to be

$$\sigma_{d_e} = \frac{\hbar}{2E_{\text{eff}}\tau} \sigma_\phi \sim \frac{\hbar}{2CE_{\text{eff}}\tau\sqrt{N}}.$$

This is the origin of the eEDM statistical sensitivity scaling. To improve sensitivity, we want large effective electric field, long coherence time, high contrast, and large detected signal.

⁴ This comes directly from the fact that for a single spin-1/2 particle in an equal superposition of $|\uparrow\rangle$ and $|\downarrow\rangle$,

$$\langle s_z \rangle = 0, \quad (\Delta s_z)^2 = \frac{1}{4}.$$

Then for N uncorrelated spins, the variances add [23, 24]

$$(\Delta S_z)^2 = \sum_{i=1}^N (\Delta s_{z,i})^2 = \frac{N}{4}.$$

1.2.5 What would a nonzero eEDM measurement actually imply?

This question is still one of the questions I struggle with the most. In many older eEDM theses from this group, the matter-antimatter asymmetry of the Universe appears as one of the central motivations for the experiment. Here, I will refer to that problem as the baryon asymmetry of the Universe (BAU).

The basic story is familiar: the fact that we exist is because the Universe contains more matter than antimatter. Sakharov pointed out that to explain such an asymmetry, there needs to be violation of C and CP symmetry, among other ingredients [10]. The hypothesized process that produced the asymmetry is called baryogenesis.

There are many proposed baryogenesis mechanisms. Among them, electroweak baryogenesis (EWBG) has been especially appealing because in principle it could be tested experimentally in lab because it involves physics at the electroweak scale. In EWBG, the basic idea is that during the electroweak phase transition in the early Universe, new CP -violating interactions could cause matter and antimatter to evolve slightly differently, and this difference would eventually turn into the observed baryon asymmetry⁵ [25, 26].

In simple realizations of EWBG, the same new CP violation also induces a nonzero electron EDM, which is why the eEDM was often viewed as a direct (or near direct) probe of the problem of baryogenesis. Even during this era of simple EWBG models, the connection between the baryogenesis and EDM had loopholes; in Ref. [27], it was already pointed out that if there exists some additional CP -violating effects, the direct correlation between the BAU and EDM signal can be weakened.

Over roughly the last decade, this optimistic picture turned into a more pessimistic one. The

⁵ In slightly more detail, the Higgs field had to go through some phase transition, from its high-temperature state in the early Universe to the state we know today. In EWBG, the way this happened was through “bubbles” that formed during this phase transition—inside the bubble, the new phase of the Higgs field appeared, whereas outside the bubbles, it was still the old phase [25, 26]. Then, these bubbles expanded and eventually the whole Universe phase transitioned into the phase of the Higgs field we see today. During the phase transition, these bubble walls can bias matter and antimatter differently, *if* there is CP -violating interactions present at the electroweak scale at the wall. Then another electroweak processes (called electroweak sphaleron) further convert this bias into a real net baryon asymmetry. This is about the limit of my knowledge in particle physics.

first reason was experimental: EDM searches improved dramatically and the present bound on the eEDM is far more stringent than the bounds that were available when many earlier EWBG reviews were written. For instance, the most recent HfF^+ result of $|d_e| < 4.1 \times 10^{-30} \text{ e} \cdot \text{cm}$ from our group is roughly 250 times stronger than the eEDM limits available during the time of earlier reviews [8, 26].⁶ As a result, some of the simple EWBG models have become severely constrained. In many minimal EWBG models, the same new CP -violating interactions must be large enough to produce the observed baryon asymmetry and these interactions also tend to generate an EDM. Therefore, these improved null EDM measurements force them to be small.

Most recently however, the story has become more subtle again. The reason is that the mapping between a specific model and how it produces the baryon asymmetry and the EDM signal is being refined. Recent work suggests that some older approximation schemes in these calculations may have been too simplistic and improved treatment and calculations can relax this conflict between the EDM bounds and certain EWBG models [28]. Therefore, now it seems only the minimal and “vanilla” versions of EWBG models appear to be in trouble⁷, but broader classes of EWBG models remain viable at the moment.

So is measuring the eEDM still useful? Yes, absolutely. First of all, the existing null results already have reshaped the viable model for EWBG. Secondly, a nonzero eEDM would still be extremely important because it would be clear evidence for new CP violation beyond the Standard Model, which is exactly one of the ingredients still needed for baryogenesis. What it would not do is prove that the observed BAU was generated by that same source of CP violation.

To really understand the origin of the BAU, EDM measurement must be combined with other complementary probes of electroweak-scale new physics, such as collider searches or gravitational wave signatures of the phase transition from the early Universe [25, 26]. A future nonzero eEDM measurement would therefore not be a complete answer to the baryogenesis problem, but it would

⁶ Technically, the JILA HfF^+ result improved on the previous best upper bound by about a factor of 2.4, relative to the ACME II ThO result of Ref. [1], but about a factor of 250 stronger compared to when these older reviews were written.

⁷ The title of Ref. [26] is literally “Bubble Trouble”, how cute.

be a significant clue.

So, returning to my original naive question: what would a nonzero eEDM measurement actually mean? Well, it turns out it's a loaded question, I cannot provide a simple answer. This is partly because my own understanding of particle physics is limited, but also partly because the field itself does not yet have a solid interpretation ready and waiting—it's constantly changing and updating as both the experimental and theoretical methods are improving. However, it's robust to say that a nonzero eEDM would be very compelling evidence for new CP violation beyond the Standard Model, and in many motivated models that CP violation is associated with the electroweak and/or TeV-scale new physics.

1.3 Thesis outline

The work presented in this thesis was carried out as part of the JILA ThF^+ eEDM experiment. This experiment is the third-generation of the JILA eEDM experiment, designed to improve the statistical sensitivity.

The main work of this thesis has two broad goals. The first is to understand and extend the coherence time in ThF^+ . The second is develop some of the basic experimental ingredients needed to move toward a continuous, high-throughput precision measurement.

Chapter 2 introduces some basic principles of the JILA eEDM experiment. I also discuss the motivation for the Generation III (Gen. III) experiment and describes the Gen. III prototype apparatus used in this thesis.

Chapters 3-7 focus on efforts to achieve long coherence times. Chapter 3 lays out the different classes of mechanisms that can limit the coherence time. Chapter 4 investigates the lifetime of the relevant quantum manifold. Chapter 5 studies the trap lifetime and examines several possible heating mechanisms in the system. Chapter 6 introduces new trapping configurations that we implemented in an attempt to improve the trap lifetime in the prototype apparatus. Finally, Chapter 7 investigates decoherence mechanisms that directly limit the measured coherence time.

Chapter 8 discusses several key ingredients and proof-of-principle experiments needed for a

future continuous, high-throughput precision measurement. Chapter 9 introduces a systematics notation that will be used for the Gen. III experiment and gives a crude estimate of the Gen. III prototype apparatus sensitivity to d_e . Finally Chapter 10 summarizes the thesis and discusses future directions.

Chapter 2

How the JILA eEDM experiment works

In this chapter, I first review the basic experimental approach of the JILA eEDM experiment. Then I motivate and introduce the Gen. III prototype apparatus, which is the main focus of this thesis.

2.1 Overview of the JILA eEDM experiment

The JILA eEDM experiment uses trapped polar molecular ions [3, 8, 29, 30]. At its core, the JILA eEDM experiment is a Ramsey spin-precession measurement: the two spin quantum states are prepared into a coherent superposition, allowed to evolve for some interrogation time and then read out by measuring the accumulated relative phase. This accumulated phase contains the contribution from the eEDM.

2.1.1 How to measure eEDM

For a bare electron in external electric and magnetic fields, the interaction Hamiltonian is given by

$$H_e = -\boldsymbol{\mu} \cdot \mathbf{B} - \mathbf{d}_e \cdot \mathbf{E} \quad (2.1)$$

Since both the magnetic dipole moment and the electric dipole moment must be collinear with the spin, this can be written as:

$$H_e = -\frac{g_s \mu_B}{\hbar} \mathbf{S} \cdot \mathbf{B} - \frac{d_e}{\hbar} \mathbf{S} \cdot \mathbf{E}. \quad (2.2)$$

The basic idea of an eEDM is a spin experiment. In a magnetic field, opposite spin projections are split by the Zeeman interaction. If the electron also has a nonzero electric dipole moment, then there is an additional spin-dependent shift in the presence of an electric field. Reversing this electric field reverses the sign of the eEDM contribution, as can be seen in Eq. (2.2). By performing a differential measurement between opposite orientations of the electric field, the dominant magnetic contribution cancels to first order, leaving a signal directly proportional to the eEDM. Figure 2.1 shows this basic experimental procedure schematically.




Figure 2.1: Schematic illustration of the basic eEDM measurement. Reversing the orientation of the electric field reverses the eEDM-induced energy shift while leaving the leading the magnetic contribution unchanged to first order.

In a real experiment, however, a free electron is not directly probed. Instead, the electron is bound inside an atom or molecule. In particular, in polar molecules, the relevant interaction is greatly enhanced because the valence electrons experience a very large effective electric field [31, 32]. This is one of the main reasons that molecules are such powerful platforms for eEDM searches.¹

2.1.2 eEDM measurement sensitivity figure of merit

The eEDM measurement is performed using Ramsey spectroscopy. As discussed in Chapter 1, the statistical sensitivity scales approximately as [12]

$$\delta d_e \propto \frac{\hbar}{2CE_{\text{eff}}\tau\sqrt{N}} \quad (2.3)$$

¹ For similar reasons, molecules are attractive platform for other types of EDM searches.

where C is the Ramsey fringe contrast, E_{eff} is the effective electric field experienced by the electron, τ is the interrogation time, and N is the total number of detected ions or counts. At the most basic level, this figure of merit tells us that the statistical sensitivity improves when the experiment has

- (1) a larger effective electric field,
- (2) more detected signal
- (3) a longer interrogation (coherence) time.

Many beam-based EDM experiments primarily take advantage of the large number of molecules. At JILA, by contrast, the main strategy is to gain sensitivity through long interrogation time [8, 33]. For that reason, the JILA eEDM experiment utilizes molecular ions, which can be easily trapped for a very long time. In the next section, I review several features of our experiment that enable us to perform such a precise measurement of the eEDM.

2.2 Why the JILA eEDM approach is powerful

There are two central experimental challenges for eEDM measurements. First, since the eEDM is such a tiny effect, one must maximize the size of this phase shift coming from the eEDM. Second, one must also suppress and/or reject technical systematic shifts that can mimic such a signal. To tackle these two challenges, at JILA, our strategies include: the use of polar molecular ions, use of the quantum $^3\Delta_1$ manifold in the molecule, differential measurements enabled by Ω -doublets, long interrogation times made possible by ion trapping, and the use of multiple experimental switches to diagnose systematic effects.

In this section, I summarize the basic measurement principle and main features of the JILA eEDM experiment.

2.2.1 Feature 1: molecules

To measure the EDM of the electron, you might naively want to use bare electrons. In practice, however, eEDM experiments are almost always performed using electrons bound in atoms

or molecules. This immediately gives us one practical advantage: the electron is naturally confined as part of the atomic or molecular system. At the same time, it turns out the internal structure of the atom or molecule can enhance the energy shift from the eEDM contribution [31, 32].

At JILA, we use polar molecules rather than atoms. The most important advantage of using polar molecules is that a polarized molecule provides a very large effective electric field experienced by the unpaired electron [31, 32, 34]. This effective field is often many orders of magnitude larger than any static electric field that can be applied in the lab; therefore it strongly amplifies the energy shift produced by an eEDM.

A second more subtle advantage is that molecules can often be polarized with comparatively modest laboratory electric fields. In an atom, polarization requires mixing opposite-parity electronic states, which are typically separate by optical-scale energies. In a molecule, by contrast, the relevant opposite-parity states are often separated by much smaller energies, because they arise from rotational structure or even closely spaced parity doublets in some special molecules [32]. As a result, a smaller laboratory electric field can already mix these states much more efficiently and produce a polarized molecule.

However, there are disadvantages of using molecules. Mainly, the level structure of molecules is far more complicated than that of atoms. In addition to electronic structure, molecules have rotational and vibrational structure, producing a much richer level diagram than atoms. This complexity must be first understood and exploited carefully before the molecule can be used for an eEDM measurement.

2.2.2 Brief background on molecular structure

Before continuing the features of the JILA eEDM experiment, I want to take a quick moment and introduce some basic information on molecular structure.

For a linear molecule, the electronic orbital angular momentum and electronic spin angular momentum projected onto the internuclear axis are denoted by Λ and Σ , respectively. Their sum,

$$\Omega = \Lambda + \Sigma,$$

is the projection of the total electronic angular momentum projected onto the internuclear axis. The total angular momentum excluding nuclear spin is denoted by $\mathbf{J}$. Then, with a nuclear spin $\mathbf{I}$, the total angular momentum is

$$\mathbf{F} = \mathbf{J} + \mathbf{I},$$

and its projection onto the *laboratory frame*² is labeled by m_F . Then, the molecular term symbol is given by [35]

$$^{2S+1}\Lambda_{\Omega}^{+/-}, \quad (2.4)$$

where Λ is written with a capital Greek letter as $\Sigma, \Pi, \Delta, \Phi, \dots$ for the corresponding $\Lambda = 0, 1, 2, 3, \dots$. The $2S + 1$ term denotes the spin multiplicity, and $+/-$ indicates the reflection symmetry along an arbitrary plane [35].

For the case relevant here, the $^3\Delta_1$ state within ThF^+ is well approximated by Hund's case (a) [33, 35, 36]. In Hund's case (a), the electronic orbital angular momentum $\mathbf{L}$ is strongly coupled to the internuclear axis, and spin-orbit coupling couples the spin to that axis as well, such that Ω is a good quantum number. This state has $S = 1$, $\Lambda = 2$, and $\Omega = 1$. For ThF^+ , the F atom has $I = 1/2$, therefore the $J = 1$ manifold used for the eEDM measurement as two hyperfine manifolds, $F = 1/2$ and $F = 3/2$.

For more information on general molecular structure, see Ref. [35]. For the structure specific for HfF^+ and ThF^+ , see Refs. [37, 38].

2.2.3 Feature 2: ions

Unlike many other eEDM experiments, the JILA experiment uses ions. The main advantage of ions is that they are relatively easy to trap. This trapping capability allows long interrogation times, which is very valuable for precision measurements based on coherent phase accumulation. In addition, it also allows us to actually vary the interrogation time, which is one of the central strengths of the JILA approach. In contrast to beam experiments where the interaction time is

² This is an important distinction, that m_F is the projection onto the lab frame, whereas Σ, Λ , and Ω are all projections within the molecular frame.

often set by the beam-line geometry and beam velocity, trapped ions allow flexible control over the Ramsey evolution time.

However, ions also introduce an important complication. In order to perform the eEDM measurement, the molecules must first be polarized by an external electric field. Because the molecules are charged, a static polarizing electric field would exert a large force on the ions and eject them from the trap. The solution used at JILA is to apply a rotating electric bias field, instead of a static one.³




Figure 2.2: Schematic illustration of the rotating electric bias field used to polarize trapped ThF^+ ions. At several times during one rotation cycle, the molecular axis follows the instantaneous direction of the applied $\mathcal{E}_{\text{rot}}$ adiabatically, while the ion undergoes circular motion in the radial plane, with orbit radius r_{rot} . In this way, the molecular ion remains polarized without being expelled from the trap.

This rotating field polarizes the molecule, so that the molecular quantization axis follows the instantaneous field direction. As the field rotates slowly compared to the relevant molecular level splittings so that the molecule's polarization follows the field adiabatically [29, 30]. However, the field rotates fast compared to the ions' motion, so the ions don't get driven out of the trap. The rotating field drives the ion into a well-defined circular orbit in the radial plane, while allowing the ions to remain trapped as shown in Fig. 2.2. This feature will come in handy as a way to produce an effective Zeeman shift in the rotating frame, as will be discussed later.

³ When my advisor Eric Cornell came up with the rotating electric field idea for precision measurement, apparently, a lot of people in the field thought this idea was crazy and never going to work. But now, several groups have proposed of using rotating electric field.

2.2.4 Feature 3: $^3\Delta_1$ manifold

The specific electronic manifold used in the eEDM experiment is the $^3\Delta_1$ manifold. This manifold has several properties that make it especially attractive for eEDM measurement. One of them is that this state has unusually small magnetic sensitivity [32, 33].

In a $^3\Delta_1$ state, the spin and orbital angular momentum contributions to the magnetic moment largely cancel. This leads to a magnetic moment that is anomalously small [32]. This suppression of the magnetic moment is a major advantage because it reduces sensitivity to magnetic field noise and magnetic systematic effects, which are often the largest contribution to systematics budget in most EDM experiments.

In earlier HfF^+ experiments at JILA, the $^3\Delta_1$ manifold was a low-lying metastable state [3, 8]. For the Gen. III, the molecule is changed to ThF^+ , in which the $^3\Delta_1$ state is the true electronic ground state [33, 36]. This is a major reason for the change of the molecular species, so that the spontaneous decay no longer sets the coherence time in ThF^+ , in the same way as in HfF^+ .

2.2.5 Feature 4: Ω doubling for differential measurement

Another very powerful feature of some molecules is the presence of very closely spaced opposite parity levels, often referred to as Ω -doublet states [35, 39, 40]. These states are a pair of opposite parity eigenstates, usually labeled e and f , formed from the nearly degenerate $\Omega = \pm 1$ components [35]

$$|e\rangle = \frac{(|\Omega = +1\rangle) - |\Omega = -1\rangle}{\sqrt{2}}, \quad |f\rangle = \frac{(|\Omega = +1\rangle) + |\Omega = -1\rangle}{\sqrt{2}}.$$

In a linear molecule with nonzero electronic angular momentum along the internuclear axis, the states with projections Ω and $-\Omega$ are nearly degenerate. Weak rotational and electronic couplings lift this degeneracy and produce the small Ω -doublet splitting, ω_{ef} .

Therefore, within the $^3\Delta_1$ manifold, the hyperfine structures $F = 3/2$ and $F = 1/2$ are further divided into $|e\rangle$ and $|f\rangle$ states, as shown in Fig. 2.4.

In zero electric field, these are states of good parity. In an applied external electric field,




Figure 2.3: Qualitative level structure of ${}^3\Delta_1, J = 1$ with the Ω -doublet. In zero electric field, the eigenstates are opposite-parity states e and f , separated by ω_{ef} . Figure adopted and modified from Ref. [33].

however, the opposite-parity states mix and become states of good orientation: aligned or anti-aligned with the external field, as shown in Fig. 2.5.

This makes the Ω -doublet structure ideal for differential measurements. By comparing signals from the two opposite-oriented states, one effectively reverses the sign of the eEDM-sensitive interaction without physically reversing any applied fields as shown in Fig. 2.5. This built-in reversal is the basis of the internal comagnetometry used in many molecular eEDM experiments and was first developed by Prof. Dave DeMille [40].

This feature is especially powerful in our experiment because the apparatus uses many pulsed lasers for ion production and readout, and they introduce significant technical noise in the total signal. If both doublets can be detected simultaneously, much of this common-mode noise cancels while preserving sensitivity to the eEDM [8, 41].

2.2.6 Feature 5: REMPAD

How do we actually detect both doublets at the same time? We use resonantly enhanced multiphoton asymmetric dissociation (REMPAD) protocol [41–43]. In this method, the first photon drives a state-selective bound-to-bound transition, to an intermediate state. The second photon excites the molecule above the dissociation threshold. For a particular combinations of the intermediate state, final state, and the laser polarization, the molecule breaks apart in an anisotropic pattern, such that the heavy fragment ion acquires recoil whose direction depends on the initial quantum state [42, 43].

The detailed physics that determines the angular distribution of the photofragments is somewhat involved and will not be reviewed here; for broader discussions, see Refs. [38, 44–46]. Instead, I summarize only the aspects most relevant for the experiments in this thesis. Because of the Stark-doublet structure, ThF^+ can be instantaneously aligned or anti-aligned with respect to the rotating bias electric field. Semiclassically, the photofragments are then ejected preferentially along the molecular axis. The resulting Th^+ fragments are extracted onto an imaging microchannel-plate (MCP) detector in a direction orthogonal to the recoil axis. As a result, ions ejected with opposite




Figure 2.4: Schematic illustration of the electronic structure of the ${}^3\Delta_1, J = 1, F = 3/2$ manifold in a presence of an external electric field, here denoted as $\mathcal{E}_{\text{rot}}$. The external electric field mixes the parity states, to produce states of good eigenstate. These states are split by the external electric field and the molecular-frame dipole moment, giving rise to Stark upper and lower doublets. In the upper doublet, the internuclear axis $\hat{n}$ is anti-aligned with the external electric field, and in the lower doublet, it is aligned. Figure adopted and modified from Ref. [29].




Figure 2.5: Schematic energy-level diagram showing the eEDM-sensitive spin precession frequencies in the two oriented Stark doublets. The magnetic interaction contributes to (almost) the same Zeeman shift, proportional to $g_F \mu_B \mathcal{B}_{\parallel}$, in both cases. The eEDM contribution changes sign between the two doublets because the effective internal electric field $\mathcal{E}_{\text{rot}}$ is parallel to $\mathcal{E}_{\text{rot}}$ in one oriented state and antiparallel in the other. As a result, the measured transition frequencies f^u and f^l acquire opposite eEDM shifts, enabling differential measurement and internal comagnetometry. Figure adopted and modified from Ref. [29].

recoil land on opposite halves of the detector, as Fig. 2.6 shows. In this way, the two oriented states, and therefore the two Stark doublets, can be distinguished from each other in a single measurement [38, 41, 43].




Figure 2.6: Geometry of the two-state REMPAD readout used for simultaneous detection of the two oriented states in ThF^+ . The rotating electric field defines the instantaneous molecular orientation, and the REMPAD lasers dissociate the molecules into Th^+ and F fragments. The Th^+ fragments are then extracted onto an imaging MCP. Depending on which Stark doublet the ThF^+ was in at the moment of dissociation, the Th^+ ions land on opposite sides of the detector. Figure from Ref. [43].

Quantum mechanically, the fragment angular distribution depends on both the initial molecular state and the polarization of the dissociation photons [44–46]. To maximize the directional asymmetry between the two orientation states, it is generally advantageous to dissociate from a state with quantum numbers satisfying $J = |\Omega| = |m_J|$. In our case, we begin in the $^3\Delta_1(v = 0, J = 1, m_J = \pm 1)$ manifold. This motivates driving an $R(1)$ transition with a circularly polarized first REMPAD photon to an intermediate state with $J = |\Omega| = |m_J| = 2$, followed by a second circularly polarized photon that excites the molecule above the threshold to a continuum state with $|\Omega| = 3$ character. In practice, this leads us to search for $|\Omega| = 2$ intermediate states that combine high dissociation efficiency with sufficiently large photofragment kinetic energy⁴.

This readout method is particularly powerful because it combines state selectivity with spa-

⁴ The photofragment kinetic energy is determined by how much of the excess energy appears as translation kinetic energy after accounting for internal excitation of the atomic fragments. In a heavy species such as Th, the large number of low-lying excited atomic states often makes this determination challenging, since much of the excess energy may instead go into atomic excitation.

tially resolved detection of the dissociation recoil. In recent work on ThF^+ , a dissociation survey identified significantly more efficient transitions than were previously available, improving the dissociation efficiency from the few-percent level to the tens-of-percent level [43]. This transition also gave us a good spatial separation between the two doublets, as shown in Fig. 2.7(c).

2.2.7 Feature 6: experimental switches

To diagnose and suppress systematic effects in our system, we use experimental switches. These switches are deliberate reversals of some of the key experimental parameters that change the sign of some contributions to the signal, while keep others unchanged.

For instance, the most obvious example is the reversal of the effective electric-field orientation relative to the spin quantization axis. This reverses the sign of the eEDM-induced energy shift. Other switches include reversal of the bias magnetic field, reversal of the rotation direction of the rotating electric field, etc.. By organizing the measured signal into different parity channels under these switches, as will be discussed in more detail in Chapter 9, we can identify which components have the symmetry expected for a true eEDM signal and which are more likely to arise from imperfect systematic effects.

2.3 Gen. III experiment Concept

So far, we have seen some features that make the JILA eEDM special. The JILA eEDM experiment has progressed through several generations. Generations I and II used HfF^+ molecular ions. Generation III is designed to improve the eEDM sensitivity through two major ways:

- (1) change of molecular ion species from HfF^+ to ThF^+
- (2) development of a higher-throughput, multiplexed experimental protocol.

In the following sections, I discuss these two upgrades in more detail.




Figure 2.7: Comparison of candidate two-state REMPAD protocols for simultaneous readout of the two oriented Stark states shown with two different colors. Each panel shows the detector image and projected histograms for a different choice of intermediate and continuum quantum numbers. Pathways that proceed through an intermediate state with $J = |\Omega| = |m_J| = 2$ produce significantly larger spatial separation on the detector, compared to the pathway through $m_J = 0$. The nominal dissociation transition used in this thesis corresponds to the right panel.

2.3.1 New molecule ThF⁺

The Gen. III experiment uses ThF⁺, rather than HfF⁺. The main motivation for the change is that the eEDM sensitive $X^3\Delta_1$ manifold in ThF⁺ is the true electronic ground state. This removes spontaneous decay as a fundamental limitation on the coherence time, in the way it appears in HfF⁺. Therefore, the molecular species ThF⁺ is attractive for long-coherence measurements.

In addition, ThF⁺ is expected to have a larger effective electric field than HfF⁺, leading to greater intrinsic eEDM sensitivity enhancement. These advantages make ThF⁺ a great species for the next-generation experiment.

Even though there are two different molecular ions, they share a lot of the molecular structures with different spectroscopic values. The different values between the two species are shown in Table 2.1.

2.3.2 Multiplexed experiment protocol

We can break down the eEDM figure of merit of Eq. (1.6):

$$\delta d_e \propto \frac{\hbar}{2CE_{\text{eff}}\tau\sqrt{N_{\text{single}}f_{\text{rep}}T_{\text{tot}}}}, \quad (2.5)$$

where N_{single} is the number of detected ions per single experimental shot, f_{rep} is the repetition rate of the experiment, and T_{tot} is the total integration time for which we take data. In a conventional EDM experiment, the repetition rate of the experiment is often set by the coherence time, roughly $f_{\text{rep}} \approx 1/\tau$. In this case, increasing the coherence time does not improve the sensitivity linearly, but instead as $\propto 1/\sqrt{\tau}$.

This means going from the coherence time of ~ 2 s achieved in HfF⁺ to ~ 20 s in ThF⁺ would only improve the statistical sensitivity by a factor of $\sqrt{10} \approx 3$, rather a factor of 10 one might have hoped for from the longer coherence time alone. We want to be greedy and recover the full factor of 10. One way to achieve this is by decoupling the coherence time from the repetition rate. In Gen. III, this is done by multiplexing our experiment.

Figure 2.8 shows the basic multiplexing schematic for Gen. III. We first create an ion cloud

Table 2.1: Spectroscopic constants for the eEDM-sensitive $^3\Delta_1$ states of HfF^+ and ThF^+ . References for individual entries are given in the final column. The HfF^+ constants are taken from Ref. [3, 30, 47–51]. The ThF^+ constants are compiled from Refs. [33, 36, 38, 50, 52, 53]; see table notes for unit conversions.

Constant	HfF^+	ThF^+	Description	Reference
B_e/h	8.983(1) GHz	7.288(2) GHz	Equilibrium rotational constant	[36, 47]
$A_{ }/h$	−62.0(2) MHz	−20.1(1) MHz	^{19}F hyperfine constant	[3, 33, 38]
d_{mf}/h	1.97(1) MHz V^{-1} cm	1.70(5) MHz V^{-1} cm ^b	Molecule-frame electric dipole moment	[33, 38, 51]
$\omega_{ef}/(2\pi)$	0.74(4) MHz	5.29(5) MHz	Ω -doubling constant	[33, 38, 47]
g_N	5.257 74(2)	5.257 74(2)	Nuclear magnetic g factor of ^{19}F	[50]
$G_{ }$	−0.0122(3)	sign ambiguous ^c	Effective electronic g factor	[33, 38, 48]
g_F	−0.0031(1)	$ g_F = 0.0149(3)$	$F = 3/2$ state g factor	[3, 30, 33, 38]
$ \mathcal{E}_{\text{eff}} /h$	5.5×10^{24} Hz e^{-1} cm ^{−1} ^a	9.0×10^{24} Hz e^{-1} cm ^{−1} ^a	Effective electric field	[49, 52, 53]

^a Converted from $|\mathcal{E}_{\text{eff}}| = 22.5$ GV/cm for HfF^+ and 37.3 GV/cm for ThF^+ using 1 GV/cm = 2.41799×10^{23} Hz (e cm)^{−1}.

^b Converted from $d_{\text{mf}} = 3.37(9)$ D of Ref. [33] using 1 D = 0.5034 MHz (V/cm)^{−1}.

^c The sign of g_F has not been determined experimentally. Neglecting the rotational contribution, the corresponding values are $G_{||} = -0.042(2)$ if $g_F > 0$ and $G_{||} = +0.048(2)$ if $g_F < 0$ [33].

and prepare it for the precision measurement. Once that cloud begins its coherent evolution, it is translated along the length of the trap. Meanwhile, a new ion cloud is created and prepared in the same upstream region. This process is repeated and in steady state, multiple ion clouds occupy the trap at the same time each at a different state of its own Ramsey evolution. The repetition rate is then set by how quickly we can create and prepare successive ion clouds. The coherence time is set by the time it takes each cloud to travel from the state-preparation region to the state-detection region. In this way, Gen. III can keep the long coherence time without slowing down the experimental repetition rate.




Figure 2.8: Schematic of the multiplexed Gen. III measurement protocol. A neutral ThF beam enters and is ionized in the state preparation region. After state preparation, each ion cloud is translated along the trap while undergoing coherent Ramsey evolution for a time τ . New ion clouds are prepared at a repetition rate f_{rep} . In steady state, many clouds are stored in the trap simultaneously, each going through its own Ramsey evolution. At the downstream end, the molecules are dissociated and detected on the main MCP.

2.4 Gen. III prototype experiment

The full Gen. III eEDM experiment is designed to operate with trapped ThF^+ ions in a multiplexed apparatus. Going directly to the final architecture from the earlier-generation apparatuses, however, requires a significant increase in experimental complexity. To systematically test major technical issues, we developed a prototype Gen. III experiment, before building the full apparatus.

Development of the prototype began during the COVID-19 period, when access to the laboratory was limited, and fabrication and assembly took place in roughly the 2021-2022 time frame.

Because the prototype apparatus has already been discussed in detail in the these of Kia Boon Ng and Noah Schlossberger, here I only summarize the aspects most relevant for the measurements presented in this thesis.

The prototype was built to address two major requirements. The first was to demonstrate continuous-mode operation: the apparatus must be able to support a large-scale ion trap geometry in which ions can be translated along the trap axis while remaining confined and furthermore while preserving the coherence needed for Ramsey spectroscopy. The second was to establish the field control, state preparation, and state readout needed for performing precision Ramsey spectroscopy of ThF^+ with long coherence time.

Figure 2.9 shows a schematic of the prototype apparatus. In the following subsections, I briefly describe the major subsystems relevant to the measurements discussed in this thesis.

2.4.1 Electrodes

In the prototype apparatus, the ion trap consists of rod and ring electrodes. The eight long rod electrodes run along the length of the apparatus and generate electric fields in the radial xy plane. These electrodes are used both to produce the rf quadrupole field for radial confinement and to generate the rotating bias field used to polarize the molecular ions.

Axial fields are produced by ring electrodes, spaced uniformly along the trap axis by 2 cm. These rings provide the static axial confinement as well as the time-dependent voltages used for ion transportation along the length of the trap. The overall electrode configuration is shown in Fig. 2.10. The electrodes were designed and fabricated in collaboration with Calvin Schwardron in the JILA instrument shop.

All eight rod electrodes are driven independently. The ring electrodes are divided into three groups: quadrant rings, special rings, and translation rings. The eight quadrant rings are segmented into four isolated quadrants in the $\pm x$ and $\pm y$ directions around the central octagon region. These rings allow us to apply controlled radial dc fields near the trap center and are used for stopping fields, static shims, and impulsive kicks in the radial plane.




Figure 2.9: Gen III prototype experiment schematic. Neutral ThF is produced in source chamber. They go through supersonic expansion. Once they arrive at the octagon, they get ionized and trapped. The necessary state preparation is performed at the octagon, where we have optical access. Once they are ready to go through quantum evolution, they are transported to the Science (cold) region. The two $\pi/2$ pulses are applied here. When they are read to be detected, they are transported back to the octagon, where state readout occurs. The ions are dissociated and extracted towards the ion detector (MCP).




Figure 2.10: Electrode geometry of the Gen. III prototype trap, showing the rod electrodes for radial confinement and rotating bias field, together with the ring electrodes used for axial confinement and transport. There are three types of rings: special (red), quadrant (yellow), and translation (green) rings. Their roles are described in the main text. The quadrant rings are divided into four quadrants in the xy -plane, and they are used to apply controlled radial electric fields near the trap center.

The ten special rings, five on each side of the quadrant-ring region, are all wired independently and can therefore be driven separately. The special rings provide the kickout field needed to detect the ions on the detector.

The translation rings, by contrast, are wired through the spines, which are square structural rods running along the assembly. Every four rings are connected to a pair of spines located 180° apart, so that the translation voltages can be applied through the spines, rather than by routing individual wires through the entire trap. In practice, this means that the 41 physical translation rings in the prototype are driven using only four distinct applied voltages distributed through two spine pairs.

The quadrant rings and special rings are driven by power op-amps and can reach voltage as high as $\pm \sim 400$ V. On the other hand, the translation rings are limited to about ± 10 V.

Prior to assembly, all electrodes were cleaned by ultrasonic bath, solvent washing, and baking at approximately 300°C . To minimize contamination and unintended magnetization, assembly was carried out in a small dedicated “clean” room using nonmagnetic tools and storage hardware. The final electrode assembly was completed by myself, Anzhou Wang, and Calvin Schwardron.




Figure 2.11: Photograph of the completed electrode assembly for the Gen. III prototype apparatus.




Figure 2.12: Photograph of the completed electrode assembly and me standing next to it for scale (I'm about 5'5"!)

2.4.2 Lasers

The prototype experiment uses ten lasers in total: five pulsed lasers for ablation, ionization, and dissociation, and five cw lasers for state preparation and readout. Within the group, these lasers are referred to by animal names. A summary of lasers used is given in Table 2.2 and Fig. 2.13.

Because thermal Doppler broadening is the dominant linewidth contribution in the present system, absolute laser-frequency stability requirements are modest compared with those in narrow-line atomic experiments. Therefore, the cw lasers are locked to wavemeters using relatively simple servo loops with update times of order ~ 1 s, giving typical stability at the level of about $\sim \pm 30$ MHz. Currently, the pulsed lasers are not locked continuously, but they are monitored during operation and manually corrected as needed to remain within roughly $\sim \pm 500$ MHz.

Table 2.2: Photons used in the Gen. III ThF⁺ experimental sequence. Unless otherwise noted, all transitions address the $v = 0$ vibrational level of the relevant electronic states, except for Monkey, which addresses $v = 1$. All the lasers propagate in the xy -plane. The cw beams have typical diameters of approximately $2\sigma \sim 10$ mm.

Name	Symbol	Transition	Power/Energy	λ (nm)	Polarization	Pulse width	Notes
Ablation	Hamster	–	3 mJ	532	Linear	10 ns	
Ionization 1	Lion	$\Omega = 3/2 \leftarrow X^3\Delta_2$	100 μ J	304.3728	Linear	10 ns	
Ionization 2	Horse	$? \leftarrow \Omega = 3/2$	500 μ J	532	Linear	10 ns	
Rotational repump	Sheep	$P(2) \Omega = 0^- \leftarrow {}^3\Delta_1$	70 mW	685.489059	Linear	CW	
Vibrational cleanup	Monkey	$P(1) \Omega = 0^- \leftarrow {}^3\Delta_1$	90 mW	717.59123	Linear	CW	
m_F pumping	Goat	$Q(1) \Omega = 0^- \leftarrow {}^3\Delta_1$	5 mW	685.443527	Circular	CW, strobed	
m_F depletion	Goat	$Q(1) \Omega = 0^- \leftarrow {}^3\Delta_1$	5 mW	685.443527	Circular	CW, strobed	
m_F depletion	Hen	$Q(1) \Omega = 0^+ \leftarrow {}^3\Delta_1$	–	954.939194	Circular	CW, strobed	Currently not used
Rotational cleanup	Chick	$P(1) \Omega = 0^+ \leftarrow {}^3\Delta_1$	180 mW	954.987809	Circular	CW	
Dissociation 1	Panther	$\Omega = 2 \leftarrow {}^3\Delta_1$	100 μ J	272.27199	Circular	10 ns	“Banana” transition
Dissociation 2	Bear	$? \leftarrow \Omega = 2$	10 mJ	355	Circular	10 ns	

Figure 2.13: Schematic zoo of all the lasers used in the Gen. III experiment, corresponding to the lasers listed in Table 2.2. The animal icons denote the laser names used in the experiment. The left panel shows the cw optical pumping, cleanup, and repump transitions, while the right panel shows the pulsed REMPI and REMPD transitions used for ionization and detection. The diagram is not drawn to scale.

2.4.3 Molecular beam

A pulsed beam of neutral ThF molecules is produced in the source chamber by ablating a piece of metal target containing Th with a 532 nm pulsed Nd:YAG laser into a pulsed supersonic expansion of Ne seeded with SF₆. The target is mounted on a combined rotation and translation stage and is stepped periodically to expose a fresh surface to the ablation laser. The target surface is also polished regularly to maintain stable production.

Chemical reactions in the ablation plume produce a “soup” of thorium fluorides, including neutral ThF. After supersonic expansion, the molecules are rotationally and vibrationally cooled to a temperature of order 10 K. Once the neutrals arrive in the main science chamber, ThF is ionized by a two-photon ionization scheme using 313 nm and 532 nm pulsed lasers. The first photon excites a state-selective transition from the neutral ground state $^2\Delta_{3/2}$. The second photon pushes the molecule above threshold through an autoionizing state only about 60 cm⁻¹ above the ionization limit.

This ionization scheme predominantly produces ThF⁺ in the $X^3\Delta_1$ electronic ground state with $v = 0$. Because the total energy lies below the next vibrational manifold (about 650 cm⁻¹ above) and below the first electronically excited state (about 300 cm⁻¹ above), the ions are produced primarily in the lowest rovibronic manifolds, with population concentrated in the first few rotational states.

At the moment of ionization, stopping fields are pulsed on using the quadrant rings to decelerate the newly formed ions, and the trapping potentials are turned on immediately to capture them. Typically, about ~ 5000 ThF⁺ are trapped.

2.4.4 Ion Trap

Comprehensive discussions on Paul traps can be found in standard references [54–56], as well as in earlier theses from this group [37, 57, 58]. Therefore, I do not review the general theory of Paul traps here.

The prototype uses a linear Paul trap geometry formed by the eight rod electrodes for radial confinement together with ring electrodes that provide axial confinement and ion transport. Axial confinement and transport are generated by modulating the voltages applied to the ring electrodes. In particular, a traveling-wave pattern of the form

$$V(z_n, t) = \frac{U_{ec,0}}{2} \sin(kz_n - 2\pi f_{\text{trav}}t), \quad (2.6)$$

is applied to the ring at axial position z_n , where $k = 2\pi/\lambda$ and λ is the spatial period of the traveling potential, i.e. the bucket size, and $U_{ec,0}$ denotes the peak-to-peak voltage applied to the rings.

This voltage pattern produces a moving sinusoidal axial potential. Near a local minimum, the potential is approximately harmonic with the form

$$U_{\text{axial}} \approx \frac{U_{ec,0}}{Z_0^2} z^2. \quad (2.7)$$

The term Z_0 is a geometry factor that relates the applied ring voltage to the axial curvature. It is determined using COMSOL, and in the prototype apparatus, $Z_0 \sim 33$ cm.

Radial confinement is provided by the electric field,

$$\mathcal{E}_{\text{rf}}(\mathbf{r}, t) = \frac{V_{\text{rf},0}}{R_0^2} (x\hat{\mathbf{x}} - y\hat{\mathbf{y}}) \cos(\omega_{\text{rf}}t), \quad (2.8)$$

where $\omega_{\text{rf}} = 2\pi \times 50$ kHz⁵, $R_0 \sim 5.36$ cm is the effective trap radius, and x and y are the laboratory-frame radial coordinates.

Using the trap-parameter convention adopted in our group and summarized in the appendix of Ref. [38], the nominal operating point of $f_{\text{rf}} = 50$ kHz, $V_{\text{rf}} = 25.6$ V, and $U_{ec,0} = 20$ V corresponds to Mathieu parameters of approximately $q = 0.141$ and $a = 0.002$. In this regime, the micromotion amplitude is small compared with the secular motion amplitude.

The rf confinement can therefore be described within the pseudopotential approximation. For an oscillating field of the form

$$\mathcal{E}_{\text{rf}}(\mathbf{r}, t) = \mathcal{E}_0(\mathbf{r}) \cos(\omega_{\text{rf}}t),$$

⁵ For some measurements, we operated at different values of ω_{rf} , but 50 kHz is the nominal value used for most of this thesis.

the fast micromotion of an ion with mass m and charge e is approximately

$$\mathbf{r}_{\text{micro}}(t) \approx -\frac{e}{m\omega_{\text{rf}}^2} \mathcal{E}_0(\mathbf{r}_{\text{sec}}) \cos(\omega_{\text{rf}} t),$$

where $\mathbf{r}_{\text{sec}}$ denotes the slowly varying secular motion. Then, averaging over one rf cycle, the kinetic energy of this micromotion acts as an effective potential,

$$U_{\text{pseudo}}(\mathbf{r}_{\text{sec}}) = \left\langle \frac{1}{2} m \dot{\mathbf{r}}_{\text{micro}}^2 \right\rangle = \frac{e^2}{4m\omega_{\text{rf}}^2} |\mathcal{E}_0(\mathbf{r}_{\text{sec}})|^2. \quad (2.9)$$

For the quadrupole rf field of the form Eq. (2.8), the pseudopotential becomes

$$U_{\text{pseudo}}(x, y) = \frac{e^2 V_{\text{rf},0}^2}{4m\omega_{\text{rf}}^2 R_0^4} (x^2 + y^2). \quad (2.10)$$

Therefore, near the trap center, the oscillating rf field produces an approximate harmonic radial confinement.

The secular frequencies are given

$$f_z = \frac{f_{\text{rf}}}{2} \sqrt{a_z}, \quad f_{x,y} = \frac{f_{\text{rf}}}{2} \sqrt{a_{x,y} + q^2/2} \quad (2.11)$$

using the conventions of a and q in Ref. [38]:

$$\begin{aligned} a_z &= 2 \frac{8e}{m\omega_{\text{rf}}^2} \frac{U_{\text{ec},0}}{Z_0^2}, \\ a_x &= -\left| \frac{8e}{m\omega_{\text{rf}}^2} \frac{U_{\text{ec},0}}{Z_0^2} \right| + \frac{8e}{m\omega_{\text{rf}}^2} \frac{V_{\text{DC},0}}{R_0^2}, \\ a_y &= -\left| \frac{8e}{m\omega_{\text{rf}}^2} \frac{U_{\text{ec},0}}{Z_0^2} \right| - \frac{8e}{m\omega_{\text{rf}}^2} \frac{V_{\text{DC},0}}{R_0^2}, \\ q^2 &= \left(\frac{4e}{m\omega_{\text{rf}}^2} \frac{V_{\text{rf},0}}{R_0^2} \right)^2 \end{aligned} \quad (2.12)$$

where again $U_{\text{ec},0}$ is the peak-to-peak axial voltage, $V_{\text{rf},0}$ is the amplitude of the rf voltage. The term $V_{\text{DC},0}$ is a static radial quadrupole voltage applied to the rods, which allows the two radial secular frequencies to be tuned relative to one another. For the nominal trap parameters of $V_{\text{rf},0} = 25.6$ V, $U_{\text{ec},0} = 20$ V, and $V_{\text{DC},0} = 0$ V, the corresponding secular frequencies are approximately 2 kHz along all three trap axes.

2.4.5 Rotating electric field

To polarize the molecular ions, we apply a rotating bias field in the radial plane,

$$\mathcal{E}_{\text{rot}}(t) = \mathcal{E}_{\text{rot}} [\cos(\omega_{\text{rot}}t)\hat{\mathbf{x}} \pm \sin(\omega_{\text{rot}}t)\hat{\mathbf{y}}], \quad (2.13)$$

where the sign specifies the rotation direction. For most of the measurements discussed in this thesis, $\omega_{\text{rot}} = 2\pi \times 147.5 \text{ kHz}$ ⁶ and the field magnitude is typically $\mathcal{E}_{\text{rot}} = 60 \text{ V/cm}$.

As mentioned previously, the molecular polarization follows the instantaneous field direction adiabatically. Therefore, the molecular quantization axis is set by $\hat{\mathcal{E}}_{\text{rot}}(t)$. The rotating field also drives the ion into a circular orbit in the radial plane, as shown in Fig. 2.2. This driven motion is given by

$$\mathbf{r}_{\text{rot}}(t) = -\frac{e}{m\omega_{\text{rot}}^2}\mathcal{E}_{\text{rot}}(t) = -r_{\text{rot}}\hat{\mathcal{E}}_{\text{rot}}(t), \quad (2.14)$$

so that the orbit radius is

$$r_{\text{rot}} = \frac{e\mathcal{E}_{\text{rot}}}{m\omega_{\text{rot}}^2}. \quad (2.15)$$

For $\omega_{\text{rot}} = 2\pi \times 147.5 \text{ kHz}$ and $\mathcal{E}_{\text{rot}} = 60 \text{ V/cm}$,

$$r_{\text{rot}} \approx 2.69 \text{ mm}.$$

The important point in Eq. (2.14) here is that the ion orbit is phase-locked to the rotating electric field; instantaneously, the molecular quantization axis follows $\hat{\mathcal{E}}_{\text{rot}}(t)$, and the ion displacement vector $\mathbf{r}_{\text{rot}}$ points in the opposite direction to $\hat{\mathcal{E}}_{\text{rot}}(t)$. As a result, quantities that co-rotate with the molecular axis can be expressed equivalently in terms of $\hat{\mathcal{E}}_{\text{rot}}(t)$ or in terms of $\hat{\mathbf{r}}_{\text{rot}}$, up to a sign convention. This will be the key to understanding how an axial magnetic field gradient produces an effective rotating-frame Zeeman shift in the next section.

2.4.6 Magnetic fields

Because the molecules are polarized by the rotating electric field, it is not straightforward to apply a magnetic field that also co-rotates with the molecular axis. One way to generate such

⁶ In the later development of the experiment, we increase the rotating-field frequency to about 300 kHz. Unless otherwise noted, however, the measurements discussed in this thesis were performed at 147.5 kHz.

an effective co-rotating (or counter-rotating) magnetic bias field is to apply an axial magnetic field gradient, using a pair of coils in an anti-Helmholtz configuration. The field generated by these coils has the quadrupole form

$$\mathbf{B}_{\text{axgrad}}(\mathbf{r}) \propto B_{\text{axgrad}}(2z\hat{\mathbf{z}} - x\hat{\mathbf{x}} - y\hat{\mathbf{y}}), \quad (2.16)$$

with the overall sign depending on the current direction.

Evaluated on the orbit $\mathbf{r}_{\text{rot}}(t)$, the radial component of $\mathbf{B}_{\text{axgrad}}$ co-rotates with the molecular quantization axis. The effective Zeeman shift is set by the gradient field projected onto the co-rotating axis,

$$B_{\text{rot}}^{\text{eff}} \propto \left\langle \mathbf{B}_{\text{axgrad}}(\mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathcal{E}}_{\text{rot}} \right\rangle. \quad (2.17)$$

Since $\hat{\mathbf{r}}_{\text{rot}}(t) = -\hat{\mathcal{E}}_{\text{rot}}$, equivalently, this is given as

$$B_{\text{rot}}^{\text{eff}} \propto \left\langle \mathbf{B}_{\text{axgrad}}(\mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \right\rangle. \quad (2.18)$$

Evaluating this for the quadrupole field given above, in the rotating frame, this field becomes

$$B_{\text{rot}}^{\text{eff}} \propto B_{\text{axgrad}} r_{\text{rot}}. \quad (2.19)$$

The details of the magnetic field will be discussed later in Chapter 7. In the remainder of this thesis, I express this idea in terms of the projection of $\mathbf{B}_{\text{axgrad}}$ onto $\hat{\mathbf{r}}_{\text{rot}}$, with the sign absorbed into the chosen convention of $\mathbf{B}_{\text{axgrad}}$.

Since the prototype apparatus has two main regions where measurements are performed, we use one set of B_{axgrad} coils near the octagon region and another set near the science region.

Currently, the prototype apparatus does not have any dedicated shim coils for independently tuning the magnetic field along the x , y , and z direction at the ion position. For the Gen. III apparatus, we plan to add shim coils in all three directions in addition to the B_{axgrad} coils. We also plan to deploy a constellation of magnetic field sensors to monitor the magnetic environment around the vacuum chamber.

2.4.7 Cryogenic system

To suppress the blackbody-radiation-induced excitation out of the quantum manifold in ThF^+ , we cool some part of the apparatus to approximately 160 K. The role of blackbody radiation excitation in limiting the population lifetime will be discussed further in Chapter 4.

Because the required temperature is modest compared with typical cryogenic experiments, the system is cooled using a single commercial cryocooler (Ametek CryoTel DS30) mounted on the far side of the vacuum chamber. In this region, the vacuum is surrounded by an aluminum “cold jacket” to improve thermal conduction along the vacuum tube, which is made of stainless steel.

To monitor the temperature along the trap, we originally planned to mount thermocouples along the spines. However, shortly before assembly, we realized that these particular thermocouples we intended to use were magnetic. To avoid introducing additional magnetic materials near the ions, we instead mounted the thermocouples on the outside of the vacuum chamber, farther from the trapping region.

More details on the thermal design and the experimental performance of the cryogenic system can be found in Ref. [38].

2.4.8 Control system

The experiment is controlled by ARTIQ, which serves as the central timing and sequencing system for the apparatus. The ARTIQ FPGA coordinates the timing of all the hardware components.

A digital delay generator (SRS DG645) provides the continuous 10 Hz trigger for the pulsed Nd:YAG lasers with the appropriate relative delays. One of the outputs from the digital delay generator goes to ARTIQ. When ARTIQ receives one of these triggers, it defines the start of the experimental sequence. From that point onward, the ARTIQ TTL system distributes the required timing signals to the rest of the hardware in the system.

All the dc voltages are generated by the ARTIQ DAC system. For the special and quadrant

rings, fast voltage pulses must also be applied. To do this, each such ring is supplied by two DAC outputs, and the fast pulse shapes are produced using high-speed switches triggered by the ARTIQ TTL signals.

The oscillating fields are generated using homebuilt DDS boards⁷ designed by Felix Viemeyer of the JILA electronics shop. The desired waveform sequence is first loaded onto the boards, and the DDS outputs begin when they receive an external trigger from ARTIQ.

More details on individual components of the control system can be found in Ref. [38, 59].

2.5 Gen. III prototype measurement

In this section, I describe how measurements are performed in the prototype apparatus. The goal is to introduce the experimental sequence and the Ramsey-fringe measurement procedure used throughout the rest of this thesis.

2.5.1 Experimental Sequence

This section describes the experimental sequence used for the Gen. III prototype measurement. The sequence was designed primarily for performing Ramsey spectroscopy after ion transport into the science region. In practice, however, the same script is also used for measurements without transport. This choice makes the experimental control flexible and keeps the timing logic consistent between different experimental configurations, which is an important feature in the prototype experiment, where we change the experiment sequences often.

The experimental sequence is shown in Fig. 2.14. The neutral molecular beam is produced as described in Sec. 2.4.3. For the present discussion, the experimental sequence begins at ionization. The main steps are as follows.

- (1) **Ionization, stopping, and trap turn-on.** The ionization lasers, the stopping pulse, and the rf trapping voltages are all turned on at the beginning of the sequence. The ions are not

⁷ These boards were originally developed for the Gen. II experiment, so some of their features are not used in the prototype apparatus.




Figure 2.14: Experimental sequence used for the Gen. III prototype Ramsey measurement. In the experimental block, “pulse” refers to the duration of each event, and “delay_A_B” refers to the delay between event A and B. The sequence begins with ionization and stopping, followed by trap-centering shim ramps, application of the rotating electric field, state preparation, ion transport to the science region, Ramsey $\pi/2$ pulses, state detection, and final ion extraction. The lower panel shows the timing structure of the two Ramsey $\pi/2$ pulses.

produced and stopped at the geometric center of the trap. Instead, their initial position depends on the ionization-laser position and the strength and timing of the stopping field. As a result, the newly created ion cloud is displaced from the rf null.

- (2) **Initial centering with quadrant-ring shims.** During the *rampShim* block, dc voltages on the quadrant rings are ramped in order to move the ion cloud toward the trap center. In this block, if the rotating electric field $\mathcal{E}_{\text{rot}}$ is used, then the quadrant-ring voltages are instead ramped to the values that center the ions, in the presence of $\mathcal{E}_{\text{rot}}$. This correction is necessary because the imperfections in the applied rotating field, along with patch charges and other static imperfections, shift the effective trap center. The goal of this block is therefore to place the ions near the rf null under the field conditions used for that particular measurement.
- (3) **Application of the rotating electric field and state preparation.** Once the ions are centered near the rf null, the rotating electric field is applied. The sequence then enters the *state-preparation* block, during which the relevant lasers are applied to prepare the ions in the desired internal state. The block usually lasts about 300 ms.
- (4) **Preparation for transport.** After state preparation, the sequence enters the *SecondTrap* block. The original purpose of this block is to increase the rf confinement before transporting the ions to the science region.
- (5) **Transport into the science region.** If the measurement uses ion transport, then the ions are moved to the science region during the *translate_in* block. For measurements without transport, this timing block still remains in the sequence, but the transport voltages are disabled.
- (6) **Trap settings for Ramsey spectroscopy.** After transport, the sequence enters the *ThirdTrap* block. In this block, both the radial and axial confinement are reduced in order to lower the ion density during Ramsey spectroscopy.

- (7) **Ramsey interrogation.** The Ramsey measurement is performed using two $\pi/2$ pulses separated by a free-evolution time. The lower panel of Fig. 2.14 shows the timing structure of these two Ramsey pulses.
- (8) **Return transport and state detection.** The second half of the sequence reverses the earlier operations. If transport was used, the ions are transported back from the science region back to the octagon. The *state-preparation* block is replaced by a *state-detection* block, during which the final molecular-state populations are measured.
- (9) **Reduction of $\mathcal{E}_{\text{rot}}$ before ion extraction.** The sequence then enters the *SecondErot* block. During this block, the magnitude of $\mathcal{E}_{\text{rot}}$ is reduced before ion extraction. In the prototype apparatus, this reduction was necessary because of limitations of the original rod-driver system.⁸ Although the rotating electric field is turned off during ion extraction, the original rod-driver system consisted of resonant circuits with slightly different quality factors. As a result, different rods ramped down at slightly different rates, giving the ions a large ponderomotive kick and preventing them from reaching the detector. Reducing the magnitude of $\mathcal{E}_{\text{rot}}$ before extraction mitigated this problem.
- (10) **Final recentering and ion extraction.** Because changing the magnitude of $\mathcal{E}_{\text{rot}}$ can also shift the effective trap center, a final *rampShim.REMPD* block is used to recenter the ions before the dissociation laser and extraction.

Several small delays are included between blocks. These delays were added because the first ARTIQ FPGA used in the prototype apparatus was an older version. Therefore, some events were not executed reliably when successive commands were placed too close together in time.

The same experimental script is used both for measurements with transport and for measurements without transport. Because the timing sequence is built by stacking blocks sequentially, a timing block remains present even when the corresponding operation is disabled. For example, in a measurement without transport, the *translate.in* block is still present in the timing sequence, but

⁸ As of December 2025, the rod-driver system has been upgraded, so this issue should no longer be present.

the transport voltages are not applied. Although this is not the most streamlined implementation, it was sufficient for the prototype apparatus and allowed the sequence to remain compatible with the existing control infrastructure. More importantly, it provided a flexible framework in which experimental blocks and control parameters could be added, removed, or disabled without creating many separate versions of the experimental script.

2.5.2 Ramsey fringe

For each Stark doublet, we detect the populations in the $|\downarrow\rangle$ and $|\uparrow\rangle$ manifolds as a function of the Ramsey time t_{Ramsey} . During state preparation, we can prepare the ions such that the population is mapped into the $|\uparrow\rangle$, which I refer to as the *in* phase configuration, or into the $|\downarrow\rangle$ manifold, which I refer to as *anti*.⁹

Similarly, during state readout (or state detection), the phase of the m_F depletion laser can also be chosen to be either *in* or *anti*. If the preparation and readout phases are chosen as *in-in*, then we effectively prepare ions in $|\uparrow\rangle$ before the first $\pi/2$ pulse and detect the population remaining in the $|\uparrow\rangle$ after the second $\pi/2$ pulse. We refer to this configuration as “In”. If the phases are chosen as *anti-in*, then we effectively detect the complementary spin manifold after the second $\pi/2$; we refer to this configuration as “Anti”.

From these two measurement, we form the spin asymmetry for each Stark doublet as [15, 51, 60]¹⁰

$$\mathcal{A}_{u/l} = \frac{N_{\text{In}}^{u/l} - N_{\text{Anti}}^{u/l}}{N_{\text{In}}^{u/l} + N_{\text{Anti}}^{u/l}}, \quad (2.20)$$

where the superscript u/l refers to the specific doublet and the subscript refers to the preparation and readout phase combination.

The resulting Ramsey fringes are shown in Fig. 2.15. We fit each doublet Ramsey fringe

⁹ In practice, *in* corresponds to the case in which $\mathcal{E}_{\text{rot}}$ is parallel to the k -vector of the m_F pumping laser, while *anti* corresponds to the antiparallel case, hence the names.

¹⁰ In a perfect world, we would do in-in and in-anti. However, the dissociation transition has different dissociation efficiency for a fixed polarization such that dissociating from $|\uparrow\rangle$ and $|\downarrow\rangle$ has different transition strength. For this reason, we use in-in and anti-in.

to [15, 51, 60]

$$C_{u/l} e^{-\gamma_{u/l} t_R} \cos(2\pi f_{u/l} t_R + \phi_{u/l}) + O_{u/l} \quad (2.21)$$

where C is the initial contrast, f is the frequency, t_R is the free evolution time, ϕ is the initial phase, O is the offset, and γ is some contrast decay rate for each doublet denoted by the subscript.




Figure 2.15: Example Ramsey fringe measurement for the upper and lower Stark doublets, showing the $|\uparrow\rangle$ and $|\downarrow\rangle$ states used in the Ramsey sequence. Left: Level structure of the ${}^3\Delta_1(v=0, J=1)$ manifold with external electric field $\mathcal{E}_{\text{rot}} \sim 60$ V/cm. Solid lines correspond to states with $\Omega = +1$, and the dashed lines to $\Omega = -1$. Right: measured spin asymmetry $\mathcal{A}_u$ (top) and $\mathcal{A}_l$ (bottom) as a function of Ramsey time t_{Ramsey} , with the fits shown in dashed lines. The fitted fringe frequencies are used to extract the frequencies for the two doublets.

Throughout this thesis, when Ramsey fringe is measured, this is the procedure it follows.

Chapter 3

Overview of Prototype Coherence Limits

3.1 Prototype goals

The prototype experiment was developed with two primary goals. First, we aimed to investigate the requirements for near-continuous-mode operations. This involved characterizing the electrode designs necessary for translating ions between regions and maintaining Ramsey coherence during transport. Our efforts towards the continuous-mode operation are detailed later in Chapter 8.

Second, the prototype served as a testbed for precision measurements in the Gen. III architecture. A central proof-of-principle goal was the experimental demonstration of long coherence times. The subsequent four chapters describe the investigations into the mechanisms currently limiting our coherence.

3.2 Categorizing Coherence Limits

A central motivation for the transition from HfF^+ to ThF^+ in the Gen. III experiment is that the eEDM-sensitive states in ThF^+ reside in the true electronic ground state. This is in contrast to the Gen. I and Gen. II HfF^+ experiments; in HfF^+ , the eEDM-sensitive states were metastable, which imposed a hard ceiling on the achievable coherence time due to spontaneous decay to about 2 seconds.

Because ThF^+ operates in the ground state, the interrogation time is not fundamentally limited by spontaneous emission, meaning we could theoretically interrogate the molecules indefinitely.

In practice, however, this interrogation (coherence) time is still limited by environmental and technical mechanisms. To fully leverage this ground-state molecule for precision measurements, we must be able to interrogate the ions for extended durations without losing either state population or phase coherence of the Ramsey superposition.

It is therefore useful to separate the observed coherence limitations into two broad categories: loss of count and loss of contrast. Loss of count refers to a reduction in the number of ions that remain detectable in the addressed states at the end of the experimental sequence. Loss of contrast refers to a reduction in the Ramsey fringe amplitude, i.e. contrast, among the ions that remain detectable. These two observables are experimentally distinct, although they can be produced by the same physical mechanisms.

3.2.1 Loss of count

The total loss of detectable signal, or “loss of count”, is determined by the rate at which ions exit our target quantum states. We perform Ramsey spectroscopy in the $^3\Delta_1, v = 0, F = 3/2, m_F = \pm 3/2$ states of ThF^+ . The total observed loss rate, Γ_{count} can be described with a sum of different contributions:

$$\Gamma_{\text{count}} = \Gamma_{\text{trap}} + \Gamma_{J=1} + \Gamma_{m_F} + \cdots, \quad (3.1)$$

where the terms here represent physically distinct ways in which ions can be removed from the detected state manifold. This expression is meant to explicitly show the relation for the different dominant contributions to the measured count decay.

Because these states exist within a nested physical architecture (the m_F states are within the $J = 1$ level, which is within the trapped-ion ensemble), we define these rates as follows:

- (1) Γ_{trap} (Trap loss): The rate at which ions physically escape the trap. This is generally independent of the internal quantum state.
- (2) $\Gamma_{J=1}$ (Rotational level loss): The rate at which ions transition out of the $J = 1$ rotational level, normalized to the trap loss. In other words, this is the rate at which ions move to a

different rotational state while remaining trapped.

- (3) Γ_{m_F} (Zeeman sublevel loss): The rate at which ions are redistributed out of the $m_F = \pm 3/2$ states, normalized to the $J = 1$ loss. This accounts for m_F mixing within the $J = 1$ level that does not necessarily result in a change of J .

By defining the rates this way, we can attribute specific experimental observations to distinct physical causes. For example, improving the trap vacuum might reduce Γ_{trap} , and this would result in overall improvement in Γ_{count} . Conversely, if we decrease the magnitude of the bias electric field, we might see an increase in Γ_{m_F} due to ion-ion collision, even if $\Gamma_{J=1}$ and Γ_{trap} remain constant. The specific mechanisms governing these rates are detailed in Chapters 4 and 5.

3.2.2 Loss of contrast

The experimental coherence time is also limited by the loss of Ramsey fringe contrast. In this case, the ions may remain in the detected state manifold, but the ensemble no longer supports high-contrast Ramsey fringe. This decay can arise from either the ensemble dephasing or decoherence. We write the corresponding contrast-loss rate as

$$\Gamma_{\text{contrast}} = \Gamma_{\text{inhomogeneity}} + \Gamma_{\text{Berry}} + \Gamma_{\text{ion-ion}} \cdots, \quad (3.2)$$

where $\Gamma_{\text{inhomogeneity}}$ is the rate of contrast loss from magnetic and electric field inhomogeneity, Γ_{Berry} is the rate of contrast loss associated with Berry-like phase accumulation, and $\Gamma_{\text{ion-ion}}$ due to ion-ion collisions. As in Eq. (3.1), this expression is intended as a useful decomposition of the observed contrast decay, and it is not intended to show that the mechanisms are mutually exclusive.

More generally, Ramsey contrast is reduced whenever different ions acquire different phases during the free-precession time, or when the phase acquired by a given ion becomes sensitive to uncontrolled experimental variables. Spatial field inhomogeneity is one example: if the applied electric or magnetic fields are not perfectly homogeneous over the ion cloud, then ions with different energies and trajectories sample different time-averaged effective fields. This results in a spread

in Ramsey precession frequencies and therefore reduces the ensemble-averaged contrast. Temporal field instabilities can also produce a similar effect.

Ion-ion collisions also provided another important contrast-loss mechanism in the prototype. During such a collision, an ion trajectory can change, modifying the geometric phase accumulated by that ion. Since the collisions occur stochastically, different ions acquire different phases over time, which can reduce Ramsey contrast. At the same time, however, ion-ion collisions can be beneficial in some operating regimes. If collisions randomize the ion trajectories sufficiently, they can make the motion more ergodic and produce motional narrowing, helping the ensemble average over some static spatial inhomogeneities. These contrast-loss mechanisms are discussed in more detail in Chapter 7.

3.3 Mechanisms vs Effects: A Structural Note

In the following four chapters, I have organized the discussion by the observable *effect* (count vs. contrast), rather than by *mechanism* (e.g., collisions, field homogeneity). I am structuring the discussion this way to emphasize our current understanding of the prototype’s operational limits: our loss of count is predominantly limited by one set of mechanisms, while our loss of contrast is limited by another. However, the reader should keep in mind that a single physical mechanism can, and often does, have multiple effects across the experiment. For instance, our current understanding is that ion-neutral collisions are the primary limitation on the $J = 1$ manifold lifetime in the prototype. Therefore, I will discuss ion-neutral collisions primarily in the context of $J = 1$ lifetime. This *does not* mean that ion-neutral collisions have zero effect on the $m_F = \pm 3/2$ state decay rate, the overall ion trap lifetime, or the loss of contrast. The categorization is meant to highlight the dominant effect of each mechanism as we currently understand it in the prototype architecture.

3.4 Roadmap of the coherence studies

In the next four chapters, I describe the present coherence limits of the prototype experiment, using the loss of count and loss of contrast framework introduced here. Each chapter focuses on

one experimentally important observable effect.

Chapter 4 focuses on loss from the internal quantum manifold. For the $J = 1$ rotational level, the measured lifetime is partly explained by blackbody-radiation excitation, but the data suggest that this mechanism is not sufficient by itself. Ion–neutral collisions, especially with H_2 , appear to provide an additional contribution to the observed $J = 1$ loss, although the interpretation of the baseline H_2 pressure requires caution. I also discuss loss from the addressed $m_F = \pm 3/2$ states. The observed dependence on ion number, trap confinement, and rotating electric-field magnitude indicate that ion–ion collisions are an important limitation on the m_F -manifold lifetime. This motivates lowering the ion density during Ramsey evolution.

Chapter 5 then discusses the consequence of reducing the ion density. In the prototype, reducing the trap confinement also reduces the trap depth, and as a result ions become more vulnerable to motional heating and can spill out of the trap before the desired interrogation time is reached. Several mechanisms can contribute to this heating or trap loss, including the large ponderomotive energy associated with the rotating electric field, rf-driven heating, and ion–neutral collisions. Differential measurements show that some of these mechanisms can measurably affect the trap lifetime. However, systematic uncertainties in the thermometry, especially its dependence with ion number, prevent us from assigning the present heating rate to a single dominant mechanism definitively.

Chapter 6 takes a complementary approach—rather than requiring a complete understanding of the heating, we investigate different trapping configurations that can be more tolerant to the increasing motional energy. These alternative trapping configurations increase the effective trap depth while maintaining relatively loose confinement. In particular, the octupole and rotating-rf trap configurations are investigated as ways to reduce ion loss at long trapping times.

Finally, Chapter 7 turns from count loss to contrast loss. Here, I discuss how temporal fluctuations and spatial inhomogeneities of the rotating electric and magnetic fields produce a distribution of accumulated Ramsey phases across the ion ensemble. The chapter also discusses ion–ion collisions, in the context of driving motional narrowing and averaging over static field

inhomogeneities.

Together, these chapters identify the main practical requirements for extending the Ramsey interrogation time in Gen. III prototype apparatus: suppressing rotational-state loss from blackbody radiation and residual gas collisions, reducing ion density enough to limit m_F -changing ion–ion collisions, maintaining sufficient effective trap depth as the confinement is weakened, and controlling field inhomogeneities well enough that the surviving ions retain Ramsey contrast.

Chapter 4

Towards Long Coherence Time: Loss of Signal from Quantum Manifold

A fundamental prerequisite for achieving extended coherence times is the ability to interrogate the molecular ensemble without significant loss of detectable signal. As mentioned, in practice, signal can be reduced in two distinct ways: ions may be physically lost from the trap, or population may be transferred out of the quantum states used for the measurement. In this chapter, I discuss the dominant mechanisms in the latter. The mechanisms discussed here reduce signal by reducing the fraction of the ions that remains in the target measurement states.

Building on Eq. (3.1), I begin with mechanisms that transfer population between the target $m_F = \pm 3/2$ Zeeman sublevels, Γ_{m_F} , with an emphasis on ion-ion interactions. Then I discuss mechanisms that drive population out of the $J = 1$ rotational manifold, $\Gamma_{J=1}$, such as blackbody-radiation-driven redistribution.

4.1 m_F manifold lifetime

We perform Ramsey spectroscopy between the two stretched Zeeman sublevels $m_F = \pm 3/2$ within the $J = 1$ manifold. For $\mathcal{E}_{\text{rot}} = 60$ V/cm, the dominant internal splitting within this manifold is set by the Stark splitting with the rotating electric field. Consequently, any mechanism that rapidly perturbs the local electric-field magnitude or direction can spoil the Ramsey evolution. In this section, I discuss mechanisms that can shorten the lifetime of the m_F manifold.

Before discussing specific mechanisms, it is useful to distinguish between two broad ways in which they can affect the measurement. First, a mechanism can change the internal state popu-

lations relevant to the Ramsey measurement. In particular, it can redistribute population among the m_F sublevels, including transfer out of the addressed $m_F = \pm 3/2$ Ramsey subspace. Such state-changing processes reduce the fraction of ions that participate in the intended Ramsey evolution and can therefore reduce the raw signal amplitude (i.e. reducing signal-to-noise ratio). More generally, they can also contribute directly to decay of the observed Ramsey contrast, because they interrupt the intended internal-state evolution. Second, a mechanism can leave the m_F populations unchanged while modifying the accumulated phase, producing dephasing without population transfer. In this section, I focus on the first category, while the latter is discussed in Chapter 7.

4.1.1 Ion-ion collisions

Because the ions are charged, each ion produces a Coulomb field that perturbs the local quantization axis seen by nearby ions. In the weak and slowly varying limit, this perturbation mainly produces an additional geometric phase called Berry's phase and therefore acts as a dephasing mechanism. Here, I focus on the opposite limit of close, relatively violent encounters. When the Coulomb field from one ion becomes comparable to $\mathcal{E}_{\text{tot}}$ on a timescale that is not slow compared with the relevant m_F splittings, the local quantization axis can be tilted nonadiabatically. In this strong-collision limit, a single close encounter can redistribute population among the Zeeman sublevels. A schematic illustration of this strong-collision picture is shown in Fig. 4.1.

As a crude order-of-magnitude estimate, we consider the classical distance of closest approach r_c between two singly charged ions given the temperature of the ions T . This distance is obtained by equating the Coulomb energy to a characteristic collision energy

$$E_{\text{coll}} \sim 2 \left(\frac{3}{2} k_B T \right) = 3 k_B T.$$

This gives

$$r_c \sim \frac{1}{4\pi\epsilon_0} \frac{e^2}{3k_B T} \approx 0.19 \, \mu\text{m} \left(\frac{30 \, \text{K}}{T} \right), \quad (4.1)$$

where ϵ_0 is the vacuum permittivity, e is the elementary charge, and k_B is the Boltzmann constant.




Figure 4.1: Schematic of ion–ion induced redistribution among the m_F sublevels. In a sufficiently close collision, the Coulomb field from a neighboring ion can become comparable to the rotating bias field $\mathcal{E}_{\text{rot}}$, tilting the local quantization axis. In this nonadiabatic limit, a single collision can transfer population of the initial m_F state and therefore contribute directly to loss of signal.

Up to such numerical constants, this is essentially the classical Coulomb distance of closest approach, i.e. the Landau length. [61, 62]

At this separation, the electric field from one ion at the location of the other is given by

$$\mathcal{E}_{\text{ion-ion}}(r_c) = \frac{e}{4\pi\epsilon_0} \frac{1}{r_c^2} = \frac{36\pi\epsilon_0}{e^3} (k_B T)^2 \approx 420 \text{ V/cm} \left(\frac{T}{30 \text{ K}} \right)^2. \quad (4.2)$$

For the ion temperature $T \sim 30 \text{ K}$, this field is already larger than the operating value of $\mathcal{E}_{\text{rot}} = 60 \text{ V/cm}$. In this simple estimate, the Coulomb field becomes comparable to the rotating bias field even before the ions reach their classical distance of closest approach.

The collision timescale is also important. Taking a relative thermal speed between the two ions, $v_{\text{th,rel}} \sim \sqrt{12k_B T/m}$, the inverse interaction timescale is roughly

$$\tau_c^{-1} \approx \frac{v_{\text{th}}}{r_c} \sim 6 \times 10^8 \text{ s}^{-1} \left(\frac{T}{30 \text{ K}} \right)^{3/2}. \quad (4.3)$$

For $T \sim 30 \text{ K}$, the interaction occurs on a rate of a few 100 MHz, which is not slow compared to the energy splitting between different m_F states. These simple estimates suggest that close ion-ion collisions can indeed drive transitions among the m_F states.

To estimate the rate of such collisions, we define a critical impact parameter b_c as a distance at which the Coulomb field during the encounter becomes comparable to the rotating field

$$\frac{e}{4\pi\epsilon_0 b_c^2} \sim \mathcal{E}_{\text{rot}}. \quad (4.4)$$

This gives the collision cross section with the scaling

$$\sigma_c \sim \pi b_c^2 \propto \mathcal{E}_{\text{rot}}^{-1}. \quad (4.5)$$

Using the hard-sphere estimate of

$$\Gamma_{\text{coll}} \sim n \sigma_c v_{\text{th}} \propto n \frac{\sqrt{T}}{\mathcal{E}_{\text{rot}}}. \quad (4.6)$$

where n is the ion density. Using the thermal distribution in a three-dimensional harmonic, this crude estimate gives a close-collision rate of order

$$\Gamma_{\text{coll}} \approx 3 \times 10^{-2} \text{ s}^{-1} \left(\frac{N}{2000} \right) \left(\frac{30 \text{ K}}{T} \right) \left(\frac{f_{\text{sec}}}{1 \text{ kHz}} \right)^3 \left(\frac{60 \text{ V/cm}}{\mathcal{E}_{\text{rot}}} \right). \quad (4.7)$$

For the nominal settings

$$\Gamma_{\text{coll}} \sim 10^{-2} - 10^{-1} \text{ s}^{-1}$$

per ion.

This scaling can be tested experimentally in two ways. First, after optical pumping, we adiabatically change the magnitude of $\mathcal{E}_{\text{rot}}$ and measure the decay rate of the Ramsey contrast. Strictly speaking, the measured quantity is a Ramsey decoherence rate, rather than a pure m_F -manifold lifetime. The Ramsey decoherence rate does include both population loss from the target manifold and any other additional dephasing¹. However, under the conditions of these measurements, the observed dependence on $\mathcal{E}_{\text{rot}}$ is consistent with strong-collision-driven redistribution being a major contribution to the decoherence rate. A more direct probe of the manifold lifetime would be the decay of the SPAM contrast measured without the $\pi/2$ pulses.

Figure 4.2 shows the observed Ramsey decoherence rate as a function of $\mathcal{E}_{\text{rot}}$. The observed decrease in the decoherence rate with increasing $\mathcal{E}_{\text{rot}}$ is qualitatively consistent with the $\Gamma_{\text{coll}} \propto 1/\mathcal{E}_{\text{rot}}$ scaling expected from model mentioned above. The dashed line marks the approximate lower bound set by the blackbody-radiation-driven redistribution of population, which is discussed in the next section.

A second test is obtained by adiabatically varying the trap confinement by changing secular frequency during the Ramsey evolution. For an adiabatic change of a harmonic trap, the temperature scales as

$$T \propto \omega_{\text{sec}},$$

as expected for a classical ensemble in a harmonic potential [63]. At the same time, the cloud size scales as² $r_{\text{cloud}} \sim \sqrt{T}/\omega_{\text{sec}} \propto \omega_{\text{sec}}^{-1/2}$, so the density scales as

$$n \propto \omega_{\text{sec}}^{3/2}.$$

¹ Strictly speaking, loss from the target manifold results in loss of contrast only when the lost population is detected, for instance, when the ions have moved to a different m_F level and are detected at the end of the Ramsey sequence

² In this strong-collision limit, I am assuming that we are in the thermal regime and the size of the cloud changes with the temperature.




Figure 4.2: Measured Ramsey decoherence rate as a function of the rotating electric field magnitude $\mathcal{E}_{\text{rot}}$. The dashed line indicates the approximate blackbody-radiation floor discussed in the next section. The decoherence rate decreases with increasing $\mathcal{E}_{\text{rot}}$, consistent with the expectation that stronger rotating fields suppress the fraction of ion-ion encounters that can strongly perturb the quantization axis. The data are taken with $f_{\text{rot}} = 147.5$ kHz, $N \sim 4000$, and $f_{\text{sec}} \sim 2$ kHz in all three directions.

Combining these scales, we get

$$\Gamma_{\text{coll}} \propto \frac{n\sqrt{T}}{\mathcal{E}_{\text{rot}}} \propto \frac{\omega_{\text{sec}}^2}{\mathcal{E}_{\text{rot}}}. \quad (4.8)$$

The measured secular-frequency dependence is shown in Fig. 4.3, after subtracting the estimated blackbody-radiation contribution. In all three datasets, the remaining decoherence rate grows approximately quadratically with ω_{sec} , as expected. Fitting the blackbody-subtracted data to $\Gamma = af_{\text{sec}}^2$ gives the efficient listed in Table 4.1. At fixed $\mathcal{E}_{\text{rot}} = 60$ V/cm, the high-ion-number dataset has a fitted coefficient larger than the low-ion-number dataset by a factor of about 3.7, which roughly corresponds to the difference in the number of ions in the data. At fixed high ion number, reducing $\mathcal{E}_{\text{rot}}$ from 60 to 30 V/cm increases the fitted coefficient by a factor of about 2. Both comparisons are consistent with the expected scaling of $\Gamma_{\text{coll}} \propto \omega_{\text{sec}}^2/\mathcal{E}_{\text{rot}}$.

Dataset in Fig. 4.3	Approx. Ion number	$\mathcal{E}_{\text{rot}}$ (V/cm)	a
Orange	~ 3000	60	$1.1(0.2) \times 10^{-7}$
Blue	~ 1000	60	$0.3(0.06) \times 10^{-7}$
Green	~ 3000	30	$2.2(0.2) \times 10^{-7}$

Table 4.1: Quadratic fit parameters for the blackbody-radiation-subtracted secular-frequency data, using the fit form $\Gamma = af_{\text{sec}}^2$. The ion number listed here is an approximate detected ion number, intended mainly to distinguish the high- and low-density conditions.

These fits should be interpreted differentially rather than absolutes, since the detected number is only a proxy for the number of trapped ions and the simple hard-sphere model neglects the full details of Coulomb scattering.

Nonetheless, the combined $\mathcal{E}_{\text{rot}}$, ion-number, and secular-frequency trends strongly suggest that ion-ion collisions are a significant contribution to the observed m_F -manifold decay. This is why it is necessary to ramp down the trap confinement during Ramsey evolution: reducing the density suppresses the collision rate and therefore allows for long interrogation time.

4.1.2 Ion-neutral collisions

As a separate test, we investigated whether collisions with neutral background gas could limit the lifetime of the m_F manifold. To test this, we intentionally flowed in additional N_2 into




Figure 4.3: **Top:** Ramsey decoherence rate as a function of secular frequency for three operating conditions. High number at $\mathcal{E}_{\text{rot}} = 60$ V/cm, low ion number at $\mathcal{E}_{\text{rot}} = 60$ V/cm, and high ion number at $\mathcal{E}_{\text{rot}} = 30$ V/cm. **Bottom:** the ion number. The blackbody-radiation contribution has been subtracted from all points. Solid curves are quadratic fits, illustrating the approximate $\Gamma_{\text{coll}} = a\omega_{\text{sec}}^2$ scaling expected for ion-ion collisions under adiabatic compression. The fitted coefficients a are reported in Table 4.1.

the vacuum chamber and monitored the SPAM contrast.

Within our sensitivity, we found no clear evidence that collisions with the neutral background gas are an important limitation of the m_F lifetime under normal operating conditions.

4.2 $J = 1$ state lifetime

The relevant m_F states used for Ramsey spectroscopy reside inside the $X^3\Delta_1(v = 0, J = 1)$ manifold. Even if the coherence within the selected m_F subspace is preserved, any process that removes population from the $J = 1$ manifold shortens the useful interrogation time. In the prototype apparatus, we have identified two most relevant mechanisms for the $J = 1$ population lifetime. The first is blackbody-radiation (BBR) excitation of the $X^3\Delta_1(v = 0)$. The second is inelastic collision with residual neutral gas. In this section, I discuss these mechanisms and their observed effect on the rotational state lifetime.

4.2.1 Blackbody radiation excitation

A natural contribution to the $J = 1$ lifetime is BBR-driven excitation out of the $X^3\Delta_1(v = 0)$ manifold. In ThF^+ , the dominant room-temperature BBR contribution is the excitation to the first excited vibrational manifold $X^3\Delta_1(v = 1)$, with a weaker additional contribution from the excitation to $a^1\Sigma^+$ [33]. Purely rotational BBR transitions within the $X^3\Delta_1(v = 0)$ are much slower on the timescales relevant here, and therefore will not be discussed.

For an incoherent thermal radiation field, the rate from the lower state l to the upper state u can be written in terms of the Einstein A coefficient as [64, 65]

$$\Gamma_{l \rightarrow u}^{\text{BBR}} = \frac{g_u}{g_l} A_{u \rightarrow l} \bar{n}(\omega_{ul}, T), \quad (4.9)$$

where

$$\bar{n}(\omega, T) = \frac{1}{e^{\hbar\omega/(k_B T)} - 1} \quad (4.10)$$

is the thermal photon occupation number, and g_l and g_u are the degeneracies of the two levels. For

an electric-dipole allowed transition, the spontaneous emission rate A is [66]

$$A_{u \rightarrow l} = \frac{\omega_{ul}^3}{3\pi\epsilon_0\hbar c^3} |\langle l | \mathbf{d} | u \rangle|^2. \quad (4.11)$$

With this equation, we can directly connect the BBR excitation rate to the measured radiative lifetimes.

For ThF^+ , the required spectroscopic constants were measured and analyzed in Ref. [33]. This work also predicts that the effective population lifetime of $X^3\Delta_1(v=0)$ is limited to only a few seconds at room temperature, but increases rapidly as the environment is cooled. Table 4.2 is from Ref. [33] and has calculations of blackbody radiation excitation lifetime out of the $^3\Delta_1, v=0$ manifold at different temperatures.

Radiative environment temperature (K)	77	120	150	180	200	300
To $X^3\Delta_1(v=1)$ (s)	32000	400	84	29	17	3.5
To $a^1\Sigma^+$ (s)	2200	270	130	80	64	33
Combined (s)	2100	160	51	21	13	3.2

Table 4.2: Predicted blackbody-radiation excitation lifetime out of the $X^3\Delta_1(v=0)$ manifold of ThF^+ , reproduced from Ref. [33].

We anticipated this BBR contribution before the construction of the prototype apparatus. We designed the cryogenic system for the prototype to reduce the effective radiative temperature seen by the ions. More details on the prototype’s cryogenics system³ are in Ref. [38].

The prototype apparatus was not designed for multiplexed operation, and as a result the temperature was intentionally designed to be not uniform along the trap. The octagon region, where optical access is required, is exposed to room-temperature surroundings (~ 300 K). The other end of the trap is connected to the commercial cryocooler and is at ~ 150 K. As a result, there is a substantial temperature gradient along the apparatus.

To test the BBR contribution directly, we measured the $J=1$ lifetime at several locations along the prototype. Figure 4.4(a) shows the relevant prototype geometry, and Figure 4.4(b) shows

³ The prototype cryogenic system is essentially spray foam that people have in their attics. I have very fond memories of going to Home Depot in Boulder to purchase many boxes of spray foam and making the mold in my advisor’s garage as a young graduate student.

the measured temperature profile on the outer vacuum chamber. From this profile, we estimate an effective temperature at each trapping location by integrating over the solid angle. This effective temperature is then used to calculate the expected BBR-limited lifetime of $J = 1$ level.

Experimentally, we first optically pump the ions into the $J = 1$ manifold. Then we measure the remaining population with state-selective dissociation on the $R(1)$ line after waiting for different amounts of trapping time. The comparison between the calculated and measured lifetimes is shown in Fig. 4.5.

The measured lifetimes follow the spatial trend expected from the BBR model, but they are systematically shorter than the calculated values. This suggests that the model is capturing the correct temperature dependence of the BBR contributions but with an additional process that is position-independent (i.e. temperature independent) limiting the $J = 1$ lifetime.

4.2.2 Ion-neutral collisions

The remaining limitation is consistent with inelastic collisions between the trapped ions and residual neutral gas. In the classical capture picture of Gioumouzis and Stevenson, the ion interacts with a neutral species of isotropic polarizability α [67]. In this model, the interaction is the ion-induced-dipole potential⁴

$$V(r) = -\frac{C_4}{r^4}, \quad C_4 = \frac{\alpha e^2}{8\pi\epsilon_0}, \quad (4.12)$$

where α is written as a polarizability volume in SI units. We calculate the critical impact parameter b_c and the Langevin cross section σ_L for a collision energy E_{coll} [67]


$$b_c = \left(\frac{2C_4}{E_{\text{coll}}} \right)^{1/4}, \quad \sigma_L = \pi b_c^2 = \pi \sqrt{\frac{2C_4}{E_{\text{coll}}}}.$$

This gives the familiar Langevin rate constant k_L :


$$k_L = \sigma_L v = e \sqrt{\frac{\pi\alpha}{\epsilon_0\mu}}, \quad (4.13)$$

where $E_{\text{coll}} = \frac{1}{2}\mu v^2$ was used and μ is the ion-neutral reduced mass. Notably, the Langevin rate constant k_L is independent of the collision energy. The corresponding per-ion collision rate at

⁴ Here, I'm assuming a singly charged positive ion.



(a)



(b)

Figure 4.4: (a) CAD drawing of the prototype apparatus showing the relevant geometry for the radiative environment seen by the ions. (b) Measured temperature profile along the outside of the vacuum chamber, and the points are extrapolated linearly in between the plots, according to our thermal model. The temperature gradient along the apparatus produces a position-dependent effective radiative temperature.




Figure 4.5: Measured $J = 1$ lifetime at several positions along the prototype apparatus compared with the lifetime expected from the calculated BBR environment. The measured lifetimes follow the predicted trend but are systematically shorter, indicating an additional loss mechanism.

neutral number density n_{gas} is

$$\Gamma_L = k_L n_{\text{gas}} = k_L \frac{P}{k_B T}, \quad (4.14)$$

where I have used the ideal-gas relation for the last equality. This Langevin model by itself does not predict the rotationally inelastic collision rate, but it sets the rate of close encounters.

More intuitively, during a sufficiently close encounter between the ion and the neutral gas, the neutral gas applies a transient torque to the ion. A simple schematic illustration of the rotational level changing ion-neutral collision is shown in Fig. 4.6. As the neutral gas passes by, its induced dipole p provides a time-dependent electric field $\mathcal{E}_{\text{coll}}(t)$ at the ion:

$$\mathcal{E}_{\text{coll}} \sim \frac{1}{4\pi\epsilon_0} \frac{2p}{b_c^3}, \quad p \sim 4\pi\epsilon_0\alpha\mathcal{E}_{\text{ion}}, \quad \mathcal{E}_{\text{ion}} = \frac{e}{4\pi\epsilon_0 b_c^2}, \quad (4.15)$$

where $\mathcal{E}_{\text{ion}}$ is the electric field that induced the dipole moment p . This field couples to the molecular-frame dipole moment through the Stark interaction

$$H_{\text{int}}(t) = -\mathbf{d}_{\text{mf}} \cdot \mathcal{E}_{\text{coll}}(t),$$

which can couple neighboring rotational levels.

Here, to assess the strength of such a collision, we look at two dimensionless quantities

$$\beta \sim \frac{d_{\text{mf}}\mathcal{E}_{\text{coll}}}{\hbar\omega_J}, \quad \omega_J T_{\text{pass}}, \quad (4.16)$$

where ω_J is the relevant rotational splitting and T_{pass} is the duration of the close encounter. The first quantity compares this transient Stark interaction to the rotational splitting. When $\beta \gtrsim 1$, the interaction strength is large enough to mix neighboring rotational levels. The second quantity compares the collision duration to the rotational timescale and determines whether the collision is sudden or adiabatic.

Classically, the transient field exerts a torque $\tau \sim d_{\text{mf}}\mathcal{E}_{\text{coll}}$ and delivers an angular-momentum impulse $\Delta L \sim d_{\text{mf}}\mathcal{E}_{\text{coll}}T_{\text{pass}}$. In the sudden-collision limit of $\omega_J T_{\text{pass}} \ll 1$, the collision acts like a broadband electric field pulse. Then the relevant order-of-magnitude parameter becomes the accumulated phase

$$\phi_{\text{coll}} \sim \frac{1}{\hbar} \int d_{\text{mf}}\mathcal{E}_{\text{coll}}(t)dt \sim \frac{d_{\text{mf}}\mathcal{E}_{\text{coll}}T_{\text{pass}}}{\hbar}. \quad (4.17)$$

When this phase is small $\phi_{\text{coll}} \ll 1$, the rotational-state mixing and thus transfer probability is small. When $\phi_{\text{coll}} \sim 1$, the collision can efficiently mix rotational states.

For an order-of-magnitude estimate, I take the collision speed to be set by the rotating-field motion of the ions, rather than by the thermal speed of the neutral gas. For N_2 , we have

$$b_c \sim 2.3 \text{ \AA}, \quad \mathcal{E}_{\text{coll}} \sim 7 \times 10^7 \text{ V/cm},$$

using the α of N_2 in Table 4.3 and $v \sim 2500 \text{ m/s}^5$. Using $d_{\text{mf}} = 3.37 \text{ D}$ and $\omega_j \approx 2\pi \cdot 30 \text{ GHz}$ [33],

$$\frac{d_{\text{mf}} \mathcal{E}_{\text{coll}}}{\hbar \omega_j} \gg 1, \quad T_{\text{pass}} \omega_j \ll 1.$$

Therefore, in this crude estimate, the collision field is easily strong enough to mix rotational levels and the corresponding pulse is sudden from the ion's point of view. Then the accumulated phase during the collision

$$\phi_{\text{coll}} \sim \frac{d_{\text{mf}} \mathcal{E}_{\text{coll}} T_{\text{pass}}}{\hbar} \sim 70, \quad (4.18)$$

again indicating that this close encounter can easily drive rotational redistribution. Repeating the same calculation for H_2 also gives similar order-of-magnitude estimates.

These estimates suggest that such close encounters with variety of different gases are strong enough to drive rotational redistribution. Given this, now we need to distinguish between these gases by looking at the frequency of such encounters, i.e. the collision rate between the ions and the gas.

I use the phrase “Langevin-like” collisions to refer to an effective state-changing collision rate expressed relative to the classical Langevin rate.⁶ Even though the Langevin rate does not describe the exact rotationally inelastic collision rate, it is still a useful parameter as the Langevin rate sets the close-encounter collision rate.

Table 4.3 lists Langevin rate constants for several relevant gases. Figure 4.7 shows the corresponding collision rates Γ_L as a function of our operating pressure. Around 10^{-9} Torr , the

⁵ This is the rotating-field velocity of the ions at 60 V/cm with 147.5 kHz.

⁶ The Langevin collision rate usually describes a more violent process, like chemical reactions. However, only a little perturbation is enough to change the rotational level.




Figure 4.6: Schematic picture of a rotationally inelastic ion–neutral collision. The neutral gas passing by the ion is polarized by the ion. The induced dipole generates a transient electric field at the ion. This field can mix neighboring rotational levels.

classical Langevin collision rates are already on the order of 10^{-2} – 10^{-1} s $^{-1}$, comparable to the additional decay channel inferred from the lifetime measurements.

Gas	α (Å 3)	k_L (10^{-9} cm 3 /s)	$\Gamma_L(P = 10^{-9}$ Torr) (s $^{-1}$)
H $_2$	0.787 [68]	1.47	0.047
N $_2$	1.710 [68]	0.61	0.020
SF $_6$	4.490 [68]	0.52	0.017
Ne	0.381 [68]	0.33	0.011
Kr	2.484 [69, 70]	0.466	0.015

Table 4.3: Langevin rate constants k_L for ThF $^+$ colliding with selected neutral gases at 300 K, using the polarizabilities from Ref. [68]. The collision rate at 10^{-9} Torr is computed from $\Gamma_L = k_L P / k_B T$.

To experimentally test the collisional hypothesis, we added a residual-gas analyzer (RGA) and a leak valve to the main chamber. We also calibrated the ion-gauge and RGA pressure readings; this procedure is described in Appendix A. We then measured the $J = 1$ decay rate while introducing known amounts of selected gases.

Figure 4.8 shows two datasets (a) and (b). In Fig. 4.8(a), we measure the bare $J = 1$ lifetime: the ions are optically pumped into $J = 1$, held for different amounts of trapping time, and then dissociated on the $R(1)$ line. In Fig. 4.8(b), the optical pumping lasers stay on during the entire trapping time. The optical pumping lasers include rotational repumps that return population from $J = 2$ and $J = 3$ back to $J = 1$ as well as vibrational repumps that return population from $v = 1, J = 1, 2$ back to the main manifold. Therefore, the second dataset suppresses the BBR contribution and collisional loss through channels that are repumped. In other words, the it is sensitive to collisional channels that remove population outside the repumped manifold.

In both panels, the first points correspond to the baseline vacuum, i.e. the partial pressure present in the vacuum without deliberately adding gas. The decay rate increases with added neutral-gas pressure, confirming that ion–neutral collisions can change the rotational level and limit the $J = 1$ lifetime. However, the decay rate responds differently with the added hydrogen and nitrogen. When additional H $_2$ is added, the decay rate rises immediately. By contrast, when N $_2$ is added, the rate remains nearly unchanged at first and only increases after the gas pressure has






Figure 4.7: Estimated Langevin close-encounter rate Γ_L for ThF^+ with H_2 , N_2 , Ne , Kr , and SF_6 as a function of partial pressure, assuming $T = 300$ K. This Langevin rate does not determine the rotationally inelastic collision rate, but it sets the rate of close ion–neutral encounters that can drive rotational-state-changing collisions.



(a)



(b)

Figure 4.8: Pressure-dependence $J = 1$ decay rate with added N_2 and H_2 . (a) Bare $J = 1$ lifetime measurement. (b) Measure with the optical pumping and repump lasers left on during the hold time, which suppresses the BBR contribution and collisional loss through channels that are repumped. The first points in both plots correspond to the baseline pressures of each gas.

been at least doubled compared to the baseline pressure. This suggests that hydrogen is a plausible dominant contributor to the baseline collisional loss in the present apparatus.

The dashed curves in Fig. 4.8 compare the data to empirical functional form

$$\Gamma_{\text{decay}}(P) = \Gamma_0 + \eta\Gamma_L(P),$$

where Γ_0 is an offset and η is a “Langevin-like” prefactor. In the bare-lifetime data of Fig. 4.8(a), the offset Γ_0 is fixed to the BBR contribution. The optical-pumping-on data in Fig. 4.8(b), the residual offset is much smaller, consistent with the suppression of the BBR contribution and of collisional channels that are efficiently repumped. The similarity between this residual offset and the additional offset inferred from the BBR measurements in Fig. 4.5 strongly suggests the interpretation that the measured $J = 1$ lifetime is limited by the collisional process.

Taken together, these measurements show that the observed $J = 1$ lifetime is not set by BBR contribution alone. Rather, ion–neutral collisions make a significant contribution, with H_2 being the most plausible dominant background gas in the current apparatus.

There is a major caveat in this interpretation; the determination of the baseline H_2 pressure, i.e. the first H_2 point in Fig. 4.8, is not completely unambiguous and not 100% conclusive. As discussed in Appendix A, no geometric correction was applied to this baseline measurement; instead, it was obtained from the raw RGA mass-2 reading using only the H_2 response-factor correction. This would equal the true background H_2 pressure in the apparatus, only if the RGA itself were not contributing additional hydrogen.

However, an active RGA ionizer can outgas or desorb hydrogen. Therefore, part of the measured mass-2 background may originate in the probe, rather than in the chamber alone. If this is the case, then the true baseline H_2 pressure in the vacuum system is smaller than the value plotted, and the first H_2 point in Fig. 4.8 would move to the left, towards lower pressure. Then, the apparent immediate rise of the decay rate with added H_2 would be less straightforward to interpret, and it would weaken the conclusion that the baseline $J = 1$ decay is hydrogen-limited.

For this reason, one should be careful when interpreting Fig. 4.8; this is still not a definitive

proof that the baseline $J = 1$ lifetime in the prototype apparatus is set by hydrogen alone. However, what is definite is that H_2 affects the decay rate even at very small partial pressures.

This caveat does not change the design implication for the Gen. III apparatus. Even if the prototype $J = 1$ lifetime is not conclusively shown here to be hydrogen-limited at baseline, both the measurements and the numerical estimates indicate that H_2 can produce a significant effect at very small partial pressure. Since the Gen. III apparatus will have larger surface area with additional electrodes and larger vacuum volume, controlling hydrogen outgassing still remains essential. This is one reason for moving away from stainless steel and towards titanium and aluminum. These materials are also advantageous from the standpoint of magnetic cleanliness, which is another important factor for precession spectroscopy.

4.2.3 Ion-ion collisions

As we have seen above, the hard-sphere ion-ion collisions can apply a Coulomb field. While this is comparable to the rotating electric bias field, it is not enough to mix the two rotational levels.

Experimentally too, we did not see any correlation between the $J = 1$ lifetime and ion number density and secular frequencies. Therefore, we conclude that at least at this point, the $J = 1$ lifetime is not limited by ion-ion collisions.

4.3 Conclusion

In this chapter, I identified the dominant mechanisms that remove signal from the quantum manifold and therefore limit the interrogation time in the prototype apparatus. For the $m_F = \pm 3/2$ Zeeman manifold, the combined dependence on the rotating-field magnitude, ion number, and secular frequency strongly supports the interpretation that ion-ion collisions are a major source of observed Ramsey decoherence. By contrast, within our sensitivity, collisions with neutral background gas do not appear to be an important contribution to the m_F -manifold lifetime under nominal operation conditions. The practical consequence is that to achieve long Ramsey interro-

gation time, the ion density has to be reduced during the free precession, for example by ramping down the trap confinement.

For the $J = 1$ population lifetime, the measurements indicate that BBR excitation is an important but not sufficient explanation of the observed loss. The measured lifetimes follow the expected trend from different effective temperature environment along the length of the prototype apparatus, but are systematically shorter than the BBR-only prediction, suggesting an additional temperature-independent loss channel. The measurements with various injected neutral gas show that ion-neutral collisions provide an explanation for this extra contribution. In particular, H_2 produces a clear increase in the $J = 1$ decay rate, even at very small inferred partial pressures, making it a highly plausible limiting background gas in the present apparatus.

At the same time, the interpretation of the baseline H_2 point requires caution. As discussed in Appendix A, the first H_2 point is obtained from the baseline measurement of RGA at mass-2. Therefore, this point may overestimate the true baseline pressure of H_2 , if the active RGA source itself is contributing hydrogen to that measurement. While we cannot make a definitive conclusion that the prototype is definitely hydrogen-limited at baseline, the measurements still show that very small H_2 partial pressures already have a significant measurable effect on the $J = 1$ lifetime. This conclusion motivated us to move towards using titanium and aluminum for the Gen. III apparatus, for reduced hydrogen outgassing and nonmagnetic properties.

Chapter 5

Towards Long Coherence Time: Loss of Signal from Trap

In the previous chapter, we have seen mechanisms that lead to the loss of signal by reducing the fraction of the ions that remain in the target measurement states. Another very important mechanism through which we can have loss of signal is by physically losing ions from the trap.

In this chapter, I discuss trap lifetime of the ions in the prototype apparatus and several mechanisms that affect the lifetime. These include ion–neutral collisions, heating in the rf ion trap, and the additional motional energy associated with the rotating electric field $\mathcal{E}_{\text{rot}}$.

5.1 Trap lifetime of ions

One of the main advantages of a molecular ion platform is that ions can, once produced and loaded, be confined for long times in ion traps. However, even in an ion-trap based experiment, long trapping times are not automatically guaranteed under the conditions required for precision measurement. In the JILA EDM experiment, trap lifetime is determined not only by the presence of trapping fields, but also by the detailed shape of the effective trapping potential and by mechanisms that increase the ions' motional energy.

The present trap for the Gen. III prototype experiment was designed primarily to generate a highly uniform rotating electric field, rather than to realize an ideal harmonic ion trap. Therefore, the current trap design consists of the eight-electrode geometry shown in Chapter 2. When quadrupole rf voltages are applied in this geometry, the resulting rf pseudopotential contains significant anharmonic terms. In particular, along some azimuthal directions, the pseudopotential

includes negative quartic contributions, which lower the effective saddle points through which ions can escape.

For the rest of the thesis, we distinguish between the *tightness* of ion confinement proportional to the secular frequency at the trap center and the *effective depth* of the trap set by the total energy required to escape the trap via the lowest saddle point.

This feature becomes especially important during Ramsey interrogation. To reduce ion density and therefore suppress ion–ion interactions, the trap confinement is typically loosened by lowering the secular frequencies. In an ideal harmonic trap, this would simply broaden the cloud while leaving the trap depth essentially unchanged. In our prototype trap, however, looser confinement also substantially lowers the effective trap depth because of the anharmonic terms in the pseudopotential, as shown in Figure 5.1.

As a result, ions become more vulnerable to any mechanism that increases their motional energy. Accordingly, trap loss in the prototype is best understood not as a sum of completely independent decay channels, but as the consequence of several effects that either increase the ions’ motional energy or reduce the effective trap depth. Mechanisms such as rf heating, electrode-noise, and ion–neutral collision increase the secular energy of the ions. Furthermore, nonuniformities in the rotating electric field can add position-dependent excess micromotion energy and distort the effective potential. The trap loss occurs once these energy scales become comparable to the escape barrier of the trap.

5.2 Estimating ion temperature

To study the origin of trap loss, we require some estimate of the ions’ motional energy. Broadly speaking, there are two ways of probing the temperature of the ion cloud. The first is to measure the width of the ion cloud on the detector and use a transfer function to relate the observed detector distribution back to the distribution in the trap. In this approach, the temperature is *inferred* from the measured cloud widths. The second method is to use Doppler spectroscopy, in which the observed transition linewidth provides a more direct estimate of the




Figure 5.1: Schematic illustration of the relation between confinement and effective trap depth. **Top:** in an ideal harmonic trap, reducing the confinement broadens the ion cloud while the potential remains a simple single well. **Bottom:** in a realistic anharmonic trap, negative quartic contributions lower the escape saddles, so weakening the confinement can reduce the effective trap depth much more dramatically. This makes ions more susceptible to loss once heating or other motional-energy sources become comparable to the available trap depth.

ion-cloud temperature. In this section, I discuss these two complementary methods for estimating ion cloud temperatures.

5.2.1 Transfer-function thermometry

To compare the motional energy of the ion ensemble under different trapping conditions, we infer an effective ion temperature from the measured cloud widths on the detector. This diagnostic should be understood primarily as a *relative* thermometer. As discussed below, uncertainties in the detector ion optics prevent us from assigning reliable absolute temperatures or absolute heating rates. Nonetheless, this method is useful for tracking relative changes in the ensemble energy scale as different experimental parameters are varied.

The transfer function relates the phase-space coordinates of an ion at the trap to those at the detector. For each Cartesian direction, we can write

$$\begin{pmatrix} i \\ v_i \end{pmatrix}_{\text{det}} = \begin{pmatrix} m_{11,i} & m_{12,i} \\ m_{21,i} & m_{22,i} \end{pmatrix} \begin{pmatrix} i \\ v_i \end{pmatrix}_{\text{trap}}, \quad i \in \{x, y\}. \quad (5.1)$$

For the radial directions x and y , we measure the position of the ion cloud at the detector. For the axial direction, which is along the time-of-flight direction, we measure

$$\begin{pmatrix} t_{\text{tof}} \\ v_i \end{pmatrix}_{\text{det}} = \begin{pmatrix} m_{11,z} & m_{12,z} \\ m_{21,z} & m_{22,z} \end{pmatrix} \begin{pmatrix} z \\ v_z \end{pmatrix}_{\text{trap}}, \quad (5.2)$$

the time-of-flight at the MCP. In practice, only the position for x, y , and time for z at the detector are measured directly, so the analysis depends primarily on the coefficients m_{11} and m_{12} .

To determine these coefficients, we apply a short impulse at $t = 0$, chosen so that

$$x_{\text{trap}}(t = 0) = 0, \quad v_{\text{trap}}(t = 0) = v_0 = \frac{eE_{\text{impulse}}\Delta t}{m},$$

where e is the elementary charge, E_{impulse} is the applied impulse field, Δt is the duration, and the m is the ion mass. For harmonic secular motion at frequency ω_{sec} ,

$$x_{\text{trap}}(t) = \frac{v_0}{\omega_{\text{sec}}} \sin(\omega_{\text{sec}} t), \quad v_{\text{trap}}(t) = v_0 \cos(\omega_{\text{sec}} t).$$

The detected position is therefore

$$\begin{aligned} x_{\text{det}}(t) &= m_{11}x_{\text{trap}}(t) + m_{12}v_{\text{trap}}(t) \\ &= v_0 \left[\frac{m_{11}}{\omega_{\text{sec}}} \sin(\omega_{\text{sec}}t) + m_{12} \cos(\omega_{\text{sec}}t) \right]. \end{aligned} \quad (5.3)$$

We fit the signal to

$$x_{\text{det}}(t) = A \sin(\omega_{\text{sec}}t + \phi),$$

which gives

$$\begin{aligned} m_{11} &= \frac{\omega_{\text{sec}}A}{v_0} \cos(\phi), \\ m_{12} &= \frac{A}{v_0} \sin(\phi). \end{aligned} \quad (5.4)$$

In practice, we vary v_0 by varying the impulse field and extract the slope of A versus v_0 , from which m_{11} and m_{12} are obtained. Figure 5.2 shows an example of the transfer function measured in this way.

Under the additional assumption that the in-trap distribution is approximately thermal in a harmonic well, with negligible x - v_x covariance, the detector width and trap width are related by

$$\begin{aligned} \sigma_{x,\text{det}}^2 &= m_{11}^2 \sigma_{x,\text{trap}}^2 + m_{12}^2 \sigma_{v_x,\text{trap}}^2, \\ \sigma_{x,\text{trap}} &= \frac{\sigma_{x,\text{det}}}{\sqrt{m_{11}^2 + (m_{12}\omega_{\text{sec}})^2}}, \end{aligned} \quad (5.5)$$

where $\sigma_{v_x,\text{trap}} = \omega_{\text{sec}}\sigma_{x,\text{trap}}$ was used in the second line. An analogous relation holds for the other axes. In a harmonic potential, the corresponding effective temperature may then be written as

$$k_B T_x = m\omega_x^2 \sigma_{x,\text{trap}}^2, \quad (5.6)$$

with similar expressions for T_y and T_z .

We refer to this implementation as the “**slosh-kick**” method. In the prototype apparatus, however, this method suffers from an important limitation: the ion cloud cannot be measured immediately after the impulse. The earliest time we can kick the ions towards the detector to be measured is roughly 2 ms after the impulse, corresponding under nominal conditions to several secular oscillations. As a result, the fitted phase ϕ is not well constrained, which introduces a

significant uncertainty. This limitation was one of the first indications that the transfer function could not be trusted at the level required for precise absolute thermometry.

To cross-check the slosh-kick measurement, we also implemented a complementary “**shim-and-kick**” method. In this approach, rather than inferring position and velocity sensitivity from the phase of the secular motion, we deliberately offset the ion position or impart a controlled additional velocity immediately before the kickout. Position and velocity dependencies are therefore measured separately. Although this method is operationally more involved, it has the advantage of comparing directly to the ion phase-space coordinates at the moment of kickout and is less sensitive to the uncertainty in the initial secular phase. Figure 5.3 shows the transfer functions measured using the shim-and-kick method.

The slosh-kick and shim-and-kick methods agree at a qualitative level, but both revealed a more serious systematic associated with the detector ion optics. Importantly, this systematic affects centroid measurements and width measurements in different ways. The transfer-function measurements probe the mapping between the ion cloud’s centroid coordinates at the trap and the centroid position observed at the detector. In other words, the transfer functions measure the first-order transport properties of the ion optics between the trap and the detector. By contrast, inferring temperature from the measured cloud width requires that the second-order moment to be transported in a controlled and well-understood way. This second requirement is considerably more difficult to satisfy, and in the current apparatus is violated by space-charge effects and enhanced focusing due to these space-charge effects during the transport. As a result, the width-based thermometry is distorted by downstream focusing and space-charge effects, even when the transfer function remain useful for the centroid measurements.

A clear indication of this problem was that the measured cloud width depended strongly on the number of ions loaded into the trap. At first sight, this trend could have been interpreted as a genuine increase in the in-trap cloud size or temperature with ion number. Indeed, the observed trend was roughly consistent with an $N^{1/3}$ scaling expected for a cold, space-charge-dominated one-component plasma in an approximately harmonic potential. Figures 5.4 and 5.5 show the same




Figure 5.2: Example of the transfer function measured with the slosh-kick method. By applying a short impulse and varying the delay before extraction, the centroid motion of the ion cloud is mapped onto the detector position. Fitting the oscillation determines the transfer coefficients relevant for centroid transport. The signs for m_{11} and m_{12} are not carried out rigorously throughout the analysis, therefore here they should be considered as the absolute values.




Figure 5.3: Transfer function measured with the shim-and-kick method. Here the centroid position and centroid velocity are varied separately immediately before extraction, allowing the position and velocity sensitivities of the detector optics to be measured more directly. The signs for m_{11} and m_{12} are not carried out rigorously throughout the analysis, therefore here they should be considered as the absolute values.

dataset in two complementary ways. In Fig. 5.4, the points are colored by the experiment, so that each color corresponds to a single experimental run in which we attempted to keep the ion number constant. The data within each run span times from early after loading (~ 100 ms) to late times (12 s). This representation makes clear that the width-ion number correlation is already present within an individual dataset. It also shows the fluctuations in ion number within one experiment. In Fig. 5.5, the same points are colored by time, with darker points corresponding to earlier times and lighter points to later times. However, in the present case, we believe that most of the observed number dependence was introduced downstream of the trap, in the ion optics, rather than reflecting a genuine temperature change in the in-trap temperature.

The key issue is that, in the detector geometry, the MCP was recessed behind the vacuum flange and surrounded by additional shielding structures, as shown in Fig. 5.6.

The electric field used to bias the MCP therefore produced a strong focus before the ions reached the detector plane. When the detector is located beyond this focus, space-charge effects during the post-focus expansion can produce a measured width that depends strongly on ion number, even if the in-trap distribution is unchanged. This mechanism provides a more plausible explanation of the observed number dependence than a large, repeatable change in the true ion temperature. Independent laser-based thermometry, discussed in the next section, also supports the conclusion that the true ion temperature does not change strongly with ion number under these conditions.

This strong focusing field can be reduced by using SYRiA, a set of accelerating electrodes near the MCP. I designed SYRiA so that we can raise the ion kinetic energy (for higher quantum efficiency) and to reduce uncontrolled focusing onto the MCP. In practice, however, the recessed detector geometry, the imperfections in the surrounding shielding structure, and likely patch-charge effects made the ion optics difficult to control. The distortions of the nearby mesh and shielding illustrated in Fig. 5.7 further limited the performance of SYRiA.

All of these issues eventually motivated us to abandon SYRiA for routine operation, partic-




Figure 5.4: Measured detector width as a function of ion number. Although the width increases with ion number, this dependence is attributed primarily to downstream focusing and post-focus space-charge expansion in the detector optics, rather than to a large change in the in-trap temperature. Each color corresponds to a single experiments. During a single experiment, we tried to keep the ion number constant, but the time span from ~ 100 ms to ~ 12 s.




Figure 5.5: Measured detector width as a function of ion number. Although the width increases with ion number, this dependence is attributed primarily to downstream focusing and post-focus space-charge expansion in the detector optics, rather than to a large change in the in-trap temperature. Here, the points are colored by time; darkest points are at early times ~ 100 ms, and the lightest points are at late times ~ 12 s.




Figure 5.6: Picture of the old MCP. The old MCP was in a 4.5-inch flange and was recessed into the flange. The surface of the front MCP was about 1.147 inches below the surface of the flange. The MCP was about 44 mm in diameter with the mesh in the front. The mesh and the MCP assembly is surrounded by the cylindrical shield, that is electrically connected to the mesh.




Figure 5.7: Example of the detector-optics asymmetry observed near the recessed MCP. Imperfections in the local electrode and shielding geometry likely contributed to uncontrolled focusing.

ularly because adequate ion width could still be obtained without it.¹

Because of these detector-optics systematics, the transfer function extracted with the original detection setup cannot be interpreted as a reliable absolute calibration of the ion temperature.² However, this method remains a useful relative diagnostic. Throughout this thesis, temperature trends inferred from the transfer function are therefore interpreted comparatively: they indicate whether a given change in the apparatus tends to increase or decrease the characteristic motional energy of the ion ensemble, provided that the ion number is kept approximately constant.

In Spring 2026, the detection system was upgraded to a new design, TUNISIA, developed by a graduate student Tuan Nguyen. This upgrade replaced the earlier disk-based accelerating structure with segmented cylindrical electrodes and incorporated a new MCP that is larger and not recessed in the vacuum flange, as shown in Fig. 5.8. The goal of this redesign was to reduce sensitivity to stray fields, patch charges, and detector-geometry-induced focusing. Most of the data discussed in this thesis, however, were acquired before that upgrade and therefore remain subject to the limitations of the earlier transfer-function calibration.

5.2.2 Laser-based thermometry

Because the temperature estimate based on the cloud widths raised doubts about whether the apparent hotter ion cloud with increasing ion number really reflected a true temperature change, we also estimated the ion cloud temperature using Doppler spectroscopy. For this measurement, I used Hen, shown in Fig. 2.13, which depletes the $J = 1$ state by driving it to the $\Omega = 0^+$ electronic state, from which population subsequently decays into the $^1\Sigma$ manifold.³

The Hen power was reduced so that the transition was not saturated, and the laser frequency

¹ At the time, I compared the measured cloud width primarily to single-particle SIMION simulations. That comparison was misleading because SIMION does not include ion-ion interactions, and therefore it could not reproduce the additional broadening associated with post-focus space-charge expansion during kickout.

² In theory, we can model the entire ion “transfer function” using a combination of SIMION and LAMMPS (which is a numerical simulator that takes the ion-ion interaction into account). I gave this a quick try, but it was hard to model the whole trajectory with only one effective single thin lens. This is definitely worth a more thorough attempt in the future with better understood ion optics in the system.

³ Normally, Hen is strobed and used as the m_F -depletion laser. In the dataset discussed here, however, the rotating electric field $\mathcal{E}_{\text{rot}}$ was off, therefore the effect of strobing is mainly a reduced duty cycle for this measurement.




Figure 5.8: Photograph of a new MCP picture. The new MCP is bigger than the old MCP and more importantly not recessed into the vacuum flange. For most of the measurements shown in this thesis however, the new MCP was not yet installed.

was scanned across the resonance. We dissociate using $R(1)$ transition, to see how many ions are left in the $J = 1$ state after the ion cloud was illuminated with Hen. After the cloud was illuminated with Hen, the remaining population in the $J = 1$ state was measured by dissociating on the $R(1)$ transition. The data were normalized in two steps; first by the off-resonant $J = 1$ ion number, and second on shot-to-shot basis using the TOF area to account for fluctuations in ion production. The resulting spectrum was fit to a Gaussian profile

$$S(f) = c + A \exp \left[-\frac{(f - f_c)^2}{2\sigma_f^2} \right],$$

where $A < 0$ describes the depletion amplitude, f_c is the fitted line center, σ_f is the Gaussian width, and c is an offset.

Assuming the observed width is dominated by Doppler broadening, the fitted width σ_f can be converted to a temperature through

$$\Delta f_{\text{FWHM}} = 2\sqrt{2 \ln 2} \sigma_f, \quad T = \frac{mc^2}{8k_B \ln 2} \left(\frac{\Delta f_{\text{FWHM}}}{f_0} \right)^2 = \frac{mc^2}{k_B} \left(\frac{\sigma_f}{f_0} \right)^2,$$

where m is the ion mass and f_0 is the transition frequency.

Figure 5.9 shows the resulting spectra. Runs 4 and 5 were taken at low ion number, approximately 1000 ions, whereas runs 6 and 7 were taken at high ion number, with approximately 8000 ions. The right panel shows the number of trapped ions differs significantly between these two conditions, but the depletion widths in the left panel remain comparable.

The fitted Gaussian widths for runs 4-7 are

$$\sigma_f = 45.9, 40.1, 36.0, 47.8 \text{ MHz},$$

which correspond to temperatures

$$T^{\text{raw}} = 55.9(12.5), 42.8(12.6), 34.3(8.5), 60.7(19.2) \text{ K},$$

where the quoted uncertainties come from propagating the fit uncertainty in σ_f . The fitted offsets remained close to one, as expected from normalization. The fitted line centers varied only at the few-MHz level, with no obvious dependence on ion number.




Figure 5.9: Hen Doppler thermometry for four loading conditions. **Left:** doubly normalized depletion signal as a function of Hen frequency, together with Gaussian fits. Runs 4 and 5 were taken at low ion number ($\sim 10^3$ ions), while runs 6 and 7 were taken at high ion number ($\sim 8 \times 10^3$ ions). **Right:** corresponding trapped-ion number versus Hen frequency. Despite the order-of-magnitude change in ion number, the fitted linewidths remain similar.

Most importantly, these measured temperatures are consistent within uncertainty and show no clear systematic increase with increasing ion number. This behavior contrasts with the detector-based width analysis shown in Fig. 5.4, where the apparent cloud width and the inferred temperature increase substantially at higher ion number. Taken together, these results suggest that the detector-width measurement is sensitive to effects beyond intrinsic thermal velocity distribution, and that the true ion temperature is not changing significantly as the number of ions in the trap is varied.

However, when this measurement is interpreted as an absolute temperature measurement, one must be more cautious. That is because the ion beam enters the chamber radially, and the Doppler width is sensitive not only to the radial secular motion but also to the rf micromotion. In the small- q limit, we can write the ion trajectory u_i in the radial direction as [71]

$$u_i(t) \approx u_{1i} \cos(\omega_i t + \phi_i) \left[1 + \frac{q_i}{2} \cos(\Omega_{\text{rf}} t) \right],$$

where u_{1i} is the oscillation amplitude, ϕ_i is a phase determined by the initial conditions, ω_i is the secular frequency, and $i \in \{x, y\}$ labels the two radial directions. From this, we see that the time-averaged kinetic energy contains both secular and micromotion contributions. Using the corresponding expression for the time-averaged kinetic energy from Ref. [71], the apparent Doppler temperature inferred from a radial probe is

$$\frac{2E_{k,i}}{k_B} = T_{\text{app}} = T_{\text{sec}} \left(1 + \frac{q_i^2 \Omega_{\text{rf}}^2}{8\omega_i^2} \right).$$

For the trap parameters used for this measurement, this factor is of order two. Using $f_{\text{rf}} = 50$ kHz, $f_{\text{sec}} \approx 2$ kHz, and $q \sim 0.1$, one finds

$$1 + \frac{q^2 \Omega_{\text{rf}}^2}{8\omega_{\text{sec}}^2} = 1 + \frac{q^2}{8} \left(\frac{\omega_{\text{rf}}}{\omega_{\text{sec}}} \right)^2 \sim 2.$$

In addition, I carried out a simple single-ion numerical simulation in which the ion was probed at random phases of the radial motion. The extracted linewidth again gave a correction factor of order two, consistent with the analytic estimate. Therefore, the temperatures extracted from the

radial Doppler spectra should be interpreted as apparent radial kinetic temperatures. The corresponding secular temperatures are smaller by approximately a factor of two. The corresponding secular temperatures of runs 4 through 7 therefore are

$$T_{\text{sec}} = 28.0(6.3), 21.4(6.3), 17.2(4.2), 30.4(9.6) \text{ K.}$$

This caveat about the apparent vs absolute temperature of the Doppler width affects the absolute temperature scale, but it does not change the main conclusion of this section: the temperature does not show a clear dependence on ion number.

Overall, this method is less sensitive to the distortions from the ion-optics downstream from the trap discussed in the previous section and therefore provides an important cross-check on the width-based diagnostic. Its interpretation is still subject to assumptions about the line shape and the extent to which Doppler broadening dominates the observed width, but it offers an independent order-of-magnitude estimate of the secular temperature scale.

5.3 Observed heating in the prototype

Taken together, the detector-width measurements and the laser-based thermometry show that the ion ensemble acquires additional motional energy over time in the prototype trap. For the reasons discussed above, it is hard to extract a reliable absolute heating rate from the cloud widths measured at the detector. However, the relative evolution of the cloud width with time is still informative, and it provides a practical way to compare how the heating trend changes as experimental parameters are varied.

For the width-based measurement, the measured detector width depends strongly on ion number, we attempted to keep the number of ions approximately fixed when studying heating under different conditions.

Figure 5.10 shows an example of a heating measurement. Again, the y -axis should not be interpreted too literally as an absolute temperature scale, but rather as a relative measure of the square of the cloud width. The main observation is that the inferred temperature increases




Figure 5.10: Example width-based heating measurement. The inferred temperatures in the x , y , and z directions are shown as a function of trapping time, together with the corresponding ion number. The absolute temperature scale on the y -axis should be interpreted cautiously and not be taken literally, since these values are inferred from detector widths and are therefore sensitive to downstream ion-optics effects. Nonetheless, the relative increase with trapping time is real: the inferred temperature rises significantly over the first few seconds and then appears to flatten.

substantially over the first several seconds in all three directions, corresponding to a clear broadening of the ion cloud at the detector. This increase does not appear to be linear in time: most of the change occurs during the first ~ 5 s, and the trend becomes noticeably flatter, although the uncertainties are still large.

To obtain an order-of-magnitude estimate of the absolute heating rate, we also performed Doppler spectroscopy on the ion cloud at an early time, around 300 ms (blue), and at a later time around 5 s (orange).




Figure 5.11: Hen Doppler thermometry at early, about 300 ms, (blue) and late, about 5 s, (orange) trapping times. **Left:** doubly normalized depletion signal as a function of Hen frequency, together with the Gaussian-fits. **Right:** corresponding trapped-ion number during the two measurements. The fitted widths imply a larger secular temperature at late time than at early time.

The corresponding secular temperatures are:

$$T_{\text{sec}} = 13(3), 33(8) \text{ K},$$

which imply an estimated heating rate of approximately

$$4(2) \text{ K/s}.$$

Within uncertainty, this estimate is consistent with the width-based observation that the apparent temperature increases significantly over the first several seconds. The uncertainty in both measurements are still fairly large, so this comparison should not be overinterpreted. Nonetheless, the two methods give a self-consistent picture: the heating is on the scale of a few Kelvins per second.

5.4 Mechanisms contributing to trap loss

Having established that the ion ensemble gains motional energy on experimentally relevant timescales, I now discuss several mechanisms that we investigated that can contribute to trap loss in the prototype. Rather than attempting a complete quantitative decomposition of the loss rate, I focus on the dominant experimental knobs that were observed to affect the trap lifetime.

5.4.1 Energy associated with $\mathcal{E}_{\text{rot}}$

The rotating electric field $\mathcal{E}_{\text{rot}}$ is required to polarize the molecules and define a quantization axis, but it also drives substantial rotational micromotion of the trapped ions. As discussed earlier, for an ideal field

$$\mathcal{E}_{\text{rot}}(t) = \mathcal{E}_0 [\cos(\omega_{\text{rot}}t)\hat{\mathbf{x}} + \sin(\omega_{\text{rot}}t)\hat{\mathbf{y}}],$$

an ion of charge e and mass m undergoes circular motion with radius

$$r_{\text{rot}} = \frac{e\mathcal{E}_{\text{rot}}}{m\omega_{\text{rot}}^2}, \quad v_{\text{rot}} = \frac{e\mathcal{E}_{\text{rot}}}{m\omega_{\text{rot}}}.$$

The associated rotational micromotion energy is therefore

$$U_{\text{pond,rot}} = \frac{e^2}{2m\omega_{\text{rot}}^2} |\mathcal{E}_{\text{rot}}|^2, \quad (5.7)$$

This energy scale is very large. For the typical experimental parameters for the prototype, this energy is of order 97000 K for $\mathcal{E}_{\text{rot}} = 60$ V/cm, $\omega_{\text{rot}} = 2\pi \cdot 147$ kHz, and $m = 251$ amu.

In a perfectly uniform rotating field, this large micromotion energy is a common effect: all ions execute the same circular motion, so the field does not reduce the trap lifetime by itself. However, the problem arises when imperfections in the rod drive distort the rotating field so that its amplitude and/or phase acquires spatial structure across the ion cloud. In that case, the rotational micromotion energy becomes position dependent and acts like an additional effective potential.

In the original rod-driver system, we implemented four resonant tank circuits, each driving one pair of opposite rods (for example rods 1 and 5 in Fig. 2.10). The differential component, $V_{\text{diff},15} = V_{\text{rod1}} - V_{\text{rod5}}$, produces the intended bias field between the two rods. The common mode

component, $V_{\text{comm},15} = V_{\text{rod}1} + V_{\text{rod}5}$, produces the desired rf voltage. In practice, we had four independent controls over the differential mode. But for the common-mode, we only had two independent controls across the four rod pairs. More detailed descriptions of the original rod-driver design and implementation can be found in the theses of Kia Boon Ng and Noah Schlossberger [38, 59]. In the original design, the differential component of each opposite rod was made very uniform across the four rod pairs, because the design emphasis was on the homogeneity of the rotating electric field for precision measurements. By contrast, the common-mode component were originally intended mainly to supply the rf voltage and therefore were not engineered for the same level of pair-to-pair control. As a result, the actual voltages appearing on the rods were not fully uniform in amplitude and phase across all eight rods.

This limitation became severe once we attempted to operate at large $\mathcal{E}_{\text{rot}}$. In the original configuration, we could not reliably retain ions in the trap for more than ~ 3 s when $\mathcal{E}_{\text{rot}}$ was ramped to 60 V/cm. It eventually became clear to us that matching only the differential drive of the opposite rod pairs was not sufficient. The differential signal on a given opposite rod pair mainly determines the intended rotating field component along that pair's axis. On the other hand, the common-mode signal shifts the average voltage of that pair. If the common mode were identical for all rod pairs, it would result in a global offset and would be harmless. However, because pair-to-pair common-mode matching had not been a design requirement of the original system, the common-mode signals were not uniform. As a result, even when the differential drive was reasonably well matched across the four opposite pairs, the pair-dependent common-mode deviations distorted the field and introduced spatially varying curvature and angular structure in the potential. In other words, the rotating field was no longer truly round around the azimuthal direction.

More quantitatively, the field experience by the ions can be written as

$$\mathcal{E}_{\text{rot}}(\mathbf{r}, t) = \mathcal{E}_0(t) + \delta\mathcal{E}(\mathbf{r}, t),$$

where $\mathcal{E}_0(t)$ is the desired spatially uniform rotating field and $\delta\mathcal{E}(\mathbf{r}, t)$ contains the unintended spatial gradients and distortions arising from rod-to-rod imbalance. Because the rotational micro-

motion energy scales as the square of the electric field, even a small distortion of $\delta\mathcal{E}$ can produce a substantial position-dependent energy shift through the cross term with the large baseline field

$$|\mathcal{E}_{\text{rot}}|^2 = |\mathcal{E}_0|^2 + 2\mathcal{E}_{\text{rot}} \cdot \delta\mathcal{E}(\mathbf{r}) + |\delta\mathcal{E}(\mathbf{r})|^2. \quad (5.8)$$

Therefore, ions at different radii and azimuthal angles experience different values of $|\mathcal{E}_{\text{rot}}|$, which increases the spatial variation of the rotational micromotion energy and can strongly shorten the trap lifetime. In the original rod-driver system, this distortion was large enough that we could not stably hold ions for a long enough times at $\mathcal{E}_{\text{rot}} = 60 \text{ V/cm}$.⁴

Looking at Eq. (5.7), the most obvious improvement is to increase the rotation frequency f_{rot} . In Gen. II, the field rotated at $f_{\text{rot}} = 375 \text{ kHz}$. However, during the design of the Gen. III prototype, we were concerned about the amount of current required to drive the much longer rod electrodes, which have substantially larger capacitance. Therefore, we decided to operate at $f_{\text{rot}} \approx 150 \text{ kHz}$ initially. In Fall 2025, we did develop a new rod-driver system capable of operating at significantly higher frequency, around 300 kHz. That new system was designed by fellow graduate student Anzhou Wang, and more details will be given in his thesis.

Before that new system was available, however, I identified a practical short-term solution.⁵

The key insight was that the “rf channels” could be used to shim away the nonuniform common mode at f_{rot} . I therefore upgraded the original rod-driver system to include two additional common-mode channels. With these added channels, we obtained independent control of both the differential and common-mode signals for each opposite rod pair, so that the amplitude and phase of all eight rod voltages could be tuned individually rather than only matching the pairwise difference. We also introduced lock-in based measurements of the differential and common-mode amplitudes and phases⁶, which made it possible to tune the rod voltages more precisely.

⁴ Technically, we could keep ions with 60 V/cm of $\mathcal{E}_{\text{rot}}$ if we increased the tightness of confinement, since the effective trap depth also gets deeper, per Fig. 5.1. However, for the Ramsey spectroscopy, the confinement is loosened, and the combination of the shallower effective trap depth and the distortion of the rotating field resulted in a very short trap lifetime of our ions.

⁵ Supposed to be “short”, but the solution worked well enough for the experiments we were running, and the development of the new rod-drive system took long enough that we ultimately used this solution for about 2 years.

⁶ Previously, these quantities had been adjusted qualitatively from oscilloscope traces.

These upgrades allowed us to be able to retain ions for long enough to operate the experiment at $\mathcal{E}_{\text{rot}} = 60$ V/cm, and in practice they enabled the measurements presented in this thesis. Although a completely new rod-driver system was eventually constructed, most of the work described in this thesis was carried out with the upgraded original rod-driver system. In that sense, identifying the importance of the common-mode imbalance and implementing the additional common-mode control was a necessary experimental step that made the subsequent prototype measurements possible.

Looking back at Eq. (5.8), we can estimate what kind of deviation is required for the cross term to become comparable to a realistic trap depth U_{trap} :

$$\begin{aligned} U_{\text{pond,rot}} &\approx \frac{e^2}{2m\omega_{\text{rot}}^2} (2\mathcal{E}_{\text{rot}} \cdot \delta\mathcal{E}) \\ |\delta\mathcal{E}_{\text{rot}}| &< \frac{m\omega_{\text{rot}}^2}{e^2|\mathcal{E}_{\text{rot}}|} U_{\text{trap}} = 0.3 \text{ V/cm}, \end{aligned} \tag{5.9}$$

where I have used U_{trap} of 800 K^7 , $\mathcal{E}_{\text{rot}}$ of 60 V/cm , and $\omega_{\text{rot}} = 2\pi \cdot 147.5 \text{ kHz}$. This means the rotating field has to be uniform down to about 0.5% over the extent of the ion cloud, not for the sake of the precision measurement itself, but simply in order to retain the ions in the trap.

To visualize the effect of the spatially varying ponderomotive potential from an “un-round” rotating electric field, we deliberately introduced a controlled nonuniformity in the applied rod voltages. Figure 5.12 shows a COMSOL calculation of the resulting $|\mathcal{E}_{\text{rot}}|^2$ inhomogeneity over the approximate cloud extend of $\pm 1 \text{ cm}$. The first row shows the nominally symmetric case. In the second row, the voltage on the top electrode is increased by 1% , and in the third row it is decreased by 1% .

As expected, between the second and third rows, the sign of the electrode imbalance reverses the sign of gradient in the ponderomotive potential across the cloud. However, the two signs are not equivalent. In the $+1 \%$ case, the distortion of $|\mathcal{E}_{\text{rot}}|^2$ is spread more broadly over the cloud. In the -1% case, however, the perturbation produces a more localized low- $|\mathcal{E}_{\text{rot}}|^2$ region near the weakened rod, so the corresponding ponderomotive “hole” is sharper and deeper. In other words,

⁷ This is the trap depth at our nominal secular frequency of 2 kHz . At lower secular frequency, which is needed for maintaining coherence for a long time, the corresponding trap depth is lower, as low as $\sim 300 \text{ K}$.




Figure 5.12: COMSOL study of how a deliberate single-electrode amplitude imbalance produces spatial variation in the ponderomotive energy associated with the rotating electric field. The left column shows the calculated spatial pattern of $|\mathcal{E}_{\text{rot}}|^2$ over the trap cross section, within the approximate ion-cloud extent of ± 1 cm. The right column shows the corresponding line cut through the cloud region, in the direction of the imbalance. **Top:** nominally symmetric voltages applied to all eight rods. **Middle:** the voltage on the top rod increased by 1 %. **Bottom:** the voltage on the top rod decreased by 1 %. Although the same fractional voltage error is applied in both directions, the resulting field distortion is not symmetric: the positive perturbation produces a broader, shallower distortion, while the negative perturbation creates a sharper and deeper local low- $|\mathcal{E}_{\text{rot}}|^2$ region on the side of the perturbed rod.

in our 8-rod configuration, even though the two perturbations have the same magnitude at the level of the applied rod voltage, the negative perturbation produces a substantially larger local variation in $|\mathcal{E}_{\text{rot}}|^2$ across the ion cloud.

The corresponding experimental result is shown in Fig. 5.13. The figure shows the ion survival as a function of the imposed fractional voltage amplitude offset on two of the electrodes, individually. The negative offsets rapidly deplete the trapped ions, whereas positive offsets are tolerated better, as the COMSOL simulation suggests. Although only the offset on Rod 1 and Rod 5 are varied in this figure, the same qualitative asymmetry was observed when scanning each of the eight rods individually. This suggests that the asymmetric sensitivity to positive and negative offsets is not primarily a rod-specific patch-charge effect, but rather a generic feature of how single-rod amplitude errors distort the rotating field in this eight-rod geometry.

The experimentally relevant rod-*voltage* offsets are about -2.5% and $+10\%$. The COMSOL calculation shows that a 1% rod-voltage offset produces about 0.1% field variation on the negative side and 0.04% variation on the positive side. With this conversion, the observed voltage fractional offset thresholds are broadly consistent with the simple estimate mentioned earlier of $\delta\mathcal{E}_{\text{rot}}/\mathcal{E}_{\text{rot}} \sim 0.5\%$, where the ponderomotive distortion becomes comparable to the trap depth.

In reality, even if we applied a perfect set of voltages to the electrodes, field inhomogeneity could still arise from geometric imperfections. In the prototype, we do not simply operate with the nominal ideal voltages; we use the ions themselves to shim the field imperfections.⁸ We also found out experimentally that different trapping regions require different sets of voltages.

In the prototype experiment, depending on where we perform the experiment, we use different sets of voltages so that the ions experience a field that is as round as possible. This is not something we can do in the full Gen. III apparatus, so it will be important to pay much closer attention to the physical tolerances of the electrode locations.

⁸ To shim the voltages, we first lower the trap depth so that ion survival becomes more sensitive to the ponderomotive distortion. Then we experimentally vary the amplitude and phase of each electrode while monitoring whether the ions survived. In this way, the ions serve as a probe of how round the rotating field actually is.




Figure 5.13: Measured ion survival while deliberately varying the fractional voltage amplitude of one rod relative to its nominal value, with the other rod voltages held fixed. The scans are shown for Rod 1 and Rod 5. The left axis shows the area under the TOF curve, and the right axis shows the corresponding ion number. For both rods, negative amplitude offsets rapidly deplete the trapped ions, whereas positive offsets are tolerated substantially better. Similar qualitative behavior was observed when scanning each of other seven rods individually, consistent with the COMSOL result that a negative single-rod perturbation produces a sharper local ponderomotive distortion than a positive perturbation of the same magnitude.

5.4.2 rf heating

Heating is relevant here primarily because it can shorten the trap lifetime, although in a precision-measurement context, increased motional energy is also relevant more broadly. However, in this section, I will focus on heating mainly as a possible contributor to trap loss.

I will first discuss basic theory behind rf heating and some numerical simulations we did. Then, I will move to the experimental results.

5.4.2.1 rf heating theory

In an ideal collisionless rf trap, the ions follow stable Mathieu trajectories. In this case, the time-dependent rf field does not by itself imply heating. However, when ions go through collisions and interrupt those trajectories, the rf field can do net work on the ions. In other words, the heating is not simply a consequence of time-dependent confinement, but of the combination of the rf drive with *collisions* that transfer energy from driven micromotion into secular motion.

A useful way to think about the physics behind rf heating is

$$\dot{T} \sim \Gamma_{\text{coll}} \times (\text{energy gain per collision}), \quad (5.10)$$

where Γ_{coll} is the ion-ion collision rate. A simple estimate based on $\Gamma = n\sigma v$, up to numerical factors and the Coulomb logarithm, gives

$$\Gamma_{\text{coll}} \propto n \frac{e^4}{\sqrt{m} T^{3/2}}, \quad (5.11)$$

where n is the ion number density, e is the elementary charge, and m is the ion mass. The Coulomb logarithm arises from integrating over impact parameters,

$$\int_{b_-}^{b_+} \frac{1}{b} db = \ln(b_+/b_-),$$

where b_- is typically defined as the Landau length, indicating the closest distance in interaction and b_+ is the Debye screening length.

Even before specifying the rf dynamics in detail, Eq. (5.11) shows that heating rate should depend strongly on how the cloud density changes with temperature.

To estimate the number density, we distinguish two limiting regimes. In the thermal regime, the cloud size along the dimension $i \in \{x, y, z\}$ is set by the thermal energy in the harmonic pseudopotential,

$$x_{\text{rms},i} \sim \sqrt{\frac{k_B T}{m\omega_i^2}}, \quad (5.12)$$

so the cloud expands as $T^{1/2}$. By contrast, in the mean-field regime, the size is set by the balance between Coulomb repulsion and the confining force. In this case, the characteristic cloud radius scales as

$$R \sim \left(\frac{e^2 N}{4\pi\epsilon_0 m\omega_r^2} \right)^{1/3}, \quad (5.13)$$

where ω_r is the confinement frequency and N is the ion number.

Equating Eq. (5.12) and Eq. (5.13) gives a crossover temperature,

$$T_{\text{crossover}} \sim \frac{(m\omega_r^2)^{1/3}}{4k_B} \left(\frac{e^2 N}{4\pi\epsilon_0} \right)^{2/3}. \quad (5.14)$$

This temperature marks the point where thermal energy and the mean-field energy are comparable. For $T \ll T_{\text{crossover}}$, the cloud is approximately mean-field dominated and its density is nearly constant. For $T \gg T_{\text{crossover}}$, the cloud expands thermally and the density falls. Putting this together, the density dependence is given by:

$$n(T) = \begin{cases} n_{\text{max}} & : T \leq T_{\text{crossover}}, \\ n_{\text{max}} \left(\frac{T_{\text{crossover}}}{T} \right)^{3/2} & : T \geq T_{\text{crossover}}. \end{cases} \quad (5.15)$$

Here, n_{max} denotes the maximum ion number density supported by the trap in the mean-field, $T = 0$ K limit

$$n_{\text{max}} = \frac{\epsilon_0 m}{e^2} (\omega_x^2 + \omega_y^2 + \omega_z^2).$$

In other words, it is the density for which the cloud's Coulomb force exactly balances out the harmonic confining force.⁹

⁹ For all these calculations, I'm implicitly assuming a spherically symmetric confining potential with $\omega_r \approx \omega_x \approx \omega_y \approx \omega_z$. Calculating with no spherical symmetry turns out to be pretty tricky.

The exact numerical factors in all of these estimates depend on precisely how quantities such as the cloud size, density, velocity, and crossover temperature are defined. For the present discussion, however, the important point is the scaling behavior.

This simple model of heating rate arrives at the same scaling used by Chen *et al.* in their treatment of the rf heating. In their notation, the secular temperature obeys [61]

$$\dot{T}_{\text{sec}} = \frac{e^4 \rho_i(T_{\text{sec}}) \ln \Lambda}{4\pi \epsilon_0^2 \sqrt{m} (k_B T_{\text{sec}})^{3/2}} \bar{\epsilon} (T_{\text{sec}} + \eta T_{\text{ex}}), \quad (5.16)$$

where ρ_i is the ion number density, $\ln \Lambda$ is the Coulomb logarithm, $\bar{\epsilon}$ is the average fractional energy increase per collision. The term T_{ex} represents excess micromotion energy, and η relates the total and secular kinetic energy. Their density model [61],

$$\rho_i(T_{\text{sec}}) = \begin{cases} \rho_{\text{max}} & : T_{\text{sec}} \leq T_p, \\ \rho_{\text{max}} \left(\frac{T_p}{T_{\text{sec}}} \right)^{3/2} & : T_{\text{sec}} \geq T_p, \end{cases} \quad (5.17)$$

has the same structure as Eq. (5.15). Although our notation differs, T_p in Ref. [61] plays the same role as $T_{\text{crossover}}$ discussed above, up to some constants, but the scaling is the same.

Written this way, the model presented in Chen *et al.* is not mysterious. The factor multiplying $\bar{\epsilon}$ is essentially the ion-ion collision rate, while the piecewise density function follows from the crossover between the mean-field and thermal regimes. The remaining rf-specific physics is contained primarily in $\bar{\epsilon}$, the average fractional increase in ion energy per collision. Interestingly, rf heating is generally strongest for dense and cold clouds even though the energy scale associated with a single collision usually decreases, because the collision rate increases rapidly as the cloud is cold and dense. Increasing the ion number or tightening the confinement tends to increase the density and therefore increase the heating, whereas increasing the temperature at fixed confinement suppresses it by lowering the collision rate.

The quantity $\bar{\epsilon}$ in Ref. [61] is intended to capture the rf-specific energy gain per collision. The concept of $\bar{\epsilon}$ itself is not new: as Chen *et al.* note, earlier work such as Refs. [72, 73] analyzed collisions in rf traps at the microscopic level. In Ref. [61], they embed that collision-induced energy gain into a many-ion plasma heating model.

Specifically, in their supplemental material, Chen *et al.* consider a one-dimensional elastic collision model for two ions, derive the fractional energy increase per collision, then numerically evaluate the resulting expression as a function of the Mathieu parameter q . For three-dimensional motion, they obtain $\bar{\epsilon}$ to be

$$\bar{\epsilon} \approx \frac{2}{3}(1 + 2q^{2.24}). \quad (5.18)$$

This $\bar{\epsilon}$ quantity is where our numerical results begin to differ from a naive reading of this expression. Over the parameter range relevant to nominal operation of the prototype, our results suggest that, over the parameter range relevant for our nominal operations, the leading rf-induced contribution to the fractional heating rate goes as $\approx q^2$, rather than as $(1 + 2q^2)$.

5.4.2.2 Numerical simulation of rf heating

To study the rf-heating scaling more directly, an undergraduate student Rohan Kompella carried out molecular dynamics simulations using LAMMPS.¹⁰ with guidance from my advisor and me. The results presented in this section were produced by Rohan.

The basic strategy was to scan a broad region of parameter space. We varied the ion number density n , temperature T , and secular frequencies ω_x, ω_y , and ω_z so as to access many different collision rates Γ_{coll} and different $T/T_{\text{crossover}}$ (for different regimes), and then varied ω_{rf} to change the Mathieu parameter q . Because we looked at a large region in parameter space, rather than a single realization of the collision rate, this allowed us to test the relationship between the observed heating rate and the collision rate in a more robust way.

Motivated by the scaling arguments above, we compare the simulated fractional heating rate (FHR)

$$\text{FHR} = \frac{\dot{T}}{T},$$

to the ansatz

$$\text{FHR}_{\text{model}} = c\Gamma_{\text{coll}}q^2, \quad (5.19)$$

¹⁰ LAMMPS is a simulation package developed by Sandia National Laboratories and stands for Large-scale Atomic/Molecular Massively Parallel Simulator.

where Γ_{coll} is the “naive”¹¹ collision rate defined in Eq (5.11), evaluated using the simulation parameters, and c is a fixed numerical normalization. Rather than focusing on the absolute numerical prefactor, we define the dimensionless ratio

$$A = \frac{\text{FHR}_{\text{sim}}}{\text{FHR}_{\text{model}}}. \quad (5.20)$$

If the ansatz $\text{FHR} \propto \Gamma_{\text{coll}} q^2$ captures the dominant dependence, then A should remain approximately constant as the simulation parameters are varied. Furthermore, by looking at whether or not A changes, we can see any dependence of A on a specific parameter.

For each run, the ions were initialized from either a thermal distribution or a mean-field distribution or a combination of both, depending on the value of $T/T_{\text{crossover}}$. The system was then evolved for 1 s, and the heating rate was extracted from the temperature evolution. An example of the temperature evolution is shown in Fig. 5.14. The top panel shows the total energy as a function of time, and the bottom panel shows the corresponding cloud temperature, averaged across the three dimensions, with a linear fit used to extract the heating rate.

To isolate the q dependence, we first compared simulations at three different rf frequencies while keeping the secular frequencies approximately fixed. For this particular dataset, the cloud size and collision rate remain nearly unchanged, and only q is varied.¹² The results are summarized in Table 5.1.

The measured FHR decreases strongly as q is reduced, and the values are consistent with an approximately quadratic dependence on q . The ratio A remains of order unity and does not show a clear trend with q , even though the FHR changes by an order of magnitude. This indicates that the leading dependence is well captured with the q^2 ansatz.

It is worth noting that the scaling factor $\bar{\epsilon}$ in Chen *et al.* of Ref. [61] would be

$$1.117, 1.025, \text{ and } 1.005$$

¹¹ We refer to this as the naive collision rate because it is based on a simple scaling model of the collision rate, and the numerical prefactor is not treated rigorously.

¹² For different rf frequencies, V_{rf} was adjusted so as to keep the secular frequencies (and thus the density, temperature, and Coulomb logarithm) fixed.




Figure 5.14: Representative example of heating-rate extraction fomr a LAMMPS simulation. The top panel shows the total energy versus time. The bottom panel shows the cloud temperature averaged over all three directions and the linear fit used to extract the heating rate.

	25 kHz	50 kHz	100 kHz
Heating rate (K/s)	3.22(22)	0.616(26)	0.17(6)
FHR (1/s)	0.644(45)	0.123(5)	0.034(13)
Model FHR (1/s)	0.3690(3)	0.09380(6)	0.02350(1)
Naive collision rate (1/s)	4.653(4)	4.729(3)	4.741(3)
Initial ion density (m^{-3})	$1.520(1) \times 10^{11}$	$1.545(1) \times 10^{11}$	$1.549(1) \times 10^{11}$
q	0.282	0.141	0.0704
$A = \text{FHR}_{\text{sim}}/\text{FHR}_{\text{model}}$	1.745	1.311	1.438

Table 5.1: Simulation results for three rf frequencies with approximately fixed secular frequencies $(f_x, f_y, f_z) = (2275, 1955, 1840)$ Hz, $T = 5$ K, $N \approx 1000$, and $T/T_{\text{crossover}} \approx 1.82$. Each point is the average of four independent random seeds and uncertainties are standard errors. The measured fractional heating rate is consistent with an appxoimately quadratic dependence on q .

for the q values of $q = 0.282$, 0.141 , and 0.0704 . Normalizing this scaling to the measured 50 kHz point, this scaling law would have given the fractional heating rate of

$$0.131, 0.123, \text{ and } 0.121 \text{ s}^{-1},$$

for 25, 50, and 100 kHz respectively. This is different from what we measured in Table. 5.1.

We next checked whether the same scaling remains valid as the ion number is changed. Table 5.2 shows the results for $N = 500, 1000$, and 2000 ions at fixed $q = 0.141, T = 5$ K, and nearly fixed secular frequencies. As expected, increasing the ion number increases the ion density and the naive collision rate and lowers $T/T_{\text{crossover}}$. The measured heating rate increases accordingly. At the same time, the ratio A remains within a factor of order unity, again indicating that changing the ion number does not introduce a qualitatively new scaling beyond the $\Gamma_{\text{coll}}q^2$ in our model, at least within the present statistical precision.

Number of ions	500	1000	2000
$T/T_{\text{crossover}}$	2.848	1.817	1.139
Heating rate (K/s)	0.61(19)	0.62(5)	1.22(36)
FHR (1/s)	0.12(4)	0.123(10)	0.24(7)
Model FHR (1/s)	0.05580(4)	0.09380(6)	0.14500(14)
Naive collision rate (1/s)	2.813(2)	4.729(3)	7.325(7)
Initial ion density (m^{-3})	$9.193(7) \times 10^{10}$	$1.545(1) \times 10^{11}$	$2.393(2) \times 10^{11}$
q	0.141	0.141	0.141
$A = \text{FHR}_{\text{sim}}/\text{FHR}_{\text{model}}$	2.186	1.311	1.683

Table 5.2: Simulation results for different ion numbers at fixed $q = 0.141, T = 5$ K, and approximately fixed secular frequencies $(f_x, f_y, f_z) = (2276, 1956, 1840)$ Hz. Each point is the average of four independent random seeds, and the uncertainties are standard errors.

To test the robustness of this conclusion over a wider range of parameters, Rohan carried out additional runs spanning roughly an order of magnitude in $T/T_{\text{crossover}}$. This was done by varying combinations of the ion number, the temperature, the secular frequencies, and in some cases, an artificially boosted ion charge.¹³ The results are summarized in Table 5.3.

Across this broader parameter scan, the measured FHR values continue to show the expected dependence on q . Figure 5.15 is a tabulated plot of the corresponding behavior of A across the

¹³ In these “boosted-charge” simulations, the ions were assigned a larger charge while keeping the same mass, so that the Coulomb interaction strength was enhanced.

FHR (1/s)	$q = 0.282$	$q = 0.141$	$q = 0.0704$
$T/T_{\text{crossover}} \approx 5.5$	0.120(19)	0.030(11)	0.0076(36)
$T/T_{\text{crossover}} \approx 1.86$	1.50(5)	0.362(29)	0.049(12)
$T/T_{\text{crossover}} \approx 1.82$	0.644(45)	0.123(5)	0.034(13)
$T/T_{\text{crossover}} \approx 0.55$	4.92(23)	1.19(6)	0.187(43)
$A = \text{FHR}_{\text{sim}}/\text{FHR}_{\text{model}}$	$q = 0.282$	$q = 0.141$	$q = 0.0704$
$T/T_{\text{crossover}} \approx 5.5$	0.89(14)	0.89(32)	0.90(42)
$T/T_{\text{crossover}} \approx 1.86$	1.04(4)	0.98(8)	0.53(13)
$T/T_{\text{crossover}} \approx 1.82$	1.74(12)	1.31(54)	1.44(54)
$T/T_{\text{crossover}} \approx 0.55$	0.94(4)	0.87(4)	0.53(12)

Table 5.3: Scan of the simulated fractional heating rate over a wider range of $T/T_{\text{crossover}}$. The rows combine simulations with the similar $T/T_{\text{crossover}}$ but with different ion number, temperature, secular frequencies, and in some cases boosted ion charge. Each entry is an average over 4 to 16 random seeds, and uncertainties are standard errors.

simulation space. The ratio A shows noticeable scatter and may retain some residual dependence, however we do not find evidence for a strong or systematic dependence of A on the microscopic simulation parameters beyond that associated with statistical fluctuations.

Taken together, these simulations support the conclusion that the leading rf-induced contribution to the fractional heating rate is well described by

$$\frac{\bar{T}}{T} \propto \Gamma_{\text{coll}} q^2, \quad (5.21)$$

especially for the parameter range relevant to the prototype apparatus.

5.4.2.3 rf heating experiment

Motivated by the simulation results above, we carried out a few experiments to see whether the heating observed in the prototype apparatus is in fact dominated by rf heating.

In principle, the model discussed above provides a way to estimate the expected rf-induced heating rate from the cloud temperature, ion number, secular frequencies, and rf frequency. In practice however, quantitatively calculating the heating rate to the present experiment is difficult. As discussed earlier, obtaining an absolute temperature for the ion cloud is challenging in our apparatus. Additionally, determining the absolute ion number is also not straight-forward because the detection efficiency is not well known. For these reasons, it is difficult to make a reliable absolute prediction for the rf-induced heating rate in the prototype.

Instead, a more robust approach is to perform a differential measurement. We vary the rf frequency, and thus the Mathieu parameter q , while keeping the secular frequencies nearly fixed. We then compare how the measured heating trend changes. Figure 5.16 shows a width-based measurement for three rf frequencies: 25, 50, and 100 kHz, corresponding to $q = 0.282, 0.141$, and 0.0704 , respectively. The trap voltages were adjusted such that the secular frequencies remained nearly the same across the three cases.

As noted above, the absolute values on the vertical axes should not be interpreted too literally as the true ion temperature. The more relevant question is whether the heating rate trend changes




Figure 5.15: Robustness of the dimensionless factor $A = \text{FHR}_{\text{sim}}/\text{FHR}_{\text{model}}$ across the simulated parameter space. The main panel shows A as a function of temperature, while the inset panels separate the data into the $T = 5, 10$, and 15 K subsets and display A as a function of ion number N , initial ion density, q , and $T/T_{\text{crossover}}$. The scatter is appreciable between points, but there is no clear evidence for a strong systematic dependence of A on any single simulation parameter.




Figure 5.16: Width-based heating measurement for three rf frequencies 25, 50, and 100 kHz, corresponding to $q = 0.282, 0.141$, and 0.0704 . The inferred temperatures in the x , y , and z directions are plotted versus trapping time, together with the corresponding ion number. The trap voltages were adjusted so that the secular frequencies remain the same between different rf frequencies. Because these temperatures are inferred from detector widths, the absolute scale should not be interpreted literally. Instead, the relevant comparison is the relative trend with trapping time: the 25 kHz data show somewhat stronger heating, especially along z , whereas the 50 and 100 kHz traces differ only slightly.

systematically as q is varied. A somewhat larger trend is visible at $f_{\text{rf}} = 25$ kHz, especially in the z -direction. Since the rf confinement is applied in the radial plane. If the transfer function in the radial direction carries an unknown systematic that changes with rf frequency, that systematic would affect the x and y widths more strongly than the axial width. Therefore, I place somewhat more weight on the z -axis result here.

At the same time, the difference between 50 and 100 kHz is comparatively small. This suggests that at our nominal operating point of $f_{\text{rf}} = 50$ kHz, rf heating is likely not the dominant contribution to the observed heating. In other words, lowering the rf frequency to 25 kHz appears to make rf-related heating more important, but increasing the frequency beyond that nominal value does not dramatically change the measured heating trend. This behavior is consistent with the idea that other mechanisms dominate the heating under normal operating conditions. This conclusion should still be regarded as qualitative, since the width-based thermometry itself is subject to detector-ion-optics systematics.

Nonetheless, operating at a higher rf frequency remains compelling, and in future work, we plan to operate $f_{\text{rf}} = 100$ kHz.

5.4.3 Ion-neutral collisions

Ion-neutral collisions can also contribute to trap heating and eventually to trap loss. As discussed in Section 5.4.2, the heating rate can be modeled heuristically as Eq. 5.10,

$$\dot{T} \sim \Gamma_{\text{coll}} \times (\text{energy gain per collision}).$$

For the collisions with the neutrals, we have

$$\dot{T} \sim \Gamma_L \frac{\Delta E}{k_B} \tag{5.22}$$

where Γ_L is the Langevin collision¹⁴ rate discussed in Section 4.2.2 and ΔE is the characteristic energy transferred to the ion in a single collision [67, 74]

¹⁴ Again, the heating rate here should also go as “Langevin-like” collisions. We don’t expect the heating rate to follow exactly like the Langevin collision rate, which tends to underestimate the total rate of elastic collisions that transfer energy to the ion.

Without the rotating electric field, the temperature of the ions is ~ 30 K. Therefore the background neutrals are much hotter than the trapped ions, with the temperature ~ 300 K. In this case, the elastic collisions tend to transfer energy into the ion cloud.

In a simple elastic two-body picture, the maximum fraction of the incident kinetic energy transferred in a head-on collision is

$$\frac{4m_{\text{ion}}m_{\text{gas}}}{(m_{\text{ion}} + m_{\text{gas}})^2}.$$

If the ion is approximately at rest and the neutral carries the incident thermal energy, then a crude estimate for the energy transferred per collision is

$$\Delta E \sim \frac{4m_{\text{ion}}m_{\text{gas}}}{(m_{\text{ion}} + m_{\text{gas}})^2} k_B T_{\text{gas}}, \quad (5.23)$$

up to some numerical factors associated with the collision angles and some thermal distribution. In this simple picture, the energy transfer is most efficient when the neutral mass is closer to the ion mass.

Figure 5.17 shows the estimated heating rate using Eqs.(5.22) and (5.23). Compared to Fig. 4.7 which shows only the Langevin collision rate Γ_L , the heating rate estimate makes clear that H_2 is relatively inefficient at heating ThF^+ because of the large ion–neutral mass mismatch, whereas heavier gases like Kr are substantially more efficient.

To experimentally test to see if the observed heating is caused by ion–neutral collisions, we deliberately inject known amounts of gas into the chamber and measure the heating rate, using the same general approach as in the the $J = 1$ lifetime measurement. Figure 5.18 shows the measured heating rates with added N_2 , H_2 , and Kr with no rotating electric field.

Again, the absolute scale on the y -axis should not be taken too literally because the heating rates were extracted using the width-based thermometry and are therefore subject to systematics. Nevertheless, differential comparisons between gases are still informative. At comparable partial pressures, Kr produces the strongest heating, while H_2 is the least effective, consistent with the simple mass-transfer estimate.

For all three gases, a change in heating rate required deliberate gas injection to pressures




Figure 5.17: Estimated heating rate of ThF^+ ions due to collisions with background neutrals, calculated from the Langevin collision rate Γ_L and the elastic collision energy-transfer model. Although H_2 has a sizable collision rate (shown in Fig. 4.7), its large mass mismatch with ThF^+ makes it inefficient at transferring energy, whereas heavier gases such as Kr and SF_6 are predicted to heat the ions more effectively given the same partial pressure.




Figure 5.18: Measured mean heating rate (averaged over all three directions) of ThF^+ ions as a function of inferred partial pressure for added N_2 , H_2 , and Kr with no rotating electric field. The first point in each dataset corresponds to the baseline partial pressure before additional gas is injected. Although the absolute heating-rate scale is affected by systematics in the width-based thermometry, the differential trend is clear. The heavier gas Kr produces the strongest heating, while H_2 is the least effective at comparable pressures. For all gas present in this plot, we see increasing partial pressure by a factor of ten beyond their baseline values does not much change heating, therefore, the baseline pressures of these gases are not the cause of heating we observe in the prototype.

much larger than the nominal operating conditions. This suggests that, in the absence of the rotating electric field, ion–neutral collisions with H_2 , N_2 , and Kr are not the dominant source of heating under normal conditions.

We also carried out several measurements with the rotating electric field applied. In this case however, the interpretation becomes less straightforward. Once the rotating field is on, the ions have substantial driven kinetic energy (~ 90000 K with $\mathcal{E}_{\text{rot}} = 60$ V/cm, $f_{\text{rot}} = 147.5$ kHz). Collisions can convert this ponderomotive energy into heating.

In practice, the data with the rotating field were difficult to interpret because there is significant ion loss in this configuration with the injected gas. Once the ions begin spilling out of the trap, it becomes difficult to separate heating from loss, and the loss rate itself is hard to quantify because of detector saturation. For this reason, the measurements with the rotating field did not lead to a clear conclusion.

However, we still have significant ion heating even *without* the rotating electric field. For the no rotating field case, the data suggest that ion–neutral collisions are not the dominant source of the observed heating.

5.4.4 Noise on electrodes

Noise on the electrodes can also drive ion heating. A useful treatment of heating from thermal noise in ion traps is given in Ref. [75]. In our system, electrode noise can appear either as a differential signal between the two rods in a given pair or as a common-mode signal on that pair. Differential-mode noise can resonantly drive the ions’ secular motion when it has noise near the secular frequency. By contrast, common-mode noise can drive parametric heating when it is present near twice the secular frequency.¹⁵

In characterizing the voltage noise on the rods, we found appreciable differential-mode noise

¹⁵ Noise can also be present on the rings. As discussed before, the special and quadrant rings are driven through the power op-amps and therefore have noise levels comparable to those of the rods. By contrast, the translation rings are driven directly by the ARTIQ DAC system and are substantially quieter. In many of the long-coherence and heating measurements, the ions are translated into the Science region, where the relevant electrodes are primarily the rods and the translation rings. For this reason, I focus on the noise performance of the rods.

on all four rod pairs and a much larger common-mode noise on one particular pair. Figure 5.19 shows the measured differential- and common- mode noise between all the rod pairs. The differential-mode noise is present on all pairs, with prominent peaks near 2.6 and 2.9 kHz. The common-mode noise is strongly dominated by the Rod 2–6 pair with a dominant peak around 2.6 kHz.

Using the measured noise amplitudes, we carried out a simple analytic estimate of the corresponding heating rates. For the differential-mode peaks near 2.6 and 2.9 kHz, the expected heating rate was of order 11 K/s. The large common-mode peak gave a much larger value, of order 1800 K/s.¹⁶ Although these estimates are crude, they suggested that the observed common-mode feature could plausibly be relevant for the measured heating.

To test whether the ions were sensitive to this noise, we scanned the secular frequencies by varying the end-cap and rf voltages. In particular, we scanned through values for which the secular frequency approached half of the common-mode noise frequency, around 1.3 kHz, where parametric heating would be expected. The left column in Fig. 5.20 shows the measured ion-cloud widths of the ion cloud in x , y , and z as the rf voltage was varied. The right column shows representative detector images (for x and y) and time-of-flight traces (for z) of the ion cloud taken on and off resonance.

A clear increase in cloud size is observed near resonance, consistent with excess heating driven parametrically by the common-mode noise on the Rod 2–6 pair. At this point, the measured secular frequencies were approximately [1332(7), 1400(20), 1360(30)] Hz for x , y , and z , respectively, which is consistent with the ions lying near half of the 2.6 kHz common-mode noise feature shown in Fig. 5.19.

On the other hand, when we scanned near the frequencies where differential-mode heating would be expected, we did not observe comparably clear evidence for resonant heating.

We eventually traced the large common-mode noise to a problem with the op-amp board driving that particular rod pair. After repairing the board, we repeated the measurement in the

¹⁶ These estimates should be interpreted only at the order-of-magnitude level. They depend on several approximate assumptions, so the precise numerical values are not meant to be taken too literally.




Figure 5.19: Measured differential- and common-mode voltage noise spectra on the four pairs. Measured by a fellow graduate student Anzhou Wang. **Top:** differential-mode noise between each rod pair. Prominent features are visible near 2.6 and 2.9 kHz on all rod pairs. **Bottom:** common-mode noise on each rod pair. The Rod 2-6 pair shows a much larger common-mode feature than the others, indicating a technical noise source specific to that channel.




Figure 5.20: Ion-cloud widths measured while scanning the rf voltage through a region where the secular frequencies approach half of the common-mode noise feature. **Left:** measured widths in x , y , and z versus rf voltage. **Right:** representative detector images and time-of-flight traces taken near and away from resonance. Near resonance, the cloud becomes significantly broader, consistent with parametrically driven heating from common-mode electrode noise. Here, the secular frequencies are about ~ 1300 Hz, in all three directions.

same region of parameter space and could no longer reproduce the anomalous heating.¹⁷ This suggests that the previously observed resonance was caused by technical common-mode noise on that electrode pair and that this particular contribution was successfully mitigated.

5.4.5 Other possible contributions

Several additional effects may also contribute to heating. In this section, I briefly discuss a few effects that we investigated at a preliminary level, but for which we found no obvious evidence that they were the dominant contributors to the heating rate observed in the current apparatus.

(1) Residual electrode noise

The measurements above show that electrode noise can, in some circumstances, produce heating in the ion cloud. However, the simple estimate based on the measured voltage-noise spectra still contain significant uncertainty, and it remained possible that smaller, broadband, technical noise on the electrodes could contribute to the nominal heating rate. To test this possibility more directly, we added additional filters to each of the rod electrodes. We then compared the measured heating trend with and without the filters in place.

In this test, filters could be added only to the rod electrodes, not to the rings, since the ring electrodes must preserve the high-frequency components needed for stop and kickout voltages. Within our experimental sensitivity, we observed no obvious change in the heating trend after adding the rod filters.¹⁸ This result suggests that residual rod-electrode noise, at least at the level addressed by these added filters, was not the dominant limitation in the nominal operating configuration.

(2) Collisions with ions of different m/z

If another ionic species with a different value of m/z is present in the trap, it can collide

¹⁷ However, we did see a new peak that seemed to be moving around a little bit in terms of the frequency. We suspect that this new peak was coming from the rings, because when we translate into the Science region, it was no longer there.

¹⁸ For the experiments with the filters, we looked at the heating trend of the ions inside the Science Region, where the rings have much less noise, compared to at the octagon.

with our ion, ThF^+ . This collision may lead to heating, since the two species have different micromotion trajectories. We know that a heavier ionic species is present in the trap, with a mass approximately 10% larger than that of ThF^+ .¹⁹ To test whether this impurity species was responsible for the observed heating, we implemented a mass filter by slowly ramping the Mathieu parameters a and q until species with m/z roughly 10-15% above that of ThF^+ were removed from the trap. We observed no obvious change in the heating rate with the filter applied, suggesting that collisions with this heavier species were not the dominant source of heating under our operating conditions.

(3) Noise associated with high-power lasers

Our experiment uses high-power YAG lasers. We have observed that when these lasers are not well isolated from the optical tables, noise at 10 Hz, the laser repetition rate, can appear in various electronic signals, including oscilloscope traces of the ring voltages. One of these YAG lasers is used for ionization and therefore cannot be turned off during normal operation. A second, higher-power YAG laser (“Cow”) is used for dissociation, and this can be turned off when dissociation is not required. We compared the heating rate with Cow on and off. No obvious difference was observed. However, this test was performed using the width-based temperature estimate, and the on/off comparison was implemented as a relatively slow chop. As a result, the measurement may have been susceptible to drifts in ion number over the course of the experiment, and therefore this null result should be interpreted with some caution.

(4) Heating rate from excess micromotion energy

As discussed in Ref. [71], ions displaced from the rf null can acquire excess micromotion energy. In our apparatus, ions can be displaced, for example, from stray dipole fields produced by patch charges. In routine operation, we make an effort to center the ions in the

¹⁹ We think that this species is either ThF_2^+ or ThFO^+ . Our detection cannot resolve the time of flight of these two candidates.

trap by checking micromotion-related signatures during ejection.²⁰ . We therefore have no strong reason to believe that excess micromotion is the dominant limitation in the nominal configuration. Nonetheless, to test this possibility more directly, I performed a simple experiment in which the ions were intentionally displaced from the trap center by applying a dipole field and allowing them to evolve for different hold times. A modest increase in the heating trend was observed as the applied dipole field was increased. However, when extrapolated back to zero intentional dipole field²¹ , the heating rate approached a nonzero asymptote. This behavior suggests that under nominal operating conditions, the ions are already located close to the rf null and that excess micromotion from static displacement is not the primary source of the observed heating.

The purpose of this section is not to claim that these effects produce no heating at all. Rather, it is to document the additional mechanisms that we considered and tested at a preliminary level, and for which we found no obvious evidence that they were the dominant limitation in the present measurements. If heating remains a concern in future work, each of these possibilities may require a more systematic and quantitative investigation.

5.5 Mitigation and lessons from the prototype

When the mechanisms discussed in this chapter become sufficiently strong, ions acquire enough motional energy to escape the trap. We have identified several knobs that affect the trap lifetime. For instance, increasing f_{rot} reduces the ponderomotive energy associated with $\mathcal{E}_{\text{rot}}$, and decreasing noise on the rf electrodes improves the lifetime by suppressing rf-driven heating. We

²⁰ To shim the ions to the rf null, we abruptly turn off the rf voltages during kickout. The switch-off phase is chosen such that the ions have either the maximum or minimum micromotion velocity, corresponding to opposite directions, e.g. $\pm\hat{x}$. This process maps the instantaneous micromotion velocity directly onto the kickout velocity. Although the cloud shape changes under this procedure, the cloud centroid provides a sensitive diagnostic of displacement from the rf null—if the ions are centered at the rf null, the centroids for the two opposite micromotion phases should land at the same point on the detector. On the other hand, if the ions are displaced, the centroid shifts in opposite directions on the detector for the two cases. We therefore use the measured centroid difference to apply a compensating dipole shim field, correcting stray fields such as patch charges and returning the ions to the rf null.

²¹ Here, zero intentional dipole field refers to the field that we apply to “shim” the patch charges through the process discussed earlier

studied each of these contributions in some detail and observed their influence on the ion lifetime. We have not yet, however, isolated a single mechanism that can account for the current ion loss.

Part of this difficulty is that the temperature diagnostic used throughout this chapter is sensitive to ion number. While the transfer-function method remains useful for relative comparisons, many of the mechanisms studies were carried out on different days and under somewhat different loading conditions. In addition, the production rate itself often fluctuated substantially from shot to shot, so the number of trapped ions could vary by a factor of two, or occasionally much more. These variations, combined with the ion-number dependence of the width-based thermometry, imply that some of our interpretations are likely affected by systematic uncertainty.

We therefore spent a significant amount of time trying to disentangle the ion optics, transfer function, and the actual temperature of the ion cloud. This effort diverted our attention from the physics goals, but it clarified an important limitation of the prototype apparatus and motivated improvements to both the detector and the ion optics. With the upgraded detection system and a better understanding of the transfer function, we expect future measurements to provide a more reliable thermometry of the ion ensemble.

With that improved thermometry in place, the ion loss channels can be revisited in the full Gen. III apparatus with much better control over systematics. I expect that the next stage of this work will be able to determine more decisively which mechanisms are limiting the ion lifetime, and in turn what changes are required to achieve the long interrogation times.

Chapter 6

Towards Long Coherence Time: Alternative Trapping Configurations

In the previous chapter, we investigated some mechanisms that limit the trap lifetime. In practice, even without a complete understanding of all the mechanisms, we can take a different strategy. Rather than focusing only on suppressing the loss channels, we can also reshape the trapping potential so that when the confinement is loosened for the quantum evolution, the effective trap depth does not decrease dramatically. In practice, this means making the trap more nearly harmonic over the spatial extent sampled by the ion cloud.

In the following sections, I discuss the effective trapping potential of the prototype quadrupole trap geometry and then describe several alternative trapping configurations that we tested in the prototype apparatus.

6.1 Effective trap depth

As shown in Fig. 2.10, the prototype trap uses eight rods to provide the rf radial confinement and a set of ring electrodes to provide static axial confinement. Because this geometry was not designed to realize an ideal harmonic trap, the resulting potential contains substantial anharmonic contributions.

As discussed in Chapter 2, for the radial confinement, the rods are driven in the 4×2 phase pattern

$$\{0, 0, \pi, \pi, 0, 0, \pi, \pi\}.$$

Naturally, this phase pattern creates pseudopotential that varies in the azimuthal direction. In

particular, the direction with the shallowest trap depth is along the bisector between two adjacent rods that are driven in phase.

For the axial confinement, voltages are applied to the rings in a cosine-like potential along z . Near $z = 0$, the potential is harmonic to lowest order. However, the finite spacing and the discreteness of the ring electrode structure introduce significant higher-order terms. In the prototype apparatus, in one bucket, there are four rings with a spacing of 2 cm, and these non-ideal terms are not negligible.

These axial fields also produce radial anti-confinement. Since the static potential Φ must satisfy Laplace's equation,

$$\nabla^2 \Phi = 0, \tag{6.1}$$

any electrostatic confinement along z must be accompanied by electrostatic anti-confinement in the radial plane. More generally, the higher-order axial terms from the discrete ring geometry result in higher-order terms in the radial direction with opposite sign, through Laplace's equation. As a consequence, the ring voltages not only confine the ions axially, but also reduce the effective radial trap depth away from the trap center.

Figure 6.1 illustrates these different contributions calculated from COMSOL simulation for the prototype trap geometry. The top-left panel shows the rf pseudopotential from the rf 4×2 configuration along the shallowest radial direction. The top-right panel shows the radial anti-confining contribution from the ring electrodes along that same direction. The bottom-left panel shows the axial confinement from the rings. The bottom-right panel shows the resulting radial effective potential after the rf confinement and ring-induced radial anti-confinement are added together. In this figure, the voltages are chosen to give approximately the nominal secular frequencies, $f_x, f_y, f_z \approx 2$ kHz with $f_{\text{rf}} = 50$ kHz.

To understand how the bucket geometry affects the anharmonicity, we recall that the axial




Figure 6.1: Components of the effective trapping potential for the prototype, shown in temperature units U/k_B . **Top left:** rf pseudopotential from the rods along the shallowest radial direction. **Top right:** radial anti-confining contribution from the ring electrodes along the radial direction. **Bottom left:** axial confinement produced by the ring electrodes. **Bottom right:** net radial effective potential obtained by summing the rf confinement and the ring-induced anti-confinement. For this example, nominal trap parameters are used: $V_{rf} = 25.6 \text{ V}$ and $V_{ring} = 10 \text{ V}$, corresponding to secular frequencies of $\sim 2000 \text{ Hz}$.

potential has the form

$$V(z) = V_0 \cos(kz), \quad k = \frac{2\pi}{\lambda}$$

where λ is the spatial period of the bucket structure. Expanding about the trap center gives a harmonic term proportional to $k^2 z^2$ and a leading anharmonic correction proportional to $k^4 z^4$. Therefore, for fixed axial confinement, increasing λ suppresses the higher-order terms and makes the axial potential more nearly harmonic. The corresponding higher-order radial distortion is likewise reduced, so the effective radial trap depth can be substantially improved even though the radial anti-confinement required by Laplace’s equation remains.

This picture is consistent with what we observed experimentally in two ways. First, the trap lifetime was generally better when f_z was chosen to be significantly smaller than f_x and f_y . Although this weakens the confinement along z , it also reduces the radial anti-confining contribution from the rings, thereby increasing the overall radial trap depth, which appears to be the more relevant quantity in the prototype apparatus. Second, in the octagon region where the electrodes can be controlled individually¹, we implemented 8-ring and 10-ring bucket configurations and observed longer trap lifetimes than with the 4-ring configuration. These bucket configurations increase λ and reduce the higher-order distortion of the trapping potential.

Motivated by these observations, the Gen. III apparatus will use a larger bucket period in order to increase the effective trap depth and make the trap less sensitive to anharmonicity.

6.2 Octupole trap

We have just seen that increasing the bucket period helps by reducing the axial anticonfinement. In the present apparatus, however, the ring geometry is fixed, especially in the cold region—changing the ring periodicity in this region would require breaking vacuum and rewiring the trap. We therefore asked a more practical question: given the existing electrode geometry, can we modify the radial confinement in a way that improves ion survival at long trapping times?

¹ We call these rings “special” rings.

For many aspects of trap dynamics and systematics studies, a (nearly) harmonic potential is desirable. For the immediate goal of retaining ions at low confinement, however, a non-harmonic potential can be advantageous. With the current eight-rod geometry of the prototype apparatus, we can implement an octupole rf configuration. In the usual quadrupole configuration, the rod phases are

$$\{0, 0, \pi, \pi, 0, 0, \pi, \pi\},$$

whereas an octupole pattern is obtained with

$$\{0, \pi, 0, \pi, 0, \pi, 0, \pi\}.$$

In the eight-rod geometry, this pattern produces an octupolar rf potential $\Phi_{\text{rf}} \propto r^4 \cos(4\theta)$ and an rf pseudopotential that scales as r^6 .

In practice, there are two broad ways of using the octupolar confinement, as shown in Fig. 6.2.




Figure 6.2: Illustrative radial pseudopotential profiles for two octupole-assisted trapping schemes in the prototype apparatus. **Left:** octupole-box configuration (blue) compared with the conventional rf configuration (orange). **Right:** octupole-wall configuration (blue) compared with the conventional rf configuration (orange). In both configurations, the outer confinement is increased while allowing relatively weak confinement near the trap center. Both configurations also include the nominal axial anti-confinement component.

In the “octupole box” configuration, the potential is more like a box than a harmonic well—ions are able to spread out more, thus reducing ion density required for the Ramsey evolution. In the “octupole wall” configuration, the octupolar pattern is used in combination with the usual rf

trap. Near the trap center, the potential is dominated by the harmonic well from the rf trap, but the octupolar trap provides steeper walls at larger radius, thereby preserving trap depth even when the secular frequencies are reduced.

Empirically, the octupole-wall configuration proved easier to optimize and more robust in operation than the octupole-box configuration.

To quantify the benefit, I measured the ion signal remaining after 12 seconds of the trapping time while scanning the rf amplitude V_{rf} . The results are shown in Fig. 6.3. The blue trace corresponds to the conventional rf trap, while the orange trace corresponds to the rf trap with an added octupole-wall component. The plotted quantity is the area under the time-of-flight signal, which is proportional to the number of ions in the trap.




Figure 6.3: Remaining ion signal after 12 s versus V_{rf} for the standard rf trap and for the rf trap with an octupole-wall component. Adding the octupole wall allows for lower confinement to be achievable by improving ion retention at lower V_{rf} .

The main effect is that the octupole-wall configuration allows for lower secular frequencies while still retaining substantial ion signal after 12 s. In contrast, the conventional rf configuration shows strong loss over the same V_{rf} range. Therefore, the octupole-wall configuration with the conventional rf trap provides a practical way to maintain ion survival in the low-confinement regime that is most relevant for long Ramsey evolution.

With the octupole-wall configuration, I was able to perform long-time Ramsey spectroscopy while retaining more ions compared to the conventional rf trap. Particularly, this configuration allowed us to directly observe Ramsey fringes at 21 s for the first time.

Looking ahead, the rod-drive upgrades and the ability to operate at higher rotating-field frequency should further suppress several of the loss mechanisms discussed in the previous section. Nevertheless, the octupole-wall configuration remains a useful and flexible tool for improving ion retention at long trapping times.

6.3 Rotating rf (rrf) ion trap

The octupole trap configuration is not the only alternative trapping configuration. At some point, our group came across the proposal of the rotating rf (rrf) ion trap by Hasegawa and Bollinger. For the parameters relevant to our experiments, the rrf configuration can increase the effective trap depth without requiring an increase in the nominal trap confinement.

Portions of this section are adapted from a manuscript in preparation. I have modified the manuscript to fit the narrative of this thesis.

6.3.1 Motivation and background

Radio-frequency (rf) ion traps, commonly referred to as Paul traps, allow for confinement in all three directions. Since three-dimensional ion confinement using purely electrostatic fields is not a solution of Laplace's equation, linear rf (lrf) ion traps utilize alternating confining and anticonfining potentials oscillating along the x and y axes as shown in Fig. 6.4a. The rapidly oscillating fields drive micromotion, and this micromotion in the field gradient results in an effective secular confining potential. Since their introduction in the 1950s by Wolfgang Paul [76], lrf ion traps have been used in many applications, including quantum computation [77–82], atomic clocks [83–87], and physical chemistry [88–90].

There is a familiar mechanical Paul trap lecture demonstration (videos available on the internet, for instance here in Ref. [91]) based on a rotating, saddle-shaped surface. An object




Figure 6.4: The landscape of the oscillating component of the potential of the (a) linear rf ion trap and the (b) “mechanical” Paul trap at different moments in time t . (a) Linear rf ion trap. The potential is confining (anticonfining) along the x -axis (y -axis) at $\omega_{\text{rf}} t = 0$. A quarter period later, the potential landscape is flat. Half a period later, the potential flips its sign. The potential flaps along the x and y axes. (b) Mechanical Paul trap. The potential rotates (in this case, about the $+\hat{z}$ direction). The same pattern appears in the rrf ion trap.

constrained to the surface by gravity and surface forces experiences a potential that rotates with time as in Fig. 6.4b. While this is not exactly how most rf ion traps work, the analogy provides useful physical intuition.

About 20 years ago, Hasegawa and Bollinger in Ref. [92] proposed an rf ion trap based on a rotating potential. In conventional lrf ion traps, the confining and anticonfining potentials flap back and forth at the rf frequency, whereas in the rotating rf (rrf) ion trap, the potential itself rotates at the rf frequency, more analogous to the mechanical lecture demonstration. Prior to this work, no experimental realization of an rrf ion trap has been reported in the literature. This could be because an rrf ion trap is experimentally more complex, requiring multiple phases of rf voltages on multiple electrodes.

In this section, we revisit this proposal and present an experimental realization of the rrf ion trap, verifying its fundamental dynamics and demonstrating the theoretical predictions outlined in Ref. [92]. Furthermore, we demonstrate practical implications of the rrf ion trap for precision metrology. Specifically, we show how the rrf ion trap configuration resolves inherent trap-depth anisotropies in eight-rod trap geometries used in precision spectroscopy on trapped molecular ions.

The sensitivity of the eEDM measurement is proportional to $\sqrt{N}\tau$ where N is the number of counts and τ is the coherence time. Therefore, to maximize sensitivity, we aim to trap as many ions as possible while suppressing unwanted ion-ion interactions by adiabatically reducing the secular frequency, which lowers the ensemble temperature and density. As mentioned earlier in the thesis, we distinguish between the *tightness* of ion confinement proportional to the secular frequency at the trap center and the *effective depth* of the trap set by the total energy required to escape the trap via the lowest saddle point. While loose confinement can be advantageous for reducing interaction-driven decoherence, it also makes the ions more vulnerable to confining-potential inhomogeneity: the effective trap depth becomes shallower, which increases ion loss. The rrf trap mitigates this tradeoff by averaging over certain static inhomogeneities and anharmonicities, yielding a deeper effective trap depth for the same nominal confinement, as will be discussed later in this paper.

6.3.2 Theory and Experimental Realization of the rrf trap

6.3.2.1 Idealized Trap Potentials

In an ideal lrf ion trap, an ion experiences a dc axial field and an oscillating radial quadrupole field. In an lrf trap, the dc axial potential Φ_z and rf radial potential Φ_{lrf} in cylindrical coordinates (r, θ, z) are given by

$$\begin{aligned}\Phi_z &= -\frac{V_{\text{dc}}}{r_a^2}(r^2 - 2z^2), \\ \Phi_{\text{lrf}} &= \frac{V_{\text{rf}}}{r_r^2}r^2 \cos(2\theta) \cos(2\omega t),\end{aligned}\tag{6.2}$$

where V_{dc} and V_{rf} are the dc and the rf voltages applied, r_a and r_r are effective axial and radial dimensions of the end caps and radial electrodes, and $2\omega = \omega_{\text{rf}}$ is the rf frequency [92]. By contrast, in an rrf trap, the rf radial potential rotates with time as [92]

$$\Phi_{\text{rrf}} = \frac{V_{\text{rf}}}{r_r^2}r^2 \cos[2(\theta + \omega t)].\tag{6.3}$$

There are two dimensionless Mathieu parameters in rf traps: q and d are defined as

$$\begin{aligned}q &= \frac{4q_0 V_{\text{rf}}}{m_0 \omega_{\text{rf}}^2 r_r^2}, \\ d &= \frac{4q_0 V_{\text{dc}}}{m_0 \omega_{\text{rf}}^2 r_a^2},\end{aligned}\tag{6.4}$$

where q_0 is the charge of the ion and m_0 is the mass of the ion.² The q parameter describes the strength of the rf electric field while d describes the applied dc quadrupole field.

Ion motion in an idealized lrf trap has been discussed extensively (see, e.g., Refs. [55, 56]). For an rrf ion trap, Hasegawa and Bollinger solved the equations of motion in the rotating, time-varying potential in Ref. [92]. In both lrf and rrf traps in the small- q limit, the transverse motion separates into a fast micromotion at ω_{rf} and a slower secular motion. In an ideal lrf potential of Eq. (6.2), the two transverse secular modes are degenerate, with the eigenfrequency $\omega_0 = (\omega_{\text{rf}}/2)\sqrt{q^2 - d}$. One convenient basis for these degenerate modes is clockwise (CW) and counterclockwise (CCW) circular motion. In practice, however, static transverse stray fields (often well approximated by

² Earlier in the thesis, a was used to describe this dimensionless parameter d . Here, I keep it as a , to be consistent with the original proposal in Ref. 92

an $x^2 - y^2$ quadrupole term) lift this degeneracy and select two orthogonal linear normal modes, typically aligned close to the x - or y -axis, resulting in two distinct secular frequencies $\omega_x \neq \omega_y$. In this case, the mode axes are fixed in the laboratory frame, and generic excitations lead to beating between the two linear modes. In the rrf trap, the degeneracy is lifted intrinsically by the rotating rf drive. For a finite value of q , the two transverse secular eigenmodes are split even in the absence of static transverse fields. To leading order in q , the CW and CCW eigenfrequencies are [92]

$$\omega_{\pm} \simeq \frac{\omega_{\text{rf}}}{2} \left(\sqrt{q^2 - d} \pm \frac{q^2}{2} \right). \quad (6.5)$$

Because the rrf eigenmodes remain CW and CCW, a linearly polarized excitation along the x - or y -axis can be viewed as an equal-amplitude superposition of these eigenmodes. The resulting transverse motion stays linear, oscillating at $\bar{\omega}_{\text{sec}} = (\omega_+ + \omega_-)/2$, but its axis rotates slowly in the xy -plane at the precession rate $\omega_{\text{prec}} = (\omega_+ - \omega_-)/2$. Exciting instead a pure CW or CCW mode produces a circular secular trajectory for which no oscillation axis is defined, and thus there is no meaningful notion of “precession.”

6.3.2.2 The JILA Eight-Rod Trap

The schematic of our ion trap is shown in Fig. 2.10. Unlike standard lrf traps, which typically use four electrodes to generate a quadrupole field, our trap utilizes an eight-rod geometry. This design is required to generate the highly uniform, rotating polarizing bias electric field necessary for our precision measurement. Each rod pair, comprised of diametrically opposed rods, is driven separately giving us full control of both amplitude and phase between rod pairs. For the rf radial confinement, the two rods in a rod pair are in phase, whereas for the polarizing field, the two rods are out of phase by π . The typical value of the rf frequency is $\omega_{\text{rf}}/2\pi = 50$ kHz.

6.3.2.3 Ion Dynamics in the rrf Trap

In the harmonic regime, ions in the rrf trap undergo secular motion with the two transverse eigenfrequencies given in Eq. (6.5). In our experiment, ions are excited with a relatively weak




Figure 6.5: Transverse mode frequencies in the rrf trap as a function of Mathieu parameter q . Blue circles show the measured mean secular frequency $\bar{\omega}_{\text{sec}}$ (top) and precession frequency ω_{prec} (bottom), extracted from the transverse motion following a weak kick along x -axis. Error bars are smaller than the marker size for all points and are not visible at this scale. Solid curves are the predictions of Eq. (6.5) evaluated using the calculated trap parameters q and d with no fit parameters. Over the range shown, ω_{prec} exhibits a quadratic dependence on q (up to higher-order corrections).

transverse pulse along the x -axis (or y -axis), which prepares a linear superposition of the two circular eigenmodes with the resulting motion well within the harmonic region of the trapping potentials. The resulting transverse motion of the ions can be expressed as an oscillation at the mean secular frequency $\bar{\omega}_{\text{sec}} = (\omega_+ + \omega_-)/2$, with a slow precession at the difference frequency $\omega_{\text{prec}} = (\omega_+ - \omega_-)/2$, as previously discussed. Figure 6.5 shows that both $\bar{\omega}_{\text{sec}}$ and ω_{prec} vary with q in quantitative agreement with Eq. (6.5). Equivalently, the transverse motion of ions in the rrf trap can be interpreted as that in an lrf trap in the presence of an effective axial magnetic field [92]. For a precession frequency of 200 Hz, the effective magnetic field is about 60 G for our ThF^+ ions.

To distinguish transverse dynamics in the lrf and rrf traps, we excite the ions with a weak kick at an azimuthal angle α , where $\alpha = 0^\circ$ corresponds to a kick along the x axis and $\alpha = 90^\circ$ along the y axis. For each value of α , we record the two detector coordinates $x_{\text{det}}(t)$ and $y_{\text{det}}(t)$ and fit each trace to a sum of two sinusoids at the transverse secular frequencies: in the lrf trap, the fit frequencies are ω_x and ω_y and in the rrf trap, ω_+ and ω_- . We extract the amplitudes for the two frequency components $a_1(\alpha)$ and $a_2(\alpha)$. To quantify the extent to which both frequencies are present in a given signal, we define the mixing fraction F

$$F(\alpha) = \frac{\min[a_1(\alpha), a_2(\alpha)]}{a_1(\alpha) + a_2(\alpha)}, \quad (6.6)$$

where $F \simeq 0$ corresponds to a single frequency motion and $F \simeq 0.5$ corresponds to comparable contributions from both frequency components.

In the lrf trap, the transverse eigenmodes are linear oscillations along fixed principal axes. Static transverse imperfections will lift the degeneracy and produce two secular frequencies $\omega_x \neq \omega_y$, as previously mentioned. A general transverse excitation therefore looks like $x(t) = a_x \cos \omega_x t + \phi_x$ and $y(t) = a_y \cos \omega_y t + \phi_y$. When the detector axes are aligned with the principal axes such that $\theta_{\text{det}} \simeq 0$, each measured coordinate contains only a single secular frequency; $x_{\text{det}}(t)$ will be dominated by ω_x and $y_{\text{det}}(t)$ by ω_y . Transverse kicks that are not along the principal axes change

the relative weights a_x and a_y and the relative phase

$$\Delta\phi(t) = (\omega_x - \omega_y)t + (\phi_x - \phi_y),$$

which cause the ion trajectory in the transverse plane to evolve between linear and circular modes, but does not introduce a second frequency component to each transverse coordinate. In contrast, in the rrf trap, the degeneracy is intrinsically lifted by the rf drive, and the two transverse circular eigenmodes (CW and CCW) oscillate at $\omega_{\pm}$. A linear kick necessarily excites a superposition of the two circular modes, so each Cartesian coordinate always contains both $\omega_{\pm}$, regardless of the azimuthal angle of the kick α . If the detector axes are rotated by a fixed angle θ_{det} relative to the trap axes, then it is possible for each Cartesian coordinate to contain *both* frequencies even in the case of the lrf trap.

$$\begin{aligned} x_{\text{det}}(t) &= x(t) \cos \theta_{\text{det}} + y(t) \sin \theta_{\text{det}} \\ y_{\text{det}}(t) &= -x(t) \sin \theta_{\text{det}} + y(t) \cos \theta_{\text{det}}. \end{aligned}$$

However, the presence of both frequencies is a strong function of α , because both linear modes have to be excited. In the rrf trap, the presence of two frequencies in each coordinate is intrinsic and remain near maximal over the full range of α . This behavior is shown in Fig. 6.6. In this figure, same data are analyzed differently: for panels (a), the ion positions in x and y are used as they are, i.e. with detector axes present in the experiment. For panels (b), the the ion positions x and y are rotated by an additional angle $\theta_{\text{det}} \simeq 45^\circ$. Panels (i) show the lrf data and panels (ii) show the rrf data. We observe that for the lrf trap, the maximum F_x occurs at $\alpha = 0^\circ$ and maximum F_y at $\alpha = 90^\circ$, indicating each detector coordinate is predominantly made of a single frequency component. However, when the detector axes are rotated by 45° , we see a strong variation of both F_x and F_y with α . However, for the rrf trap, F_x and F_y are constant at 0.5 across all α and θ_{det} .

We can also directly look at the ion dynamics at the detector. Figure 6.7 shows the transverse ion trajectory in the rrf trap. As before, the ions are given a weak transverse pulse along the x -axis. The plotted trajectory begins approximately 4 ms after the transverse kick, and the color darkens




Figure 6.6: The mixing fraction $F(\alpha)$ extracted from the transverse detector signals as a function of azimuthal kick angle α . For each value of α , we fit $x_{\text{det}}(t)$ and $y_{\text{det}}(t)$ to a sum of two sinusoids at the transverse secular frequencies and define F as in Eq 6.6. Panels (a) use detector axes present in the experiment ($\theta_{\text{det}} \simeq 0^\circ$), and panels (b) use detector axes rotated by an additional angle ($\theta_{\text{det}} \simeq 45^\circ$). Panels (i) show the lrf data, fit to frequencies $\omega_x \simeq 2\pi \cdot 1700$ Hz and $\omega_y \simeq 2\pi \cdot 1900$ Hz. Panels (ii) show the rrf data with frequencies $\omega_+ \simeq 2\pi \cdot 1700$ Hz and $\omega_- \simeq 2\pi \cdot 1900$ Hz. In the lrf trap with $\theta_{\text{det}} = 0^\circ$, each detector coordinate is predominantly has a single frequency component, indicated by $F_x \simeq 0$ at $\alpha = 0^\circ$ and $F_y \simeq 0$ at $\alpha = 90^\circ$. With $\theta_{\text{det}} = 45^\circ$, both lrf eigenmodes contribute to each detector coordinate, and $F_x(\alpha)$ and $F_y(\alpha)$ increase and vary strongly with α . On the other hand, the rrf data show $F_x \simeq F_y \simeq 0.5$ across all α and θ_{det} .

with increasing time between the kick and detection. For improved visualization, the data are linearly interpolated between successive points and smoothed with a 2-point moving average. The precession of the transverse motion is clearly visible. The precession frequency is expected to be $\omega_{\text{prec}}/2\pi \approx 120$ Hz with the secular frequency $\bar{\omega}_{\text{sec}}/2\pi \approx 2000$ Hz. Just before ejecting the ions for imaging, we quasi-adiabatically ramp down the rf amplitude, which suppresses the micromotion to the point where it is not visible in Fig. 6.7. The ~ 17 -petal-flower pattern in Fig. 6.7 is what one would expect from the beating between simultaneously excited clockwise and counterclockwise eigenmodes with frequencies that differ by about 12%, and quite different from the alternating-X-and-O pattern one observes in beating between $\hat{x}$ and $\hat{y}$ eigenmodes.

6.3.3 Implications for Precision Measurement in Non-ideal Geometries

6.3.3.1 Pseudopotential in Non-Ideal Trap Geometry

Our ion trap, as shown in Fig. 2.10, was carefully designed to provide a rotating polarizing bias electric field required for the eEDM measurement with 0.01% uniformity. However, it was **not** optimized to give perfectly harmonic rf confinement. To operate this 8-rod structure in a traditional lrf quadrupole configuration, the rods are driven with a phase pattern of $\{0, 0, \pi, \pi, 0, 0, \pi, \pi\}$. Because adjacent rods must be grouped in phase to approximate a four-pole geometry, this inherently produces a trap which is cylindrically symmetric for only small displacements. The anharmonic terms are highly anisotropic, characterized by a weak axis between adjacent electrodes held at the same instantaneous potential. When the trap confinement is loosened to minimize ion-ion interactions during precision spectroscopy, our ion cloud expands and samples regions where the fields deviate from the perfectly harmonic, idealized rf potentials discussed in Section 6.3.2. Instead of using the idealized rf potential, one can use a pseudopotential approximation to compute an effective potential. This approximation is widely used in rf trapping and is valid when the drive frequency is large compared to the secular frequencies and the micromotion amplitude is small on the spatial scale over which the field varies [55, 56, 93]. A convenient way of expressing the rf drive




Figure 6.7: Transverse ion trajectory in the xy -plane, starting approximately 4 ms after a weak transverse kick along the x -axis. The color darkens with increasing time delay between the kick and detection. For visualization, the data are linearly interpolated between successive points and smoothed with a moving average of two points. The trajectory shows a clear precession motion, a distinctive feature of ion dynamics in the rrf trap. The ion optics used to focus ejected ions onto an imaging detector have magnification 1.5 times larger in the $\hat{x}$ than in the $\hat{y}$ direction. We have correspondingly stretched the x -axis of the plot so that the image on the page is more faithful representation of the trajectory in the xy -plane.

is

$$\mathbf{E}(\mathbf{r}, t) = \Re[\mathbf{E}_0(\mathbf{r})e^{i\omega_{\text{rf}}t}],$$

where $\mathbf{E}_0(\mathbf{r})$ can be a complex spatial amplitude. In the pseudopotential approximation, the time-averaged micromotion leads to a time-independent effective potential that is proportional to the square of the rf field amplitude [54, 55, 93]

$$\begin{aligned} U_{\text{pseudo}}(\mathbf{r}) &= \frac{q_0^2}{4m_0\omega_{\text{rf}}^2} \mathbf{E}(\mathbf{r}) \cdot \mathbf{E}_0(\mathbf{r})^* = \frac{q_0^2}{4m_0\omega_{\text{rf}}^2} |\mathbf{E}_0(\mathbf{r})|^2 \\ &= \frac{q_0^2}{4m_0\omega_{\text{rf}}^2} |\nabla\Phi_0(\mathbf{r})|^2, \end{aligned} \quad (6.7)$$

where $\mathbf{E}_0 = -\nabla\Phi_0$ and $\Phi(\mathbf{r}, t) = \Re[\Phi_0(\mathbf{r})e^{i\omega_{\text{rf}}t}]$. With the rf drive $\mathbf{E}(\mathbf{r}, t)$ written in a complex form, it is clear that the pseudopotential approximation applies not only to the lrf potential, but to the rrf potential where the difference between them is the rrf potential having a complex amplitude $\Phi_0(\mathbf{r})$ with azimuthally varying phase. More specifically, the rrf potential of Eq. (6.3) can be written as

$$\Phi_{\text{rrf}}(r, \theta, t) \propto r^2 \cos(2\theta + \omega_{\text{rf}}t) = \Re[r^2 e^{i2\theta} e^{i\omega_{\text{rf}}t}].$$

With $\Phi_0(r, \theta) \propto r^2 e^{i2\theta}$, the corresponding pseudopotential is

$$U_{\text{pseudo}}(\mathbf{r}) \propto |\nabla\Phi_0|^2 = (|\nabla\Phi_c|^2 + |\nabla\Phi_s|^2), \quad (6.8)$$

where $\Phi_c \propto r^2 \cos(2\theta)$ and $\Phi_s \propto r^2 \sin(2\theta)$. This decomposition provides a useful way to interpret the rrf pseudopotential discussed below.

Using the pseudopotential approximation, the relevant measure of confinement and effective trap depth is the quantity $|\nabla\Phi_0|^2$. We now evaluate this quantity for our eight-rod geometry used in the experiment. Experimentally, we drive our eight radial electrodes with fixed phase offsets. For the lrf configuration, the electrode phases are $\{0, 0, \pi, \pi, 0, 0, \pi, \pi\}$, while for the rrf configuration they are $\{0, \pi/2, \pi, 3\pi/2, 0, \pi/2, \pi, 3\pi/2\}$, as shown in Fig. 6.8a and 6.8d, respectively. Figure 6.8b shows the transverse profile of $|\nabla\Phi_{\text{lrf}}|^2$, a quantity to which the pseudopotential is proportional, for the lrf configuration. The effective trap depth is not azimuthally uniform. This anisotropy is seen more directly in Fig. 6.8c, which shows $|\nabla\Phi_0|^2$ traces along the three cuts (i),




Figure 6.8: Transverse field curvature, proportional to the pseudopotential $U_{\text{pseudo}} \propto |\nabla\Phi|^2$, of the lrf and rrf trap configurations. (a) Electrode phase patterns and the three geometric cuts, labeled (i)–(iii), for the lrf configuration. (b) Surface plot of $|\nabla\Phi_{\text{lrf}}|^2$ in the transverse plane. (c) Corresponding one-dimensional cuts along (i)–(iii). The grey curve shows a quadratic reference fit near the trap center. (d) Electrode phase pattern and the three geometric cuts, labeled (iv)–(vi), for the rrf configuration. (e) and (f) Surface plots of $|\nabla\Phi_c|^2$ and $|\nabla\Phi_s|^2$, respectively. (g) Average of the two quadratures as the total rrf field curvature $|\nabla\Phi_{\text{rrf}}|^2 = (|\nabla\Phi_c|^2 + |\nabla\Phi_s|^2)/2$. (h) Corresponding one-dimensional cuts along (iv)–(vi). Note that the shallowest cut in this trap is deeper than that in the lrf trap.

(ii), and (iii), corresponding to angles of 0° , 22.5° , and 45° with respect to the x -axis. The effective trap depth is shallowest along cut (i), which passes between the two electrodes held at the same instantaneous potential. Cut (ii) passes through the physical rods, spanning from $r = \pm 5.215$ cm to $r = \pm 6.485$ cm, while cut (iii) lies midway between electrodes of opposite phase.

For the rrf configuration, as discussed above, the total pseudopotential can be decomposed into the cosine and sine quadratures, shown in Fig. 6.8e and Fig. 6.8f, respectively. These two quadratures are related in space by

$$\Phi_c(r, \theta) = \Phi_s(r, \theta + \pi/4).$$

It is therefore tempting to think of the rrf pseudopotential as a simple azimuthal average of a single quadrature. However, the relevant quantity to the total pseudopotential is Eq. (6.8), the average between the two quadratures, and is shown in Fig. 6.8g. Figure 6.8h shows the total rrf $|\nabla\Phi_{\text{rrf}}|^2$ traces along the cuts. Because of the averaging over the two quadratures, the cuts at $\pm 22.5^\circ$ are identical, leaving only two unique traces. These traces are substantially more uniform than those of the lrf configuration in Fig. 6.8c. More importantly, the shallowest direction in the rrf trap is deeper overall than in the lrf trap.

6.3.3.2 Enhancement of Effective Trap Depth

To experimentally probe the azimuthal dependence of the trap depth predicted in Fig. 6.8, we adiabatically apply a static dipole field that adds an additional potential

$$U_{\text{dc}}(\mathbf{r}) = -q_0 \mathbf{E}_{\text{dc}} \cdot \mathbf{r},$$

which tilts the potential landscape. As $|\mathbf{E}_{\text{dc}}|$ increases, ions are displaced from the trap center and eventually escape, so by looking at how many ions survived as a function of tilt direction, we measure the direction dependence of the effective trap depth. Here, we express the strength of the applied tilting field by the magnitude of the offset in the trap equilibrium location if the potential were fully harmonic, $x_{\text{eq}} \propto q_0 E_{\text{dc}} / (m\bar{\omega}_{\text{sec}}^2)$.




Figure 6.9: Ion survival as a function of adiabatic dc tilt for different transverse cuts. The ion signal is recorded after the trap has been adiabatically tilted by a static dipole field. In both panels, cuts (i)-(vi) correspond to the transverse directions defined in Fig. 6.8 (with 0° along the x -axis). The horizontal axis is the applied tilt expressed as an equivalent equilibrium shift $x_{eq} = q_0 E_{dc} / (m \bar{\omega}_{sec}^2)$, defined using the harmonic approximation. The dotted vertical lines are guides to the eye. Top: cuts (i), (ii), and (iii) of the lrf configuration correspond to azimuthal angles 0° , 22.5° , and 45° , respectively. Ion loss occurs at smaller $|x_{eq}|$ for cut (i), consistent with this direction having the shallowest effective confinement. Bottom: cuts (iv), (v), and (vi) of the rrf configuration correspond to 0° , 22.5° , and -22.5° , respectively. The loss curves occur at more similar $|x_{eq}|$ across angles, indicating a more azimuthally uniform effective trap depth in the rrf configuration.

In Fig. 6.9, in the lrf configuration, ion loss occurs at different x_{eq} for different azimuthal angles (0° , 22.5° , and 45° corresponding to cuts (i)–(iii) in the left column of Fig. 6.8), indicating a strong angular dependence of the effective trap depth. In particular, ions are lost at smaller tilts along the weakest direction, cut (i), consistent with the reduced trap depth for that cut predicted in Fig. 6.8.

By contrast, the rrf configuration effectively reduces this asymmetry down to just two distinct axes. To explicitly test this, we measured the rrf ion survival at azimuthal angles of 0° and $\pm 22.5^\circ$ (corresponding to cuts (iv)–(vi) in the right column of Fig. 6.8). Because the quadrature averaging maps the $+22.5^\circ$ and -22.5° directions onto the exact same averaged pseudopotential profile, the experimental loss curves for these two angles are very similar. The loss occurs at similar x_{eq} across all measured angles, visually highlighting how the rrf configuration symmetrizes the trap depth and verifies the theoretical reduction to only two distinct effective axes. Crucially, not only does ion loss occur at more uniform x_{eq} values in the rrf configuration, but the absolute minimum x_{eq} required for ions to escape is noticeably larger than that of the weakest axis in the lrf configuration.

This increased minimum escape barrier is significant because the overall effective trap depth is determined by the shallowest available axis. In our ion cloud, Coulomb interactions redistribute energy and couple their translational degrees of freedom, leading to ions exploring a wide range of trajectories over time. As a result, long-time trap depth is governed by the lowest available escape barrier of the effective potential, not by a single trajectory. The anisotropy of the lrf configuration therefore introduces “easy escape” directions that reduce this minimum barrier and yield a comparatively shallow effective trap depth. By contrast, the rrf configuration raises this minimum barrier through quadrature averaging, resulting in a deeper overall trap depth.

6.3.4 Discussion

Our ultimate goal is to maximize sensitivity to the eEDM measurement, which is proportional to $\sqrt{N}\tau$ as mentioned. The coherence time is determined largely by ion-ion collisions which depend in turn on the number of ions, their temperature, and the nature of the confining potential. We note

that for holding ions in a loose yet deep potential well, there are options other than the rrf trap, including driving our eight rods in hexadecapolar pattern, for a potential more nearly approaching a square well. We are in the process of investigating which approach is most promising for maximum eEDM sensitivity. To date all our efforts to understand systematic error in our measurement of the eEDM has been in the context of the lrf ion traps [51]. In Section 6.3.2.3, we describe the secular motion of an ion in an rrf trap being equivalent to the secular motion of an ion in an lrf trap with the addition of an effective axial magnetic field of some tens of Gauss, a value at least 1000 times larger than the physical magnetic fields present at the center of our trap during our precision measurement efforts. Of course the fictitious field couples to the center-of-mass motion of an ion and does not directly couple to the spins of the electrons within the ion. The direction of the rotation of the rrf potential and the resulting effective magnetic field do however change sign with the time reversal operation; thus we intend to explore the implications very thoroughly, given that our overarching science goal is to search for evidence for a tiny fundamental violation of time reversal symmetry.

Chapter 7

Towards Long coherence time: Loss of coherence

In our system, because the ions are in the rotating frame, the Ramsey precession frequency is not set by a static laboratory-frame magnetic bias alone. Rather, it depends on both the magnetic field structure and the rotating electric field.

Long Ramsey coherence therefore requires both temporal stability and spatial homogeneity of the magnetic and rotating electric fields. Temporal fluctuations lead to phase noise within a shot or from shot to shot. Spatial inhomogeneities cause different ions to precess at different rates, leading to dephasing across the ensemble. In this chapter, I focus specifically on how these effects limit coherence, leaving related systematic shifts to later discussion.

As a simple scale, if we require the ions to accumulate less than 0.1 cycle of additional phase at the end of a 20 second interrogation time, this corresponds to an allowed detuning error of only 5 mHz, for the Ramsey frequency of about 50 Hz. Fractionally speaking, this is at the 10^{-4} level.¹

¹ It is worth noting that this simple phase-slip estimate is distinct from the projection-noise-limited frequency sensitivity in a Ramsey measurement. On the side of the Ramsey fringe, the frequency sensitivity is approximately

$$\delta\omega \simeq \frac{1}{C\tau\sqrt{N}}.$$

Assuming we have $C \approx 1$, $\tau \sim 16$ s, and $N \sim 100$ ions, this comes out to be ~ 1 mHz. Therefore, the shot-to-shot fluctuation would have to be stable at roughly ~ 1 mHz level to avoid contributing appreciable to the statistical uncertainty. However, in our measurement, simultaneous detection of both Stark doublets greatly relaxes this requirement on shot-to-shot stability because common-mode fluctuations are largely rejected.

7.1 Temporal inhomogeneity

A long Ramsey coherence time requires the transition frequency to remain stable over the duration of a single interrogation time and reproducible from shot to shot. In our experiment, the Ramsey frequency is set by the Zeeman splitting in the rotating-field environment and therefore depends on both the magnetic field and the rotating electric field. Therefore, temporal fluctuations in either field reduce the observed coherence.

There are two distinct relevant effects. First, there can be slow drifts during an individual Ramsey sequence. These produce a time-dependent detuning and ultimately an additional phase error. Second, the fields can vary shot to shot. In this case, the different shots evolve at slightly different precession frequencies, so that combining many shots into a single Ramsey fringe can reduce the apparent contrast, even if each shot remains coherence on its own.

7.1.1 Temporal drift of the rotating electric field

The rotating electric field sets the micromotion radius r_{rot} . In the presence of an axial magnetic field gradient B_{axgrad} , this micromotion gives rise to the Ramsey precession frequency f_{Ramsey} . A temporal drift of the rotating field changes r_{rot} , which changes the magnetic field gradient sampled by the ions and therefore shifts f_{Ramsey} . Both slow drifts and shot-to-shot fluctuations are relevant here.

A useful diagnostic for this situation is a Lissajous plot of the populations of the two states $|\uparrow\rangle$ and $|\downarrow\rangle$. If the two populations remain phase correlated, the points should trace an approximately elliptical shape. This diagnostic is valuable for identifying residual coherence at late times, but it does not by itself provide a precise eEDM measurement.

We know there exists shot-to-shot fluctuations in the rotating electric field in the current system. This is mostly from the power op-amps used to generate the rotating field are turned on and off between measurements.

Another useful diagnostic tool is to measure the rotating-field magnitude during each shot

and use it to “correct” the effective Ramsey evolution time. For the Ramsey frequency that scales with the rotating field as $f_{\text{Ramsey}} \propto \mathcal{E}_{\text{rot}}$, then shot i accumulates the phase

$$\phi_i = 2\pi f_i t_i.$$

We can then define a corrected time t_i^{corr} relative to a reference frequency $\bar{f}$ by requiring that it corresponds to the same accumulated phase,

$$\phi_i = 2\pi \bar{f} t_i^{\text{corr}}.$$

This gives

$$t_i^{\text{corr}} = t_i \frac{f_i}{\bar{f}} \approx t_i \frac{\mathcal{E}_{\text{rot},i}}{\mathcal{E}_{\text{rot}}}.$$

In other words, in the data fitting process, shots with a slightly larger $\mathcal{E}_{\text{rot}}$ are assigned a proportionally larger effective Ramsey time, and vice versa. If fluctuations in $\mathcal{E}_{\text{rot}}$ are important source of shot-to-shot phase scatter, then this procedure should partially remove that scatter and recover fringe contrast.

This idea is illustrated in Fig. 7.1. Consider two sine waves with the same initial phase but slightly different frequencies f_1 and f_2 , where f_1 is the reference frequency. At the nominal Ramsey time indicated by the vertical dashed line, the two curves have different phases because they have evolved at slightly different rates. However, if the point from the f_2 curve is replotted at the correct time

$$t^{\text{corr}} = t \frac{f_2}{f_1},$$

then it is mapped to the time at which the reference-frequency oscillation would have accumulated that same phase. In this way, the correction removes the apparent phase shift associated with the frequency difference.

Figure 7.2 shows the Ramsey fringes before and after this correction. During this measurement, we monitored $\mathcal{E}_{\text{rot}}$ and saw $\sim 0.05\%$ fluctuations between each shot. For the corrected data, each point is replotted using the rescaled time t_i^{corr} . In the dataset shown here, the correction does not increase the fitted contrast. This suggests that shot-to-shot fluctuations in the magnitude




Figure 7.1: Illustration of the Ramsey-time correction. The orange curve shows represents a shot with precession frequency f_i , while the blue curve represents the reference oscillation frequency at $\tilde{f}$. A measurement taken at the nominal Ramsey time t_i is replotted at the corrected time $t_i^{\text{corr}} = t_i f_i / \tilde{f}$, so that it corresponds to the same accumulated phase on the reference oscillation.

of $\mathcal{E}_{\text{rot}}$ were not the dominant source of decoherence in this measurement. Even so, this procedure remains a useful diagnostic for determining whether the amplitude noise in the rotating field contributes to the observed dephasing.

Because we know that one dominant source of temporal fluctuation in $\mathcal{E}_{\text{rot}}$ is thermal drift associated with switching on and off the high-voltage electronics, the Gen. III apparatus will operate with both the rotating electric and rf field continuously on. Keeping on these electronics continuously should result in a steady operating temperature, and therefore should reduce this source of drift.

7.1.2 Temporal drift of the magnetic field

As discussed earlier, in this rotating-field configuration, the Ramsey precession is especially sensitive to the axial magnetic field gradient sampled by the ion micromotion. In the prototype apparatus, the dominant temporal magnetic noise relevant here is slow drift of the magnetic field gradient of the environment, rather than shot-to-shot fluctuations of a static bias magnetic gradient field.²

For the present operating conditions in the prototype apparatus with $f_{\text{rot}} = 147.5$ kHz and $\mathcal{E}_{\text{rot}} = 60$ V/cm, an axial magnetic field gradient of 3 mG/cm produces approximately 50 Hz of Ramsey precession. Equivalently, a gradient change of 1 $\mu\text{G}/\text{cm}$ shifts the Ramsey frequency by about 17 mHz. For Ramsey interrogation times of order 20 s, sub $\mu\text{G}/\text{cm}$ -level drifts can lead to substantial accumulated phase.

Unfortunately, the magnetic environment around the experiment is not static. For example, the door to the laboratory is made of a large piece of steel, and a freight elevator operates immediately adjacent to the lab. We monitored the spatial magnetic gradient near the prototype apparatus. The measurements were taken about 1.5 m away from the Science region, but in the

² Strictly, shot-to-shot fluctuations and slow drifts are not fundamentally distinct; they refer to magnetic field variations on different timescales relative to the experimental cycle. The distinction I want to make here is the following: the rotating electric field is generated by power op-amps that are switched on and off between experimental shots, so it can have true shot-to-shot reproducibility errors. By contrast, the bias magnetic field gradient is produced by DC currents that remain on continuously. Therefore, its variations are expected to appear mainly as slow drifts or low-frequency noise from the environment around the trap, rather than as discrete turn-on transients from one shot to the next.




Figure 7.2: Ramsey fringes for doublets A and B before (top) and after (bottom) correcting the effective Ramsey time of each shot based on the measured rotating-field magnitude. For this dataset, $\mathcal{E}_{\text{rot}}$ fluctuated $\sim 0.05\%$ between each shot. For the uncorrected data, the green markers show the mean asymmetry of many shots taken at the same nominal Ramsey time, with the standard error. The solid curves in the top panels and dashed curves in the bottom panels are sinusoidal fits. In this dataset, applying the time correction does not produce a clear improvement in the recovered fringe contrast.

direction that is closer to the door and the elevator. These measurements show that motion of the door and the elevator can change the magnetic field gradient by as much as $40 \mu\text{G}/\text{cm}$ at the place of measurement, far above the level desirable for long-coherence measurements.

Restricting use of the door or the elevator during data taking is not a practical long-term solution.³ Instead, we are developing shim coils to monitor and actively compensate slow changes in the ambient magnetic field gradient near the chamber.⁴ Such feedforward stabilization is planned for the Gen. III apparatus. For the measurements present in this thesis however, no such shim system was available, so the longest coherence measurements were taken during relatively quiet periods, such as nights and weekends.

7.2 Spatial inhomogeneity of the rotating electric field

In the previous generations, it was assumed that the rotating electric field was spatially uniform across the cloud. However, in practice, this need not be true. Spatial inhomogeneity of the rotating electric field can contribute to decoherence.

A spatially varying rotating electric field produces a spatially varying micromotion orbit. In other words, the rotating-field radius can depend on position

$$r_{\text{rot}}(\mathbf{r}) = \bar{r}_{\text{rot}} + \delta r_{\text{rot}}(\mathbf{r}), \quad (7.1)$$

where $\bar{r}_{\text{rot}}$ is the mean orbit radius and $\delta r_{\text{rot}}(\mathbf{r})$ is the residual spatial variation across the cloud. The spatial variation in the micromotion also implies a spatial variation in the ponderomotive energy.

If the rotating field inhomogeneity becomes too large, some ions acquire enough ponderomotive energy from the rotating field, as discussed in Sec. 5.4.1, resulting in those ions leaving the trap. In this sense, the trap itself suppresses the most extreme regions of the rotating field inhomogeneity:

³ We want to keep good and friendly relationships with the rest of JILA.

⁴ Because the effective g -factor of our ions is still small compared to that of a bare electron, and because our measurement is primarily sensitive to magnetic field gradients rather than the absolute field magnitude, we have chosen not to use magnetic shielding. It would also add substantial complexity to an already complicated system. In addition, unless designed with great care, the magnetic shielding can actually make the spatial gradient worse.

ions that experience the largest distortions are less likely to remain trapped and contribute to the measurement. In other words, the surviving ensemble does not sample arbitrarily large rotating field inhomogeneity. However, this still does not imply that the residual inhomogeneity within the trapped ensemble is negligible; moderate spatial variation in r_{rot} can remain and still matter in precision spectroscopy.

This effect is especially important in situations where the effective rotating-frame bias field is generated by the laboratory-frame B_{20} gradient. In this case, we have

$$f_{\text{Ramsey}} \propto B_{20} r_{\text{rot}}(\mathbf{r}),$$

so any spatial variation in the rotating-field radius maps directly onto a spatially varying effective Zeeman shift, and thus the Ramsey frequency. Therefore, even in the absence of higher-order magnetic curvature (which will be discussed in the next section), inhomogeneity of the rotating electric field alone can produce inhomogeneous broadening through the position dependence of r_{rot} .

It is useful to distinguish this mechanism from dephasing caused by spatial inhomogeneity of the magnetic field itself. In the presence of the rotating electric field inhomogeneity, the position dependence enters because the ions sample different rotating-field radii. On the other hand, in the presence of the magnetic field inhomogeneity, the position dependence enters because the magnetic field itself is spatially nonuniform. In the rotating frame, these two effects become coupled, but it is still conceptually useful to introduce them separately.

7.3 Spatial inhomogeneity of the magnetic field

I now turn to spatial inhomogeneity of the magnetic field. To isolate the magnetic contribution, I first assume that the rotating-field radius is uniform across the cloud, so that the position dependence of the bias field arises entirely from the magnetic field itself.

Spatial inhomogeneity of the magnetic field causes ions at different locations in the trap to acquire different Ramsey phases. However, not every magnetic field gradient contributes in the same way. The ions are confined in an approximately harmonic potential with oscillatory motion

about the trap center. As a result, magnetic field gradients that vary linearly with position do not by themselves produce ensemble dephasing for the cloud. The leading dephasing terms therefore contain higher-order spatial variation of the magnetic field.

7.3.1 Magnetic multipoles

As discussed in Chapter 2, the molecular quantization follows the rotating electric field and therefore follows the rotating micromotion trajectory of the ion, $\mathbf{r}_{\text{rot}}$. We are interested not in the full laboratory-frame magnetic field itself, but in the component of that field projected into the instantaneous rotating-field direction and averaged over one cycle of the rotation. For instance, for an ion with coordinate $\mathbf{r} = (x, y, z)$, I define the scalar effective magnetic field experienced by the ion as

$$B^{\text{eff}}(\mathbf{r}) = \langle \mathbf{B}(\mathbf{r} + \mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \rangle_t, \quad (7.2)$$

where $\mathbf{r}_{\text{rot}}(t)$ is the micromotion orbit associated with the rotating electric field and $\hat{\mathbf{r}}_{\text{rot}}$ is the corresponding unit vector along the instantaneous molecular quantization axis. The average $\langle \cdots \rangle_t$ is taken over one rotation period. Here, x , y , and z denote the ion secular coordinates, i.e. coordinates defined on timescales slow compared to the rotation period. In the discussion below, I first assume that r_{rot} is uniform across the cloud, so that the position dependence only enters through the magnetic components.

To describe the laboratory magnetic field, it is convenient to expand it in magnetic multipoles. In the current-free region near the ions, the magnetic field can be written in terms of a scalar magnetic potential,

$$\mathbf{B}(\mathbf{r}) = -\nabla \Phi_B(\mathbf{r}), \quad \nabla^2 \Phi_B(\mathbf{r}) = 0. \quad (7.3)$$

The scalar magnetic potential can be expanded using multipole expansion [94, 95]

$$\Phi_B(\mathbf{r}) = \sum_{l=1}^{\infty} \sum_{m=-l}^l \mathcal{B}_{l,m} r^l Y_{lm}(\hat{\mathbf{r}}), \quad (7.4)$$

where Y_{lm} are the spherical harmonics. The corresponding magnetic field is

$$\mathbf{B}(\mathbf{r}) = -\nabla\Phi_B(\mathbf{r}) = \sum_{l,m} \mathbf{B}_{lm}(\mathbf{r}), \quad (7.5)$$

with

$$\mathbf{B}_{lm}(\mathbf{r}) \equiv -\mathcal{B}_{lm} \nabla \left[r^l Y_{lm}(\hat{\mathbf{r}}) \right]. \quad (7.6)$$

Here, the coefficients $\mathcal{B}_{lm}$ characterizes the strength of the laboratory-frame multipoles, and $\mathbf{B}_{lm}$ denotes the corresponding field component.

It is natural to define effective rotating-frame contribution from *each* multipole as

$$B_{lm}^{\text{eff}}(\mathbf{r}) \equiv \langle \mathbf{B}_{lm}(\mathbf{r} + \mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \rangle_t, \quad (7.7)$$

such that

$$B^{\text{eff}}(\mathbf{r}) = \sum_{l,m} B_{lm}^{\text{eff}}(\mathbf{r}). \quad (7.8)$$

For instance, for the axially symmetric $l = 2, m = 0$ term, using the same convention as in Appendix B, we have

$$r^2 Y_{20}(\hat{\mathbf{r}}) = \sqrt{\frac{5}{16\pi}} (2z^2 - x^2 - y^2),$$

so the corresponding scalar potential is

$$\Phi_{20}(\mathbf{r}) = \mathcal{B}_{20} r^2 Y_{20}(\hat{\mathbf{r}}).$$

Then, the laboratory-frame magnetic field is then

$$\mathbf{B}_{20}(\mathbf{r}) = -\nabla\Phi_{20} \propto \mathcal{B}_{20} (x\hat{\mathbf{x}} + y\hat{\mathbf{y}} - 2z\hat{\mathbf{z}}),$$

up to some overall normalization.

Now evaluating this field on the rotating orbit using $\mathbf{r}_{\text{rot}}(t) = r_{\text{rot}}(\cos(\omega_{\text{rot}}t)\hat{\mathbf{x}} + \sin(\omega_{\text{rot}}t)\hat{\mathbf{y}})$ and Eq. (7.7) gives

$$\begin{aligned} & \mathbf{B}_{20}(\mathbf{r} + \mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \\ & \propto \mathcal{B}_{20} \left[(x + r_{\text{rot}} \cos(\omega_{\text{rot}}t)) \cos(\omega_{\text{rot}}t) + (y + r_{\text{rot}} \sin(\omega_{\text{rot}}t)) \sin(\omega_{\text{rot}}t) \right]. \end{aligned}$$

l	m	$B_{lm}^{\text{eff}}(\mathbf{r})$	Contribution to coherence
2	0	$\propto \mathcal{B}_{20}r_{\text{rot}}$	Common-mode bias field only
3	$\pm 1, 0$	$\propto \mathcal{B}_{3m}r_{\text{rot}}x, \mathcal{B}_{3m}r_{\text{rot}}y, \mathcal{B}_{3m}r_{\text{rot}}z$	Vanishes for a centered cloud
4	$-2, \pm 1$	$\propto \mathcal{B}_{4m}r_{\text{rot}}xy, \mathcal{B}_{4m}r_{\text{rot}}yz, \mathcal{B}_{4m}r_{\text{rot}}xz$	Averages away for nondegenerate secular motion
4	0	$\propto \mathcal{B}_{40}r_{\text{rot}}[r_{\text{rot}}^2 + 2(x^2 + y^2 - 2z^2)]$	Common-mode shift + dephasing
4	2	$\propto \mathcal{B}_{42}r_{\text{rot}}(x^2 - y^2)$	Dephasing from transverse anisotropy

Table 7.1: Leading magnetic multipoles that survive the rotating-frame average. The table is adapted from the multipole expansion presented in Ref. [59], but reorganized here according to the role each term plays in the coherence problem.

Averaging this over one rotation cycle, the terms proportional to x and y vanish, while the term proportional to $\cos^2$ and $\sin^2$ survive:

$$B_{20}^{\text{eff}} = \langle \mathbf{B}_{20}(\mathbf{r} + \mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \rangle_t \propto \mathcal{B}_{20}r_{\text{rot}},$$

which is independent of x , y , and z . Therefore a pure $l = 2, m = 0$ laboratory-frame field $\mathbf{B}_{20}(\mathbf{r})$ does not by itself cause dephasing across a centered cloud; it only produces a common-mode shift.

The same averaging procedure can be applied to the other multipoles. Noah Schlossberger performed this calculation in his thesis; Table 7.1 is adapted from Ref. [59], but reorganized here to emphasize the relevant terms and their role in the coherence problem.⁵

In Table 7.1, I have implicitly assumed that the spatial coordinates of the multipoles coincide with the principal axes of the secular motion in the trap. In other words, x , y , and z are taken to be the normal-mode coordinates of the trap with secular frequencies f_x , f_y , and f_z . This convention becomes especially important when we classify the terms as common-mode, dephasing, or averaging to zero under secular motion. For instance, we can say that xy term would average away to zero for $f_x \neq f_y$ only when the trap normal modes are along the multipole coordinates. If the trap normal modes are rotated with respect to the multipole axes, then the fields would be expressed as a linear combination of the terms listed in the table.

⁵ The mapping between a Cartesian form and a specific m value depends on what convention is used for the spherical harmonics, e.g. real versus complex harmonics, as well as the associated normalization and phase conventions. Numerical prefactors and $\pm m$ labels may therefore differ between conventions, while the physical content remains the same.

The $l = 3$ components that survive the rotating-frame average are linear in secular coordinates,

$$B_{3m}^{\text{eff}}(\mathbf{r}) \propto \mathcal{B}_{3m} r_{\text{rot}} x, \quad \mathcal{B}_{3m} r_{\text{rot}} y, \quad \mathcal{B}_{3m} r_{\text{rot}} z.$$

Since ions undergo harmonic motion in the trap, these terms average to zero over many cycles of secular motion, unless the ions are not sitting at the center of the trap. Under normal operating conditions, the ion cloud is centered near the trap minimum, therefore, these terms with order $l = 3$ do not contribute to dephasing, although they are useful diagnostics of cloud displacement.

The first multipole components that generate genuine spatial variations in the effective transition frequency across a centered ensemble are the $l = 4$ terms. In particular,

$$B_{40}^{\text{eff}}(\mathbf{r}) \propto \mathcal{B}_{40} r_{\text{rot}} [r_{\text{rot}}^2 + 2(x^2 + y^2 - 2z^2)],$$

$$B_{42}^{\text{eff}}(\mathbf{r}) \propto \mathcal{B}_{42} r_{\text{rot}} (x^2 - y^2).$$

These terms are quadratic in the secular coordinates and therefore do not average out over many secular cycles. Therefore, they represent the leading spatial magnetic inhomogeneities relevant for ensemble dephasing in a centered cloud.

The remaining $l = 4$ contributions with cross terms such as xy , xz , and yz are more subtle. For nondegenerate secular frequencies ($f_x \neq f_y \neq f_z$), they average to zero. However, they can become relevant if the two secular frequencies are equal or nearly commensurate or if the magnetic field multipole axes are not aligned with the trap normal mode axes.

Among the $l = 4$ terms, the $\mathcal{B}_{40}$ term is especially instructive since it contains both a common-mode shift and a dephasing contribution. The corresponding frequency shift due to the $\mathcal{B}_{40}$ term has the form

$$\delta f_{40}(\mathbf{r}) \propto B_{40}^{\text{eff}}(\mathbf{r}) \propto \mathcal{B}_{40} r_{\text{rot}} [r_{\text{rot}}^2 + 2(x^2 + y^2 - 2z^2)]. \quad (7.9)$$

If all ions experience the same rotating-orbit radius r_{rot} , then Eq. (7.9) naturally separates into two pieces,

$$\delta f_{40}(\mathbf{r}) \propto \mathcal{B}_{40} r_{\text{rot}}^3 + 2\mathcal{B}_{40} r_{\text{rot}} (x^2 + y^2 - 2z^2).$$

The first term, proportional to r_{rot}^3 is common mode; this term only contributes to the overall frequency shift to the whole ensemble, but not to the dephasing within the ensemble. By contrast, the second term varies across the extent of the ion cloud and therefore generates inhomogeneous dephasing.

This example illustrates an important point that a magnetic multipole can survive the rotating-frame average without necessarily contributing to decoherence. A term contributes to dephasing only when the surviving term is position-dependent and is even in x , y , and z .

This same example also explicitly shows how electric and magnetic field homogeneity become coupled in the rotating frame. For instance, if the rotating-orbit radius is not the uniform across the cloud, but has position dependence, we have Eq. (7.1)

$$r_{\text{rot}}(\mathbf{r}) = \bar{r}_{\text{rot}} + \delta r_{\text{rot}}(\mathbf{r}).$$

Then, the corresponding shift from $\mathcal{B}_{40}$ becomes

$$\delta f_{40}(\mathbf{r}) \propto \mathcal{B}_{40} [\bar{r}_{\text{rot}}^3 + 2\bar{r}_{\text{rot}}(x^2 + y^2 - 2z^2) + 3\bar{r}_{\text{rot}}^2 \delta r_{\text{rot}}(\mathbf{r}) + 2(x^2 + y^2 - 2z^2) \delta r_{\text{rot}}(\mathbf{r})], \quad (7.10)$$

to first order in δr_{rot} .

The expansion in Eq. (7.10) highlights two effects. First, the spatial variation in r_{rot} turns the nominally common-mode term $\mathcal{B}_{40}\bar{r}_{\text{rot}}^3$ into a dephasing contribution through the term $3\mathcal{B}_{40}\bar{r}_{\text{rot}}^2 \delta r_{\text{rot}}(\mathbf{r})$. Second, it also modifies the curvature term itself.

Similarly, for the $l = 3$ terms,

$$B_{3m}^{\text{eff}}(\mathbf{r}) \propto \mathcal{B}_{3m} r_{\text{rot}}(\mathbf{r}) x_i,$$

where x_i denotes one of the secular coordinates x , y , or z . Expanding about $r_{\text{rot}}(\mathbf{r}) = \bar{r}_{\text{rot}} + \delta r_{\text{rot}}(\mathbf{r})$, we have

$$B_{3m}^{\text{eff}}(\mathbf{r}) \propto \mathcal{B}_{3m} \bar{r}_{\text{rot}} x_i + \mathcal{B}_{3m} x_i \delta r_{\text{rot}}(\mathbf{r}).$$

The first term is linear in the secular coordinate and averages away. The second term, however, need not average away. If the rotating electric field has a inhomogeneity that was linear in position,

the product $x_i \delta r_{\text{rot}}(\mathbf{r})$ contains quadratic terms such as x^2 , xy , or xz . The square terms survive the secular averaging and can contribute to dephasing.

This example makes it clear that in the rotating frame, electric and magnetic field inhomogeneity are not completely separate issues. A magnetic multipole that would be harmless for a spatially uniform rotating electric field can become relevant if the rotating-orbit radius varies spatially across the ion cloud. Therefore, once we start worrying about long Ramsey coherence, the practical question is no longer only what ideal multipole we apply, but also how uncontrolled higher-order magnetic multipoles combine with spatial inhomogeneity in the rotating electric field.

7.3.2 Residual magnetization and magnetic-field homogeneity

The discussion above shows that long Ramsey coherence requires not only temporal stability of the magnetic environment, but also a high degree of spatial magnetic homogeneity across the ion cloud. In practice, this means that stray fields near the trap must be minimized. Magnetized hardware, including the trap electrodes themselves, can generate uncontrolled higher-order spatial structure in the field and produce both frequency shifts and inhomogeneous dephasing.

While assembling the prototype trap, we found that several electrode pieces had measurable residual magnetization. The electrodes were fabricated from stainless steel parts, SS304 and SS316L. These materials are nominally nonmagnetic or only weakly magnetic. However, machining processes with magnetized tools left a noticeable residual magnetic remanent. For this reason, all electrode components, including rods, rings, and spines, were degaussed prior to installation. After degaussing process, each piece was screened individually with a magnetometer before assembly to make sure that the residual magnetization after the degaussing was small.

It is useful to estimate the level of B_{40} -type magnetic inhomogeneity that can be tolerated before it limits the Ramsey coherence. As discussed in the next section, I define τ_{innate} to be the $1/e$ coherence time expected for a noninteracting thermal ensemble in the presence of a quadratic frequency curvature. For nominal trap parameters $T \approx 15$ K and $f_{\text{sec}} \approx 2000$ Hz along all three

axes, requiring $\tau_{\text{innate}} \approx 20$ s corresponds to a quadratic frequency curvature

$$f_2 \approx 0.185 \text{ Hz/cm}^2,$$

for the convention

$$\delta f(\mathbf{r}) = f_2 \left(-\frac{x^2}{2} - \frac{y^2}{2} + z^2 \right).$$

Using $f_{\text{rot}} = 147.5$ kHz and $\mathcal{E}_{\text{rot}} = 60$ V/cm, this curvature corresponds to a tolerated magnetic field variation of order

$$|\partial_z^3 B_z| \approx 4.4 \times 10^{-5} \text{ G/cm}^3, \quad |\partial_x^2 \partial_z B_z| \approx 2.2 \times 10^{-5} \text{ G/cm}^3.$$

These numbers should be understood as order-of-magnitude tolerances for the third spatial derivatives of the magnetic field near the ions.⁶

Near the octagon region, many nearby components can contribute to residual magnetic field and gradients. This region contains viewports and nearby optics that are mounted on the vertical breadboard, all of which retain some residual magnetization. This is also the region where the Ramsey coherence is not expected to be the longest, due to the BBR excitation at room temperature. For precision spectroscopy, we therefore transport the ions to the science region, where the environment is colder, quieter, and magnetically cleaner.

7.4 Collisional narrowing

The previous section identified the leading spatial magnetic inhomogeneities that can generate dephasing across the ion cloud. A natural next question is how we can quantify such dephasing given some inhomogeneities.

For a noninteracting thermal ensemble in a harmonic trap, each ion follows a well-defined secular orbit, and therefore samples a well-defined time-averaged detuning set by its oscillation amplitudes. Different ions have different amplitudes, and as a result, they accumulate different

⁶ The precise numerical conversion depends on the convention used for the magnetic multipole normalization and for the definition of the quadratic frequency curvature. The values quoted here are therefore intended as order-of-magnitude design targets.

Ramsey phases leading to the decay in the ensemble coherence. However, in a real ion cloud, Coulomb interactions between ions perturb the single-particle trajectories and redistribute motional energy among the ions. As a result, the dephasing expected from a noninteracting model may not map directly to the observed coherence decay.

This leads to the phenomenon usually called motional narrowing; here, we refer to it as *collisional narrowing*. If ion–ion interactions redistribute motion rapidly compared to the noninteracting dephasing time, then the effect of the spatial inhomogeneity can be partially averaged during the Ramsey evolution, leading to a longer observed coherence time. Conversely, if ions accumulate significant phase spread before collisions substantially modify the motion, the coherence decay is expected to remain close to the noninteracting prediction.

In this section, I first establish the noninteracting model for dephasing from spatial magnetic inhomogeneity and define the corresponding timescale τ_{innate} . Then, I introduce several thermalization rates extracted from LAMMPS simulations for ion–ion collisions. Then, I compare these thermalization timescales against the reference timescale τ_{innate} to examine whether these thermalization timescales serve as useful empirical proxies for the “onset” of collisional narrowing.

7.4.1 Inhomogeneous dephasing in noninteracting gas

To build intuition for the inhomogeneous dephasing due to spatial inhomogeneity, it is useful to first consider a noninteracting thermal ensemble of ions in a harmonic trap. Suppose the local transition frequency is given by

$$f(\mathbf{r}) = f_0 + f_2 \left(-\frac{x^2}{2} - \frac{y^2}{2} + z^2 \right), \quad (7.11)$$

where f_0 is the spatially averaged transition frequency and f_2 parametrizes the curvature of the inhomogeneous shift. This particular form captures the quadratic frequency variation associated with the $\mathcal{B}_{40}$ contribution discussed above.⁷

⁷ Here, a little different than above, I have defined the curvature to be $-x^2/2 - y^2/2 + z^2$. In other words, the curvature here is different by a factor of $-1/2$, compared to the previous section.

For a single ion under secular motion in a harmonic trap, their motion follows

$$x(t) = x_0 \cos(\omega_x t + \phi_x), \quad y(t) = y_0 \cos(\omega_y t + \phi_y), \quad z(t) = z_0 \cos(\omega_z t + \phi_z), \quad (7.12)$$

where ω_x, ω_y and ω_z are the secular frequencies and x_0, y_0 , and z_0 are the corresponding oscillation amplitudes. For interrogation times long compared with the secular periods, the oscillatory terms average over the orbit, and the single-ion time-averaged detuning is

$$\Delta f \equiv \frac{1}{t_R} \int_0^{t_R} [f(\mathbf{r}(t)) - f_0] dt = \frac{f_2}{4} (-x_0^2 - y_0^2 + 2z_0^2). \quad (7.13)$$

Thus, ions with different secular amplitudes accumulate different Ramsey phases—this is the basic origin of inhomogeneous dephasing from the $\mathcal{B}_{40}$ curvature.

Since ions are in a harmonic trap, it is convenient to rewrite Eq (7.13) in terms of the secular energies

$$E_x = \frac{1}{2} m \omega_x^2 x_0^2, \quad E_y = \frac{1}{2} m \omega_y^2 y_0^2, \quad E_z = \frac{1}{2} m \omega_z^2 z_0^2. \quad (7.14)$$

Then, we get

$$\Delta f = a_x E_x + a_y E_y + a_z E_z \quad (7.15)$$

with

$$a_x = -\frac{f_2}{2m\omega_x^2}, \quad a_y = -\frac{f_2}{2m\omega_y^2}, \quad a_z = \frac{f_2}{m\omega_z^2}. \quad (7.16)$$

For a thermal ensemble at temperature T , each secular mode has an exponential energy distribution,

$$p(E_j) = \beta e^{-\beta E_j}, \quad \beta \equiv \frac{1}{k_B T},$$

with mean $\langle E_j \rangle = k_B T$ and the corresponding variance $\text{Var}(E_j) = (k_B T)^2$ [96, 97]. Then, the ensemble-averaged detuning is

$$\langle \Delta f \rangle = f_2 k_B T \left(-\frac{1}{2m\omega_x^2} - \frac{1}{2m\omega_y^2} + \frac{1}{m\omega_z^2} \right), \quad (7.17)$$

and the variance of the frequency distribution is

$$\sigma_f^2 = (f_2 k_B T)^2 \left(\frac{1}{4m^2\omega_x^4} + \frac{1}{4m^2\omega_y^4} + \frac{1}{m^2\omega_z^4} \right). \quad (7.18)$$

Equation (7.17) shows the mean shift vanishes if the three secular frequencies are equal. However, even with the special choice of the secular frequencies, Eq (7.18) shows that the width of the frequency distribution generally remains nonzero.

The spread in the time-averaged frequency directly determines the decay of the Ramsey contrast:

$$C(t_R) = \left| \langle e^{i2\pi\Delta f t_R} \rangle \right|, \quad (7.19)$$

where Δf is given by Eq. (7.17). Using the factorized thermal distribution for E_x, E_y, E_z , this becomes

$$C(t_R) = \prod_{j=x,y,z} \left| \int_0^\infty \beta e^{-\beta E_j} e^{i2\pi a_j E_j t_R} dE_j \right| = \prod_{j=x,y,z} \frac{1}{\sqrt{1 + (2\pi a_j k_B T t_R)^2}}. \quad (7.20)$$

Equation 7.20 is not Gaussian. However, at short times $|2\pi a_j k_B T t_R| \ll 1$, expanding the equation gives

$$C(t_R) \approx e^{-2\pi^2 \sigma_f^2 t_R^2},$$

so the initial contrast decays as a Gaussian with variance given in Eq (7.18). At long times, the exact result is a power-law tail; if all three coefficients a_x, a_y, a_z are nonzero, then $C(t_R) \propto t_R^{-3}$.

In the discussion below, I define the $1/e$ time of the noninteracting model contrast as τ_{innate} , obtained by numerically solving

$$\prod_{j=x,y,z} \frac{1}{\sqrt{1 + (2\pi a_j k_B T \tau_{\text{innate}})^2}} = e^{-1}, \quad (7.21)$$

given some f_2, T , and secular frequencies for a_j . This timescale provides the natural reference against which collision narrowing is assessed.

7.4.2 Thermalization timescales from LAMMPS

To characterize the ion-ion collision dynamics that may be relevant for collisional narrowing, the undergraduate student Rohan Kompella carried out additional LAMMPS simulations as an extension of the rf-heating study. From these simulations, we extract two thermalization timescales: the longitudinal thermalization (LGT) time and the cross-dimensional thermalization (CDT) time.

For the LGT analysis, we first define the fluctuation of the single-ion energy for ion j about the ensemble mean along a particular axis $k \in \{x, y, z\}$ as

$$E_k^{j*}(t) = E_k^j(t) - \frac{1}{N} \sum_{i=1}^N E_k^i(t). \quad (7.22)$$

We then compute the normalized autocorrelation function for each ion,

$$f_{k,j}(\tau) = \frac{\int_0^{t_f} E_k^{j*}(t) E_k^{j*}(t + \tau) dt}{\int_0^{t_f} E_k^{j*}(t) E_k^{j*}(t) dt}. \quad (7.23)$$

After averaging over all ions, this gives us the axis-dependent correlation function $f_k(\tau)$, shown in Fig. 7.3. We fit its decay to an exponential form

$$f_k(\tau) \propto e^{-\tau/\tau_{\text{LGT},k}},$$

and define the corresponding LGT rate as

$$\Gamma_{\text{LGT},k} = \tau_{\text{LGT},k}^{-1}.$$

The LGT rate quantifies how quickly an individual ion loses memory of its excess energy along a




Figure 7.3: Example extraction of the longitudinal thermalization time from a representative LAMMPS dataset. The normalized autocorrelation functions $f_k(\tau)$ for $k = x, y, z$ are obtained from the fluctuations of the single-ion energies about the ensemble mean along each axis, as in Eq. (7.23). Exponential fits to the decay are used to define the corresponding LGT timescales $\tau_{\text{LGT},k}$.

given axis.

By contrast, the CDT rate characterizes redistribution of energy between different axes. To extract the CDT rate, we initialize the ensemble with excess energy along the x direction and monitor the relaxation of the temperature anisotropy defined as

$$T_x^*(t) = T_x(t) - \bar{T}(t), \quad \bar{T} = \frac{1}{3} [T_x(t) + T_y(t) + T_z(t)]. \quad (7.24)$$

We then fit

$$T_x^*(t) = Ae^{-t/\tau_{\text{CDT}}},$$

as shown in Fig. 7.4 and define the CDT rate as

$$\Gamma_{\text{CDT}} = \tau_{\text{CDT}}^{-1}.$$




Figure 7.4: Example extraction of the cross-dimensional thermalization time from a representative LAMMPS dataset. After initializing excess energy along the x direction, the temperature anisotropy $T_x^*(t)$ relaxes toward zero as energy is redistributed among the three motional axes. An exponential fit to this decay is used to define τ_{CDT} .

As in the rf-heating studies discussed in Section 5.4.2, we evaluated the LGT and CDT rates over a broad range of simulation parameters, including $T/T_{\text{crossover}}$ and naive collision rate. In addition, to reduce any systematic dependence of these timescales on certain parameters, we varied

the naive collision rate by moving through different regions of parameter space, rather than by changing only one parameter. This makes the extracted scaling less sensitive to a particular choice of density, temperature, or trap parameters. The resulting dependence of the LGT and CDT rates on the naive collision rate is shown in Fig. 7.5.

Over the range of parameters explored in these simulations, the LGT rate is systematically faster than the CDT rate, especially at low naive collision rate. At the moment, the origin of the different scaling laws between the LGT and CDT rates is not fully understood. Additional simulations in a box-like potential suggest that the difference between the LGT and CDT rates is reduced in a nonharmonic potential geometry; this may suggest a possibility of resonantly enhanced ion-ion dynamics that couples the individual ion’s harmonic secular motion and mean field from the rest of the ion cloud in the Paul trap. Because a separate manuscript on the ion dynamics is in preparations, I do not attempt a more detailed analysis here. For the purposes of the rest of this chapter, the important point is that these simulations provide an empirical way to estimate LGT and CDT rates as functions of the trap and cloud parameters.

7.4.3 LGT rate, CDT rate, and actual “narrowing” rate

The thermalization rates extracted above are related to more standard quantities discussed in the literature, although our terminology differs slightly. The quantity we call the CDT rate is closely analogous to the cross-dimensional relaxation or rethermalization, extracted by looking at the decay of an imposed temperature or momentum-space anisotropy [98–100]. In plasma physics language, the CDT rate is also analogous to the temperature-anisotropy relaxation rate [101]. The quantity we call the LGT rate here is less standard, but it is closest to a single-particle, same-axis energy decorrelation or memory-loss rate.

It is not expected that either the LGT or CDT rate should equal to the Ramsey narrowing rate. In the literature, the timescale that is most relevant for motional narrowing is the autocorrelation time of the instantaneous detuning,

$$C_{\delta\omega}(\tau) = \langle \delta\omega(t) \delta\omega(t + \tau) \rangle, \quad (7.25)$$




Figure 7.5: Longitudinal thermalization (LGT) and cross-dimensional thermalization (CDT) rates as functions of the naive collision rate (NCR) for box-potential and Paul-trap simulations on log-log axes. The points show the rates obtained from numerical fits to the simulated trajectories, while the dashed lines are power-law fits used to summarize the scaling with NCR. For the Paul trap case, the LGT rate is systematically faster than the CDT rate, with different scaling to NCR.

or equivalently the timescale over which a particle loses memory of its motion that is responsible for the detuning [102]. Here, we do not directly compute this detuning autocorrelation function. Instead, we use the CDT and LGT timescales as proxies for different aspects of the collisional dynamics.

This distinction is especially important when comparing to some work in the cold-atom system such as Ref. [103]. Coslovsky *et al.* defines an energy-decorrelation rate Γ to be related to an average collision rate Γ_{coll} through the collision softness [103],

$$s = \Gamma_{\text{coll}}/\Gamma. \quad (7.26)$$

In this framework, $s \sim 1$ corresponds to hard collisions; collisions with $s \gg 1$ indicates that many collisions are required to decorrelate the energy.

The physical picture of collisional narrowing is analogous to our trapped-ion case, but the mapping need not be exactly the same. In particular, our LAMMPS simulations show that in a Paul trap, the LGT rate is not simply linear with the naive collision rate over the parameter range of interest. This suggests that the decorrelation dynamics may be influenced not only by the collision rate Γ_{coll} , but also by trap-specific effects such as secular motion and mean field dynamics. For this reason, I treat Γ_{LGT} and Γ_{CDT} as empirical proxies for the narrowing rate, rather than with Γ_{coll} .

In Refs. [103, 104], the authors show that in the collisional-narrowing regime, the Ramsey coherence envelope can often be appropriated by a Kubo-type form

$$C_R(t) \sim \exp \left[-\frac{\alpha}{\Gamma^2 \tau^2} (e^{-\Gamma t} + \Gamma t - 1) \right]. \quad (7.27)$$

Here, τ is the bare dephasing time, in our case the dephasing time we would expect in the noninteracting model $\tau = \tau_{\text{innate}}$. The value α is a constant whose value depends on the model and on the definition of τ ; the exact prefactors are not important in this discussion, therefore below, I use $\tau = \tau_{\text{innate}}$

We can look at the short time and long time behaviors of Eq.(7.27). At short time such that

$\Gamma t \ll 1$, we have

$$e^{-\Gamma t} + \Gamma t - 1 \approx (\Gamma t)^2/2 + O((\Gamma t)^3).$$

Then, the envelope reduces to

$$C(t) \approx \exp \left[-\frac{\alpha}{2} \left(\frac{t}{\tau_{\text{innate}}} \right)^2 \right],$$

which corresponds to the no-collision result shown in Eq. (7.20). Again, the precise numerical prefactor here depends on the convention used to define the bare dephasing timescale τ . In other words, in this short time limit, collisions have not had the time to reshuffle the detuning yet. Therefore, the initial decay is still set by the inhomogeneity.

At long times $\Gamma t \gg 1$, the envelope crosses over to an exponential form,

$$C(t) \approx \exp \left[-\frac{\alpha}{\Gamma \tau_{\text{innate}}^2} t \right],$$

corresponding to a narrowed coherence time

$$t_{\text{narrowed}} \approx \Gamma \tau_{\text{innate}}^2 / \alpha = \tau_{\text{innate}} \left[\frac{\Gamma \tau_{\text{innate}}}{\alpha} \right].$$

For $\Gamma \tau_{\text{innate}} \ll 1$, collisions are too slow to “reshuffle” the detuning before the innate dephasing happens, and little narrowing is expected. For $\Gamma \tau_{\text{innate}} \gtrsim 1$, the system enters the crossover regime, and for $\Gamma \tau_{\text{innate}} \gg 1$, the detuning is rapidly decorrelated, leading to strong collisional narrowing. In particular, the narrowed $1/e$ coherence time can be much larger than the innate dephasing time.

7.4.4 Numerical Simulations of Collisional Narrowing

From the LAMMPS simulations, we extracted the LGT and CDT times rather than the detuning-decorrelation rate Γ that appears in Eq. (7.26). We do also have a naive collision rate, but the discussion above shows that it is not obvious that this rate alone should govern the Ramsey narrowing. For these reasons, I directly compare the bare noninteracting dephasing time τ_{innate} to the LGT and CDT times extracted from LAMMPS, and ask which of these is more useful empirical predictor for the onset of narrowing.

To make this comparison, I numerically accumulated the Ramsey phase along simulated ion trajectories. In one case, I used trajectories from the LAMMPS simulations. In the other, I generated a noninteracting ensemble with the same number of ions and temperature T , initialized each ion from the corresponding thermal distribution and then evolved in a purely harmonic trap.

To both trajectory sets, I introduced the same model inhomogeneous shift

$$f(\mathbf{r}) = f_0 + f_2 \left(\frac{x^2}{2} + \frac{y^2}{2} - z^2 \right),$$

and computed the ensemble Ramsey coherence as a function of interrogation time t_R . The noninteracting $1/e$ time τ_{innate} was obtained from Eq. (7.20).

I performed this comparison for three representative datasets from LAMMPS with substantially different LGT and CDT times, as summarized in Table 7.2.

	Naive collision time	CDT time	LGT time
Dataset 1	2.1 s	80 ms	30 ms
Dataset 2	4.2 s	260 ms	40 ms
Dataset 3	240 s	11.6 s	400 ms

Table 7.2: Representative LAMMPS datasets used in the numerical study of collisional narrowing. For each dataset, the table lists the naive collision time, which is (naive collision rate) $^{-1}$, together with the extracted cross-dimensional thermalization time τ_{CDT} and longitudinal thermalization time τ_{LGT} . The three datasets were chosen to span a broad range of collisional dynamics, from the case where τ_{CDT} and τ_{LGT} differ only modestly to the case where they are strongly separated. This allows a direct comparison of which timescales more closely tracks the onset of collisional narrowing.

Figures 7.6–7.8 show the resulting ensemble coherence decay. In each figure, the inhomogeneity strength f_2 increases from upper left to lower right, and the title of each panel reports the corresponding τ_{innate} . The available LAMMPS trajectories extend to 1 s.

The simulations show a clear qualitative crossover; when τ_{innate} is long compared to the relevant collision-induced decorrelation time, the coherence obtained from the LAMMPS trajectories decays significantly more slowly than the noninteracting prediction, demonstrating collisional narrowing. As f_2 is increased and τ_{innate} becomes correspondingly shorter, the difference is progressively reduced, and eventually the LAMMPS and noninteracting curves become nearly




Figure 7.6: Simulated Ramsey coherence decay for Dataset 1. In each panel, the blue curve shows the ensemble coherence obtained by numerically accumulating Ramsey phase along noninteracting ion trajectories initialized from the corresponding thermal distribution and evolved in a purely harmonic trap. The orange curve shows the same numerical phase accumulation performed along the LAMMPS trajectories. The inhomogeneity strength f_2 increases from upper left to lower right, and the title of each panel gives the corresponding noninteracting $1/e$ time τ_{innate} . For this dataset, $\tau_{\text{CDT}} \approx 80$ ms and $\tau_{\text{LGT}} \approx 30$ ms. When τ_{innate} is long compared to the τ_{LGT} and τ_{CDT} , the LAMMPS-based coherence decays more slowly than the noninteracting prediction, indicating collisional narrowing; as f_2 is increased and τ_{innate} becomes shorter, the difference between the two curves is progressively reduced. The LAMMPS-based curves are shown only to 1 s, set by the available trajectory length.




Figure 7.7: Simulated Ramsey coherence decay for Dataset 2. As in Fig. 7.6, the blue curve is obtained by numerically accumulating phase along noninteracting trajectories drawn from the thermal distribution, while the orange curve is obtained from the LAMMPS trajectories. The inhomogeneity strength f_2 increases from upper left to lower right, and each panel title reports the corresponding noninteracting $1/e$ time τ_{innate} . For this dataset, $\tau_{\text{CDT}} \approx 260$ ms and $\tau_{\text{LGT}} \approx 40$ ms. Because τ_{CDT} and τ_{LGT} are more strongly separated than in Dataset 1, this dataset more clearly shows that the crossover from narrowed to nearly non-narrowed behavior occurs when τ_{innate} becomes comparable to the shorter LGT timescale, rather than the much longer CDT timescale.




Figure 7.8: Simulated Ramsey coherence decay for Dataset 3. The blue curve shows the coherence obtained by numerically accumulating Ramsey phase along noninteracting thermal trajectories in a harmonic trap, and the orange curve shows the corresponding result obtained from LAMMPS trajectories. The inhomogeneity strength f_2 increases from upper left to lower right, and the panel titles report the corresponding noninteracting $1/e$ time τ_{innate} . For this dataset, $\tau_{\text{CDT}} \approx 11.6$ s and $\tau_{\text{LGT}} \approx 400$ ms. Even though the CDT time is far longer than the 1 s LAMMPS trajectory window, the crossover from narrowed to non-narrowed behavior still occurs on the sub-second scale. This therefore provides the strongest evidence among the three datasets that an LGT-like timescale is the more useful empirical predictor for the onset of collisional narrowing.

indistinguishable.

Dataset 1 already demonstrates the presence of collisional narrowing, but since τ_{CDT} and τ_{LGT} differ only by a factor of $\sim$ three, it is difficult to determine which timescale is more relevant from this dataset alone. Datasets 2 and 3 are more informative because the two timescales are much more different. In both cases, it seems like the crossover between the narrowed and non-narrowed behavior appears when τ_{innate} becomes comparable to the τ_{LGT} , up to some factor of order unity. In particular, for Dataset 3, the CDT time is far longer than the entire 1 second trajectory window, yet we still observe the crossover on the few hundred millisecond scale, which is close to the τ_{LGT} . This suggests that the onset of collisional narrowing is more closely associated with an LGT-like timescale than with the CDT time.

I emphasize here that these simulations do not prove that the LGT rate is the unique or fundamental narrowing rate. The more fundamental quantity would be the autocorrelation time of the actual instantaneous detuning experienced by each ion, as in Eq. (7.25). The results from this section instead support the conclusion that the LGT time is a better empirical predictor than the CDT time, for estimating when collisional narrowing begins to matter in our system.

This conclusion also highlights the difference of our system from the cold-atom picture as in Refs. [103, 104]. It appears that in a Paul trap, the rate relevant for narrowing is not determined solely by the naive collision rate. Instead, our LAMMPS simulations suggest that trap-specific dynamics can modify the decorrelation process. From a practical point of view, however, the present analysis is still useful—it provides a way to estimate when collisional narrowing should be expected for a given set of trap and cloud parameters. Therefore, we can estimate how much magnetic-field inhomogeneity can be tolerated while preserving long Ramsey coherence times.

7.5 Dephasing with applied B_{40}

To test collisional narrowing experimentally, we intentionally apply a controlled magnetic inhomogeneity with B_{40} character to the trap. This provides a way to experimentally vary the curvature of the effective magnetic field. By increasing the applied B_{40} field, we can deliberately

shorten the innate dephasing time τ_{innate} and then compare the observed Ramsey coherence time against the expectation from the noninteracting gas model and the collisional narrowing picture discussed above.

7.5.1 B_{40} gradient coils

We wound a pair of coils with 15 turns in an anti-Helmholtz configuration, as shown in Fig. 7.9. The different multipole terms from an ideal circular-loop in an anti-Helmholtz configuration is given in Appendix B. Near the trap center, the applied field contains the desired B_{40} -like curvature, as well as a “parasitic” B_{20} -like component.

Using the circular-loop model and the rotating-frame average discussed in the previous section, the expected frequency shift from these coils can be parameterized as

$$f(x, y, z) = f_0 + f_2 \left(-\frac{x^2}{2} - \frac{y^2}{2} + z^2 \right). \quad (7.28)$$

The corresponding estimates for $f_{\text{rot}} = 147.5$ kHz and $\mathcal{E}_{\text{rot}} = 60$ V/cm are

$$\left| \frac{f_0}{I_{B_{40}}} \right| = 640 \text{ Hz/A}, \quad \left| \frac{f_2}{I_{B_{40}}} \right| = 20 \text{ Hz/cm}^2/\text{A}.$$

Here f_0 denotes the common-mode shift, here dominated by the parasitic B_{20} -like component of the B_{40} coils⁸

To experimentally calibrate these coefficients, we park ions at different axial locations z and measure the corresponding Ramsey frequency f_{Ramsey} . Then, for each applied B_{40} current, the

⁸ The intrinsic common-mode shift arising from the B_{40} as shown in Eq (7.9) is much smaller than the parasitic B_{20} -like shift. In the convention

$$f(x, y, z) = f_0 + f_2 \left(-\frac{x^2}{2} - \frac{y^2}{2} + z^2 \right),$$

the B_{40} contribution has the form

$$\delta f_{40} \propto r_{\text{rot}}^3 + 2r_{\text{rot}}(x^2 + y^2 - 2z^2) \propto f_0^{(40)} + f_2^{(40)} \left(-\frac{x^2}{2} - \frac{y^2}{2} + z^2 \right),$$

so the common-mode part is

$$f_0^{(40)} = -\frac{f_2^{(40)} r_{\text{rot}}^2}{4}.$$

This corresponds to $f_0^{(40)}/I_{B_{40}} \approx 0.36$ Hz/A. This number, compared to the parasitic B_{20} -like contribution is small, therefore we ignore this contribution here.




Figure 7.9: Photograph of the magnetic-field coils mounted around the vacuum chamber near the octagon. The B_{axgrad} coils provide the axial magnetic field gradient used to generate the effective rotating-frame bias field, while the B_{40} coils provide a controllable higher-order magnetic curvature used in the collisional-narrowing measurements. In the experiments described in this chapter, current in the B_{axgrad} coils was also used to compensate the large parasitic B_{20} -like common-mode shift produced by the B_{40} coils.

measured spatial profile is fit to

$$f_{\text{Rmasey}}(z) = f_2 z^2 + az + c$$

where c is an overall offset, a is a residual linear term, and the f_2 is the quadratic curvature of interest. Here, the quadratic coefficient f_2 provides a robust measure of the applied B_{40} -like inhomogeneity.

Because the B_{40} coils also generate a large B_{20} -like shift, we simultaneously apply current to the B_{axgrad} coils to keep the Ramsey frequency near the trap center approximately fixed at ~ 200 Hz, as the B_{40} current is varied.

Figure 7.10 shows the measured quadratic coefficient f_2 from the B_{40} coils as a function of B_{40} . The extracted curvature varies linearly with the applied coil current, indicating that the coils provide a controllable B_{40} -like inhomogeneity. From a linear fit, I obtain

$$\frac{f_2^{\text{meas}}}{I_{B_{40}}} \approx 12 \text{ Hz/cm}^2/\text{A}.$$

The required B_{axgrad} compensation current also provides an experimental estimate of the parasitic value of f_0 from the B_{40} coils. Using the independently measured B_{axgrad} calibration, the f_0 from the B_{40} coils is inferred to be

$$\frac{f_0^{\text{meas}}}{I_{B_{40}}} \approx 600 \text{ Hz/A}.$$

The measured parasitic f_0 of the B_{40} coils is in reasonable agreement with the circular-loop estimate of 640 Hz/A. By contrast, the measured curvature f_2 is smaller than the ideal prediction by roughly a factor of 0.6. The discrepancy could be that the curvature depends on the third spatial derivative of the field and is therefore more sensitive to non-ideal coil geometry and alignment errors. It could also come from the misalignment in B_{axgrad} coils as these coils could also in theory have “parasitic” B_{40} -like contribution from imperfections.




Figure 7.10: Measured quadratic frequency curvature f_2 as a function of the applied B_{40} coil current. For each current, ions are parked at several axial positions and the measure Ramsey frequency is fit to a quadratic function of z . As the B_{40} coil current is varied, the current for the B_{axgrad} coils are adjusted to keep the Ramsey frequency near 200 Hz, thereby canceling most of the parasitic B_{20} -like shift from the B_{40} coils. The extracted f_2 varies linearly with B_{40} current, as expected, with a slope of magnitude $\sim 12 \text{ Hz/cm}^2/\text{A}$.

7.5.2 Experimental data

Using the B_{40} coils, I applied controlled values of the quadratic curvature f_2 and measured the corresponding Ramsey decay time. The applied values were

$$3.6, 7.5, 15, 22.5, 30.7 \quad \text{Hz/cm}^2.$$

To compare the measurements with the noninteracting model, I converted each applied f_2 into the corresponding bare dephasing time τ_{innate} using the model of Sec. 7.4.1. Because the absolute ion temperature is not very well measured with our current thermometry, this conversion was repeated for several assumed temperatures. Figure 7.11 therefore shows the same measured coherence data replotted against the inferred τ_{innate} for different values of T .

The top and bottom rows correspond to two experimental datasets with different ion numbers, approximately $N \sim 8000$ and $N \sim 1500$, respectively. The grey diagonal line indicates $\tau_{\text{obs}} = \tau_{\text{innate}}$, as a reference line. Points above this reference line are consistent with collisional narrowing, since the measured coherence time persists longer than expected from the bare dephasing time τ_{innate} . By contrast, points on or below this line may suggest little to no narrowing; however, one must be careful to not exclude narrowing right away—the measured decay can also be shortened by other mechanisms not included in the noninteracting model.

With this caveat, Fig. 7.11 suggests that the interpretation is strongly temperature dependent. For the $N \sim 1500$ dataset, assuming $T \gtrsim 10$ K, most of the data are above the $\tau_{\text{obs}} = \tau_{\text{innate}}$ line, which is consistent with a substantial degree of collisional narrowing. At $T = 5$ K, however, the same dataset sits closer to the crossover between narrowed and non-narrowed behavior. The $N \sim 8000$ dataset is even more sensitive to the assumed temperature: for low assumed T , the data are consistent with little to no narrowing, whereas for higher assumed T they move into the narrowed regime. Therefore, the present data are consistent with collisional narrowing over the plausible temperature range, but they do not establish it unambiguously without an independent estimate of the cloud temperature.




Figure 7.11: Measured Ramsey decay time τ_{obs} plotted against the inferred noninteracting dephasing time τ_{innate} for several assumed ion temperatures. The top row shows the dataset with approximately $N \sim 8000$ ions, and the bottom row shows the dataset with approximately $N \sim 1500$ ions. Each column corresponds to a different assumed temperature used to convert the applied curvature f_2 into τ_{innate} . The solid gray diagonal indicates $\tau_{\text{obs}} = \tau_{\text{innate}}$. Points above this line are consistent with collisional narrowing. The dashed horizontal line in each panel marks the baseline coherence time, i.e. with no applied additional f_2 .

Another big complication is that ion–ion collisions can influence the Ramsey signal in more than one way. On the one hand, more collisions mean more average over the spatially varying detuning and thereby lengthen the observed coherence time. On the other hand, more collisions also mean more redistribution among m_F states, as discussed in the previous chapter, which shortens the observed contrast lifetime. The measured Ramsey decay time therefore reflects a competition between narrowing and collision-induced decoherence, and it should not be expected to vary monotonically with ion number or collision rate.

For this reason, the decay times extracted here should be viewed as empirical summary parameters, rather than as direct measurements of the narrowing rate. In the present analysis, each dataset was fit with a simple exponential envelope to obtain a characteristic τ_{obs} . This is useful for comparing datasets on the same footing, but it is not the optimal lineshape model for collisionally narrowed Ramsey decay, as we see in Refs. [102–104]. A more complete measurement would be recoding Ramsey fringes at many interrogation times, extract the contrast point-by-point, and then fit the resulting envelope to the Kubo form of Eq. (7.25). Such an analysis could provide a more quantitative comparison between the experimental narrowing dynamics and the proxy rates inferred from simulation.

A complementary way to probe collisional narrowing would be to use different observables, other than the fitted coherence time which can have a complicated lineshape. For instance, we could perform a Hahn-echo measurement, as in Ref. [103], which may isolate the collision-induced reshuffling more cleanly. Another approach is to measure the Ramsey frequency shift, rather than coherence time. In such a protocol, we would selectively optically pump only part of the cloud while applying a large f_2 . In this case, initially, the subensemble that was illuminated with the optical pumping would have a different average Ramsey frequency. By varying the delay between optical pumping and the first $\pi/2$ pulse, one could then watch this frequency shift relax toward the ensemble average as ion–ion interactions redistribute the motional amplitudes. This measurement would directly probe how the ions lose memory of their initial spatial inhomogeneity, which will be very useful insight for systematics.

Chapter 8

Towards continuous-mode operation

The final Gen. III apparatus for the JILA eEDM experiment will be operable in two possible modes: *throughput* mode and *batch* mode operation, as illustrated in Figure 8.1.

In the *throughput* operation mode, the experiment would run in a nearly continuous fashion. Ion packets are produced near the source chamber where the neutral molecules are generated, and state preparation is performed shortly after ionization. The ions then travel along the trap while undergoing Ramsey evolution. State detection occurs near the detector at the downstream end of the apparatus. In this configuration, the Ramsey free-precession time is determined by the physical travel time of the ion packet between the preparation and detection regions.

In *batch* mode operation, we instead sequentially load the trap with multiple ion clouds until all of the trapping “buckets” along the rod structure are filled. Once the trap is fully populated, we simultaneously perform Ramsey evolution on all stored ion packets. After the evolution period, the ion clouds are detected sequentially, beginning with the packet that was loaded last. In this mode the Ramsey time is set by the time that the ions remain stored in the trap.

At the time of writing, it is not clear which will be the preferred operating mode for the final Gen. III machine. Both approaches present distinct technical challenges that must be addressed before full implementation. In this chapter, we investigate several of these challenges and demonstrate techniques that move the experiment toward realizing either throughput or batch-mode operation.




Figure 8.1: Conceptual illustration of the two possible operating modes of the Gen. III eEDM apparatus. **Top:** In throughput mode, ion packets are continuously produced near the source and travel along the trap while undergoing Ramsey evolution before being detected downstream. **Bottom:** In batch mode, multiple ion packets are loaded sequentially into spatial “buckets” along the trap. After the trap is filled, all packets undergo Ramsey evolution simultaneously before being detected one by one.

8.1 Ramsey fringes during translation

In the throughput-mode operation described above, the ions undergo Ramsey evolution while they are physically translating along the trapping region. It is therefore essential that the quantum phase coherence of the ions be preserved during transport. Any additional phase noise or decoherence introduced by the translation process would directly reduce the contrast of the Ramsey fringe and degrade the sensitivity of the eEDM measurement.

Since the current prototype apparatus does not support full throughput-mode operation, we performed a simplified test of this requirement. In this experiment, we apply the first Ramsey $\pi/2$ pulse while the ions are located in the cold region of the trap. Immediately after the first pulse, we begin translating the ions back toward the octagon region while allowing the system to undergo free evolution.

Figure 8.2 shows the resulting Ramsey fringes: without translation and with translation during Ramsey evolution. In this measurement we varied the time that the ions spend in the cold region prior to translation. Therefore, we expect the fringe frequency to correspond to the spin-precession frequency f associated with the Science-region magnetic field, but with an phase offset relative to the case in which the ions remain stationary during the Ramsey evolution.

This is consistent with what we observe. The oscillation frequency matches the expected value determined by the Science-region field, while the overall phase of the fringe differs slightly from the stationary case. Importantly, at these relatively short evolution times (approximately 300 ms), we do not observe any additional loss of contrast. This indicates that translating the ions during their quantum evolution does not introduce measurable decoherence at early times. These results provide initial evidence that Ramsey evolution during ion transport—an essential requirement for throughput-mode operation—is feasible in the final apparatus.




Figure 8.2: **Top:** Ramsey fringe measured while the ions are stationary at the Science region. **Bottom:** Ramsey fringe while the ions are translated during the free-precession period. The asymmetry signal is plotted as a function of Ramsey evolution time. Since the Ramsey evolution time at the Science-region is varied, the spin-precession frequency is the same as the case without translation, but a phase difference compared to the no translation case, as expected.

8.2 Catching and ejecting ions with $\mathcal{E}_{\text{rot}}$ continuously on

In the envisioned continuous or batch operation modes, we plan to keep the rotating electric field $\mathcal{E}_{\text{rot}}$ on between the preparation of successive ion packets. This choice is motivated by two primary considerations. First, the stability requirements on the rotating electric field are stringent. In order to achieve long coherence time and to control systematic effects in the eEDM measurement, the amplitude and phase of $\mathcal{E}_{\text{rot}}$ must remain stable at the level of 0.01%. If the field is repeatedly turned on and off, thermal effects in the rod drivers, particularly the op-amps, lead to small but measurable changes in both the amplitude and phase of the field. These changes exceed the required stability level.

The second motivation, and the more important one, arises from the long and continuous rod structure of the trap. In either throughput or batch mode, multiple ion packets may be simultaneously present in different regions of the trap and undergoing quantum evolution. If $\mathcal{E}_{\text{rot}}$ were ramped down while some ions are already evolving, their quantization axis would suddenly change, causing rapid loss of phase coherence. Maintaining $\mathcal{E}_{\text{rot}}$ continuously therefore ensures that previously prepared ion packets remain undisturbed.

However, loading and extracting ions while $\mathcal{E}_{\text{rot}}$ remains on is technically challenging. As discussed earlier in this thesis, the rotating electric field generates a large ponderomotive potential. With the new rod drivers we are able to operate at a rotation frequency of $f_{\text{rot}} = 300$ kHz, but even at this frequency the ponderomotive energy corresponds to roughly 23000 K for a field amplitude of 60 V/cm. This energy scale is many orders of magnitude larger than the typical kinetic energy of our trapped ions. In the following sections we investigate how ions can nevertheless be loaded into, and extracted from, the trap while $\mathcal{E}_{\text{rot}}$ remains active.

8.2.1 Catching ions

The neutral beam of ThF enters the trapping region with a forward velocity of approximately 800 m/s along the $-\hat{x}$ direction. After photoionization produces ThF^+ ions, we apply a stopping

electric field in the $+\hat{x}$ direction to cancel this forward velocity and bring the ions into the trapping region.

When $\mathcal{E}_{\text{rot}}$ is present, successful loading of ions corresponds to ensuring that the newly created ions are born onto the stable trajectory associated with the rotating electric field. In practice, this means matching both the velocity and phase of the ion motion to the rotating field.

To accomplish this, we first operate with a reduced rotating-field amplitude of $\mathcal{E}_{\text{rot}} \sim 40$ V/cm. At this field strength, the associated ponderomotive velocity v_{rot} approximately matches the forward velocity of the supersonic molecular beam. We then scan the starting phase of the rotating electric field in order to match the initial trajectory of the ions.




Figure 8.3: Number of trapped ions as a function of the starting phase of the rotating electric field when loading with reduced field amplitude ($\mathcal{E}_{\text{rot}} \approx 40$ V/cm). Ion survival occurs only when the field phase matches the direction of the rotating trajectory at the moment of ion creation.

With the phase parked near the optimal value (335°), we then attempt to load ions at the full rotating-field strength $\mathcal{E}_{\text{rot}} = 60$ V/cm. At this field amplitude the associated ponderomotive velocity is approximately $v_{\text{rot}} \sim 1200$ m/s, which is larger than the forward velocity of the molecular beam.

As a result, the ions must be given an additional impulse in the direction of the supersonic beam in order to “boost” them into the orbit of the rotating electric field. We accomplish this by applying an “anti-stop” electric field.

Although ions survive over a relatively broad range of electric field magnitudes, we observe




Figure 8.4: Number of trapped ions as a function of the magnitude of the applied anti-kick used to boost ions into the rotating-field trajectory when $\mathcal{E}_{\text{rot}} = 60$ V/cm. The broad peak indicates a range of kick values for which ions can be successfully captured.

that in some of these field magnitudes, the imperfect velocity cancellation leads to residual motion of the ions inside the trap. In particular, the ions can exhibit a sloshing motion corresponding to oscillation of the center-of-mass position within the trapping potential.

To minimize this residual motion, several parameters must be optimized simultaneously: the starting phase of $\mathcal{E}_{\text{rot}}$, the stopping voltage, and the DC position shimming fields in the x , y , and z directions.




Figure 8.5: Observation of ion sloshing in the y direction as a function of trap hold time. The vertical axis corresponds to the mean pixel position of the ion image on the camera, while the horizontal axis represents the time that the ions are held in the trap. The two panels correspond to rotating-field phases offset by $\pm 2^\circ$ from the optimal value.

As we have seen, the starting phase of $\mathcal{E}_{\text{rot}}$ is a critical parameter for successful ion loading. At $\mathcal{E}_{\text{rot}} = 60 \text{ V/cm}$ and $f_{\text{rot}} = 300 \text{ kHz}$, the corresponding rotating-field velocity is $v_{\text{rot}} \approx 1.2 \times 10^3 \text{ m/s}$. A phase error of 2° therefore corresponds to a velocity mismatch of approximately $\Delta v \sim 40 \text{ m/s}$, which is larger than the nominal thermal velocity of the trapped ions. This sensitivity arises because trapping with $\mathcal{E}_{\text{rot}}$ requires the ions to be launched onto the rotating-field orbit at the correct phase of the electric field. If the ions are created at the wrong phase, their initial velocity does not match the instantaneous direction of the rotating-field trajectory. The resulting velocity mismatch appears as residual secular motion in the trap, which manifests as slosh of the ion cloud. Since the neutral beam propagates along the x -axis, a mismatch in the launch phase primarily introduces a transverse velocity component, which appears mostly in the y direction.

Figure 8.5 shows the resulting slosh motion in the y direction. In this plot, the vertical axis

corresponds to the ion position on the camera (in pixels), while the horizontal axis shows the time for which the ions are held in the trap. The optimal starting phase is around 335° . When the phase is shifted by $\pm 2^\circ$, the initial phase of the slosh oscillation flips by approximately π . This indicates that the direction of the residual velocity has reversed: in one case the ions are launched with a residual velocity in the $+\hat{y}$ direction, while in the other case the residual velocity points toward $-\hat{y}$.

At a rotation frequency of 300 kHz, a phase shift of 1° corresponds to roughly 9 ns in timing, while a 2° shift corresponds to about 19 ns. This implies that the timing of the rotating-field turn-on must be controlled at the level of $\lesssim 10$ ns in order to minimize residual velocities.

In practice, timing jitter at this level can naturally arise in the experimental control chain. The experiment begins when the ARTIQ control system receives a trigger pulse from the DDG (Digital Delay Generator), which continuously triggers the YAG lasers at 10 Hz. ARTIQ then sends a TTL signal to our homemade DDS system (Direct Digital Synthesis) that generates the rotating electric field (see K. B. Ng's thesis for details of the control system [38]). Jitter at the level of ~ 10 ns within this chain of signals is therefore plausible. In addition, the ionization pulse laser itself has a pulse width of about 10 ns, which effectively introduces an additional timing uncertainty in the moment at which the ions are created.

Given these unavoidable timing uncertainties, we optimize the trap parameters to minimize the residual velocities in all three spatial directions. Figure 8.6 compares the ion dynamics after optimizing the loading procedure with and without $\mathcal{E}_{\text{rot}}$. Within each panel, the first column shows the mean ion position in the x , y , and z directions. The second column shows the spatial widths in x , y , and z .

We observe that the slosh behavior in the x and z directions is comparable between the two cases, indicating that no significant additional residual velocity is introduced along these axes when $\mathcal{E}_{\text{rot}}$ is present. In contrast, while a clear slosh oscillation is not observed in the y direction with $\mathcal{E}_{\text{rot}}$ on, the spatial width of the ion cloud along y is noticeably larger. This behavior is consistent with a distribution of small residual transverse velocities arising from timing jitter at the ~ 10 ns level.




Figure 8.6: Comparison of ion dynamics after optimizing the loading procedure with and without the rotating electric field $\mathcal{E}_{rot}$. For each case, the first column shows the mean ion position in the x , y , and z directions. The second column shows the spatial widths of the ion cloud in each direction.

8.2.2 Ejecting ions

Once ions have been loaded onto the stable trajectory associated with $\mathcal{E}_{\text{rot}}$, we perform the Ramsey spectroscopy sequence. After the quantum evolution time, the molecules are dissociated into Th^+ kicked towards the detector. As discussed above, we wish to keep the rotating electric field on during this process. In the prototype geometry, this means that the ions are ejected from the octagon region where the rods are long and $\mathcal{E}_{\text{rot}}$ is very uniform into a region near the rod ends, where the field amplitude correspondingly decreases.

At first, one might expect that this simply “turns off” $\mathcal{E}_{\text{rot}}$, leaving the ions with a large residual velocity associated with the micromotion driven by the rotating field. In practice, however, this geometry provides the spatial ramp-down of the field. As the ions move through the rod-end region, the shrinking rotating-field orbit can be described by an effective ponderomotive potential. If the field amplitude changes quasi-adiabatically compared with the rotation frequency, this ponderomotive potential is ramped smoothly to zero, and the micromotion energy is not converted into excess kinetic energy of the ions. The leading term in this potential produces an axial kick, while the next term, which is quadratic in the transverse radius, acts as a lens on the ion cloud.

To make this more quantitative, let z be the ejection direction, and let x and y the transverse directions. Near the trap axis, the rotating field may be written as :

$$\mathbf{E}_{\text{rot,eject}}(z, t) = E_0(z)\mathbf{e}(t) \quad \mathbf{e}(t) = [\cos(\omega_{\text{rot}}t + \phi)\hat{\mathbf{x}} + \sin(\omega_{\text{rot}}t + \phi)\hat{\mathbf{y}}], \quad (8.1)$$

where $E_0(z)$ is the spatial envelope of the rotating field along the ejection direction. The simulated profile of $E_0(z)$ in the current electrode geometry is shown in Fig. 8.7. The field is approximately constant in the central rod region and then falls rapidly near the rod ends.

To describe the rod-end region, it is useful to write the field in terms of a near-axis electrostatic potential. Let $\mathbf{r}_{\perp} = x\hat{\mathbf{x}} + y\hat{\mathbf{y}}$ and $r_{\perp}^2 = x^2 + y^2$. Near the trap axis, the rotating field is an instantaneously a dipole in the transverse direction ($m = 1$). For Laplace’s equation in cylindrical coordinates, the regular $m = 1$ solution is proportional to $-(\mathbf{e} \cdot \mathbf{r}_{\perp})E_0$ when there is no axial




Figure 8.7: SIMION simulation of the rotating-electric-field amplitude along the ejection direction. The horizontal axis is referenced to the octagon center. The dashed and dash-dotted vertical lines indicate the axial locations of the ring and rod ends, respectively. The field is approximately uniform in the central rod region and then decreases rapidly beyond the rod ends.

variation on the field. Once the field envelope is allowed to vary along z (i.e. $E_0(z)$), the regular radial solution becomes an order-1 Bessel (or modified-Bessel) function [105]. Then, for small $r_\perp$, we can expand the potential to the lowest order:

$$\Phi(\mathbf{r}, t) = -(\mathbf{e} \cdot \mathbf{r}_\perp)[E_0(z) + a(z)r_\perp^2 + \dots].$$

Requiring $\nabla^2 \Phi = 0$, we get $a(z) = -E_0''(z)/8$, and therefore we have:

$$\Phi(\mathbf{r}, t) \approx -E_0(z)\mathbf{e} \cdot \mathbf{r}_\perp + \frac{E_0''(z)}{8}(\mathbf{e} \cdot \mathbf{r}_\perp)r_\perp^2, \quad (8.2)$$

where

$$E_0''(z) \equiv \frac{d^2}{dz^2}E_0(z),$$

is the second derivative of E_0 with respect to z .

From this potential, we see that near the rod ends, the field is no longer uniform across the transverse plane. Averaging over the fast rotating motion gives the effective ponderomotive potential

$$U_{\text{eff}}(r_\perp, z) = \frac{q^2 E_0^2(z)}{2m\omega_{\text{rot}}^2} + \frac{q^2}{4m\omega_{\text{rot}}^2} \left[(E_0')^2 - E_0(z)E_0''(z) \right] r_\perp^2 + \dots, \quad (8.3)$$

where q and m are the ion charge and mass, and $E_0'(z)$ is the first derivative of $E_0(z)$ with respect to z .

The first term in Eq. (8.3) depends only on z and therefore acts as an axial potential. As the ions move from the high-field region into the rod-end region where $E_0 \rightarrow 0$, this term also decreases, and the ions experience an average force in the ejection direction:

$$F_z = -\frac{dU_{\text{eff}}}{dz} \approx -\frac{q^2}{m\omega_{\text{rot}}^2} E_x(z)E_0'(z).$$

In other words, part of the micromotion energy associated with $\mathcal{E}_{\text{rot}}$ is converted into axial kinetic energy. Since the ions go from a region of $E_{0,\text{in}}$ to $E_{0,\text{out}} = 0$, the corresponding kinetic-energy gain is approximately

$$\Delta KE \approx \frac{q^2}{2m\omega_{\text{rot}}^2} E_{0,\text{in}}^2.$$

The second term in Eq. (8.3) is quadratic in the transverse radius and therefore acts as a lens. The lens behavior is determined by the spatial derivatives of the field envelope

$$(E'_0(z))^2 - E_0(z)E''_0(z).$$

If this quantity is positive, the effective potential increases away from the axis and results in focusing.

Writing the quadratic term as

$$U_{\text{lens}}(r_{\perp}, z) = \frac{q^2}{4m\omega_{\text{rot}}^2} [(E'_0(z))^2 - E_0(z)E''_0(z)] r_{\perp}^2,$$

the transverse force is

$$\mathbf{F}_{\perp} = -\nabla_{\perp} U_{\text{lens}}.$$

We can solve for the equation of motion

$$m\ddot{y} = -\frac{q^2}{2m\omega_{\text{rot}}^2} [(E'_0(z))^2 - E_0(z)E''_0(z)] y.$$

If the ions traverse the rod-end region with approximately constant axial speed v_z , then using $d/dt = v_z d/dz$ gives

$$\frac{d^2 y}{dz^2} \approx -\frac{q^2}{2m^2\omega_{\text{rot}}^2 v_z^2} [(E'_0(z))^2 - E_0(z)E''_0(z)] y.$$

In the thin-lens approximation, where the transverse position does not change substantially while the ions cross the lens region, one can set $y \approx y_0$ in the integral. The resulting change in the transverse angle is

$$\Delta y' = -\frac{q^2}{2m^2\omega_{\text{rot}}^2 v_z^2} \left[\int [(E'_0(z))^2 - E_0(z)E''_0(z)] dz \right] y_0. \quad (8.4)$$

Using the thin-lens definition, $\Delta y' = -y_0/f$, the focal strength is then given by

$$\frac{1}{f} = \frac{1}{v_z^2} \frac{q^2}{2m^2\omega_{\text{rot}}^2} \int [(E'_0(z))^2 - E_0(z)E''_0(z)] dz. \quad (8.5)$$

Therefore, the sign of the lens is determined by the integral in this expression. For the field profile shown in Fig. 8.7, a rough estimate gives a positive integral, so the rod-end taper acts like a focusing lens.

In fact, we can evaluate the sign analytically. Note that

$$\frac{d}{dz}(E_0 E'_0) = (E'_0)^2 + E_0 E''_0.$$


Using this, the integral becomes

$$\begin{aligned} & \int_{z_i}^{z_f} \left[(E'_0(z))^2 - \left(\frac{d}{dz} E_0(z) E'_0(z) - (E'_0(z))^2 \right) \right] \\ &= 2 \int_{z_i}^{z_f} (E'_0(z))^2 dz - E_0(z) E'_0(z) \Big|_{z_i}^{z_f} \end{aligned}$$


The second term, $E_0(z)E'_0(z)$ evaluated at the starting and ending z positions will be approximately zero, as long as ions always begin from a region with uniform field to a region with zero field. Therefore, this term will always act as a focusing lens. In order to reduce the lensing, we can reduce $E'_0(z)$, in other words, a gradual spatial taper of $E_0(z)$. Additionally, the lensing is reduced for larger v_z and higher ω_{rot} .

There is another very important detail about ejecting ions with the rotating electric field on. When we first attempted to eject ions with $\mathcal{E}_{\text{rot}}$ on, we were uncertain whether the field at the rod ends would be sufficiently adiabatic to reduce the rotating-field micromotion without leaving a large residual transverse velocity. Initially, no ions were observed at the detector, and we suspected that the spatial ramp might not be adiabatic enough. However, when I revisited the experiment a couple years later, the problem was traced to an error in the DDS control code rather than to the rod-end ramp itself. In particular, if $\mathcal{E}_{\text{rot}}$ was left on indefinitely after the sequence, the experiment failed for reasons unrelated to the ion dynamics during ejection. When instead $\mathcal{E}_{\text{rot}}$ was turned off about 10 ms after the kickout, ions were observed at the detector. Since the total time of flight is only about 180 μs , this indicates that keeping $\mathcal{E}_{\text{rot}}$ on during the actual ejection and flight to the detector is not itself problematic.

To confirm this interpretation, I performed a simple numerical simulation of a single ion in the measured field envelope. After the kickout, the ions have about $v_z \sim 4500$ m/s. Using this velocity, the spatial profile in Fig. 8.7 may be converted into a time-dependent envelope for the rotating field seen by the ions during ejection.



(a)



(b)

Figure 8.8: Adiabaticity check for ion ejection with $\mathcal{E}_{\text{rot}}$ left on. (a) Simulated transverse components of the rotating electric field seen by an ion during ejection, obtained by converting the measured spatial envelope into a time-dependent envelope using the measured kickout velocity. The field amplitude is approximately constant before the kickout and then decreases to zero over several cycles after the ion reaches the rod-end region. (b) Simulated transverse position and velocity of a single ion in this time-dependent field. The transverse velocities in both x and y decrease to nearly zero by the end of the ramp, while the z components are not included in this simplified calculation. This indicates that the rod-end falloff is sufficiently adiabatic to remove the rotating-field micromotion without leaving a large residual transverse velocity.

At $f_{\text{rot}} = 150$ kHz, the field amplitude decreases from 60 V/cm to nearly zero over roughly four rotation cycles. Figure 8.8a shows the corresponding time-dependent field components and the simulated transverse trajectory of a single ion. In this simulation, I do not propagate the axial motion itself. Instead, I impose the time-dependent field envelope corresponding to an ion moving through the rod-end region at the measured kickout velocity. Within this simplified model, the transverse velocities in both x and y decrease to nearly zero by the end of the ramp, indicating that the rod-end falloff is adiabatic enough to revive the rotating-field micromotion without leaving a large residual transverse velocity.

The time-of-flight and the ion cloud widths also agree qualitatively with the picture developed above. Figure 8.9 shows the effect of leaving $\mathcal{E}_{\text{rot}}$ on during kickout while scanning its magnitude. For the orange trace, no $\mathcal{E}_{\text{rot}}$ is applied during kickout, so any variation as the horizontal axis is scanned is attributed to shot-to-shot fluctuations. For the blue trace, the magnitude of $\mathcal{E}_{\text{rot}}$ during kickout is scanned and shown on the horizontal axis. As the rotating field is increased, the ion time of flight shifts earlier by about $10 \mu\text{s}$, while the TOF width remains approximately constant. In the plot, the TOF axis is shown with a $145 \mu\text{s}$ offset removed. This behavior is consistent with a largely common-mode axial acceleration of the ion packet, as expected from the z -dependent term in the effective ponderomotive potential discussed above.

At the same time, the transverse widths in both x and y decrease with increasing $\mathcal{E}_{\text{rot}}$, consistent with focusing by the rod-end lens. However, the cloud centroids in both transverse directions also shift, although by an amount smaller than the cloud width. A purely centered lens would change the cloud size without shifting its center of mass, so the centroid motion indicates that the kickout optics also contain a transverse steering components. Furthermore, the amount of focusing is different for x and y , suggesting that the kickout optics are not perfectly symmetric. This asymmetry could arise from a misalignment between the ion cloud and the trap axis, imperfect rod-end geometry, or residual transverse DC fields in the kickout ion optics. In practice, however, this steering is not very large and can be compensated by optimizing the kickout parameters for a chosen value of $\mathcal{E}_{\text{rot}}$ so that the ions still arrive at near the center of the detector.




Figure 8.9: Effect of leaving $\mathcal{E}_{\text{rot}}$ on during kickout while scanning its magnitude. The left column shows the cloud centroid position in x , y , and the time-of-flight coordinate z , together with the total ion number inferred from the area under the TOF curve. The right column shows the corresponding cloud widths in x , y , and z , together with the mean image intensity. For the orange trace, no rotating electric field is applied during kickout; for the blue trace, the magnitude of $\mathcal{E}_{\text{rot}}$ during kickout is scanned. A time offset of $145 \mu\text{s}$ was subtracted from the time-of-flight coordinate z . As the kickout rotating field is increased, the TOF shifts earlier and the transverse cloud widths decrease, consistent with an axial acceleration together with transverse focusing during ejection. The simultaneous shift of the cloud centroids indicates an additional “steering” component in the extraction optics; this could be either from the imperfect adiabatic ramping of the field, or from an inherent offset in the cloud, which is enlarged by this focusing process.

8.3 “Batch” mode operation in the prototype experiment

In the current prototype experiment we are not yet configured to implement the full continuous throughput operation. The primary limitation is that the apparatus contains only a single location where both optical pumping and state detection are performed (see Figure 2.9). As a result, ions cannot simultaneously undergo state preparation in one region while another packet is being detected elsewhere.

However, the prototype apparatus can still emulate several key aspects of the batch-mode operation. In this approach, we sequentially create and trap ion packets, allow them to undergo Ramsey evolution, and then detect them using the same preparation and detection region. While this does not reproduce the spatially separated preparation and detection envisioned for the throughput operation, it allows us to test the timing infrastructure required for operating with multiple ion packets.

Implementing this sequence requires careful coordination of the timing between ion production, trapping fields, Ramsey pulses, and detection. In particular, the pulsed lasers and rod drivers must all operate with well-defined, synchronized timings.

In practice, achieving reliable operation requires maintaining strict “timing hygiene” throughout the experiment. The relative timing between ion creation, the application of electric fields, and optical pulses (especially the mF pumping laser) used for state manipulation must remain stable at the level of $\lesssim 10$ ns. Small timing errors can translate into phase errors in the rotating electric field or imperfect velocity cancellation during ion loading, both of which degrade the trapping efficiency and coherence of the ions.

Although these constraints arise here in the context of the prototype batch-mode experiment, they are also directly relevant for the final throughput-mode apparatus. In that configuration, multiple ion packets will be simultaneously present in different regions of the trap, and precise timing control will likewise be required to ensure that each packet experiences the correct sequence of fields and pulses during its evolution.

8.3.1 Timings

If we want to operate in batch mode (or eventually in throughput mode), we need to carefully manage the experimental timing. The relevant constraint is not simply that the ionization and dissociation pulses occur at the right absolute times, but they occur at the right phases of both the rotating electric field and the rf trapping voltage. In particular, we would like to satisfy the following conditions :

- (1) At the time of ionization, the rotating electric field $\mathcal{E}_{\text{rot}}$ must have the appropriate phase, $\phi_{\text{rot}}^{\text{ionization}}$, so that the ions are launched onto the rotation micromotion.
- (2) At the time of ionization, the rf trapping voltage V_{rf} must have the appropriate phase, $\phi_{\text{rf}}^{\text{ionization}}$, ideally near either a maximum or a minimum, so that the micromotion velocity from the rf trapping potential is minimized.
- (3) At the time of dissociation, the rotating electric field $\mathcal{E}_{\text{rot}}$ must have the appropriate phase, $\phi_{\text{rot}}^{\text{dissociation}}$, for REMPAD. In the prototype configuration, this corresponds to dissociating when the rotating electric field is along the $\pm\hat{y}$ direction.
- (4) At the time of dissociation, it is also desirable, though not strictly necessary, for the rf trapping voltage V_{rf} to have an appropriate phase, $\phi_{\text{rf}}^{\text{dissociation}}$, again near either a maximum or a minimum, in order to minimize the micromotion velocity at detection.
- (5) Finally, the ionization and dissociation laser pulses must occur at times that are properly synchronized with both ω_{rot} and ω_{rf} , so that additional ion packets can be created and detected while satisfying the phase conditions above.

These requirements are all part of the same phase-matching problem. The phases evolve as

$$\phi_i = \phi_{i,0} + 2\pi f_i t \pmod{2\pi}, \quad i \in \{\text{rot}, \text{rf}\}. \quad (8.6)$$

Therefore, once the DDS start phases and the two frequencies are chosen, the phase pair

$$(\phi_{\text{rot}}(t), \phi_{\text{rf}}(t))$$

evolves through a particular set of values. The DDS start phases are used to choose the desired phase pair at ionization.

$$\phi_{\text{rot}}^{\text{ionization}} = \phi_{\text{rot},0} + 2\pi f_{\text{rot}} t_{\text{ionization}}, \quad \phi_{\text{rf}}^{\text{ionization}} = \phi_{\text{rf},0} + 2\pi f_{\text{rf}} t_{\text{ionization}} \pmod{2\pi}, \quad (8.7)$$

In other words, we can choose the starting phases of DDS outputs such that we have $\phi_{\text{rot},0}$ and $\phi_{\text{rf},0}$ at the moment of ionization.

However, the phase pair at dissociation is determined by the trap time

$$\tau = t_{\text{dissociation}} - t_{\text{ionization}}. \quad (8.8)$$

Then, the phases are

$$\phi_{\text{rot}}^{\text{dissociation}} = \phi_{\text{rot}}^{\text{ionization}} + 2\pi f_{\text{rot}} \tau, \quad \phi_{\text{rf}}^{\text{dissociation}} = \phi_{\text{rf}}^{\text{ionization}} + 2\pi f_{\text{rf}} \tau \pmod{2\pi}, \quad (8.9)$$

This means that the dissociation phases cannot, in general, be chosen independently. In other words, varying the trap time τ does not produce arbitrary phase pairs, but instead traces out a one-dimensional trajectory in the two-dimensional phase space $(\phi_{\text{rot}}, \phi_{\text{rf}})$.

In a perfect world, none of these requirements would present any particular difficulty. In practice, however, the exact values of f_{rot} and f_{rf} are constrained by the frequency resolution of the DDS system. Our current DDS chips operate with a clock frequency $f_{\text{clk}} = 100$ MHz and a 32-bit tuning word, so the frequency resolution is

$$\Delta f = \frac{f_{\text{clk}}}{2^{32}} = 0.023283064365 \text{ Hz}.$$

The output frequency of the DDS is determined by the tuning word K according to

$$f = \frac{K}{2^{32}} f_{\text{clk}}.$$

Therefore, when we request nominal frequencies of $f_{\text{rot}} = 300$ kHz and $f_{\text{rf}} = 100$ kHz, the DDS realizes the nearest frequencies allowed by this discretization. For the rotating field,

$$K_{\text{rot}}^{\text{near}} = \text{round}\left(\frac{300000}{\Delta f}\right) = 12884902,$$

$$f_{\text{rot}} = \frac{K_{\text{rot}}^{\text{near}}}{2^{32}} f_{\text{clk}} = 300000.0026077032 \text{ Hz},$$

and similarly for the rf field,

$$K_{\text{rf}}^{\text{near}} = \text{round}\left(\frac{100000}{\Delta f}\right) = 4294967,$$

$$f_{\text{rf}} = \frac{K_{\text{rf}}^{\text{near}}}{2^{32}} f_{\text{clk}} = 99999.9931082130 \text{ Hz}.$$

This frequency digitization immediately leads to two important consequences: first, it determines the recurrence time over which the phases of $\mathcal{E}_{\text{rot}}$ and V_{rf} return to the same values; second, it affects how well these phases can be synchronized with the externally triggered laser pulses.

Since both frequencies are generated by the same DDS system, their ratio is exactly

$$\frac{f_{\text{rot}}}{f_{\text{rf}}} = \frac{K_{\text{rot}}}{K_{\text{rf}}} = \frac{p}{q}, \quad (8.10)$$

where p/q is the ratio reduced to lowest terms, and p and q are integer numbers. Then, a particular phase pair $(\phi_{\text{rot}}, \phi_{\text{rf}})$ then repeats after the common recurrence time

$$T_{\text{recurrence}} = \min\{\tau > 0 : f_{\text{rot}}\tau \in \mathbb{Z}, f_{\text{rf}}\tau \in \mathbb{Z}\}$$

$$= \frac{2^{32}}{f_{\text{clk}} \text{gcd}(K_{\text{rot}}, K_{\text{rf}})}, \quad (8.11)$$

where gcd denotes the greatest common divisor of the two DDS tuning words. Equivalently, during one recurrence time, the rotating field completes p cycles and the rf field completes q cycles. If a desired dissociation phase pair is reached at some trap time τ_0 , then other trap times that reach the same phase pair are

$$\tau = \tau_0 + mT_{\text{recurrence}}, \quad m \in \mathbb{Z}. \quad (8.12)$$

Therefore, $T_{\text{recurrence}}$ determines both the packet-to-packet recurrence time that makes sure that the ion loading conditions are the same and also the equivalent trapping times that make sure that the dissociation conditions are the same.

As an example, for the nearest-to-target tuning words above, $\text{gcd}(K_{\text{rot}}, K_{\text{rf}}) = 1$, which gives

$$T_{\text{recurrence}} = 42.94967296 \text{ s}.$$

This means that if we insist on exactly the same values of both $\phi_{\text{rot}}^{\text{ionization}}$ and $\phi_{\text{rf}}^{\text{ionization}}$, then the next identical loading condition occurs only about 43 s later. In practice, there may be some

tolerance in the allowed phase errors $\Delta\phi_{\text{rot}}^{\text{ionization}}$ and $\Delta\phi_{\text{rf}}^{\text{ionization}}$, but this result already illustrates that the naive nearest-frequency choice is not ideal for batch-mode operation.

There is a separate but related constraint that comes from the external laser trigger. These lasers are triggered by a DG645 digital delay generator, which provides a frequency resolution of 1 μHz . Nominally, the ionization laser runs at 10 Hz. If the DDG repetition rate is f_{DDG} , then the phase slip from shot to shot is

$$\Delta\phi_i^{\text{shot}} = 2\pi \left[\frac{f_i}{f_{\text{DDG}}} - \text{round} \left(\frac{f_i}{f_{\text{DDG}}} \right) \right], \quad i \in \{\text{rot}, \text{rf}\}. \quad (8.13)$$

For instance, the number of rotating-field cycles between two adjacent ionization pulses is

$$N_{\text{rot}} = \frac{f_{\text{rot}}}{10 \text{ Hz}} = 30000.0002607703.$$

Equivalently, the phase of $\mathcal{E}_{\text{rot}}$ slips from shot to shot by

$$\Delta\phi_{\text{rot}}^{\text{shot}} = 2\pi \times 0.0002607703 = 0.093877^\circ.$$

Although this phase slip is small on a single shot, it accumulates over repeated cycles, reaching nearly 2° after 20 shots, or about 2 s at a 10 Hz repetition rate. As discussed in the previous section, a phase error at this level can already produce substantial residual velocity during ion loading. For stable operation, this shot-to-shot phase slip must remain small over the full experimental sequence.

Therefore, there are three practical frequency-design goals.¹ First, $T_{\text{recurrence}}$ should be short enough so that multiple ion packets can be produced under equivalent loading conditions and also the trapping time resolution is short enough. Second, the phase trajectory in the 2d-phase space of $(\phi_{\text{rot}}, \phi_{\text{rf}})$ should go through the desired dissociation phase pair at useful trap times. Third, the frequencies should be compatible with the available DDG timing resolution, so that the laser pulses remain synchronized with both $\mathcal{E}_{\text{rot}}$ and V_{rf} .²

¹ These considerations are not fully independent, since they all arise from the same underlying phase-matching problem, but it is still useful to separate them conceptually.

² At the time of writing, we are developing a new DDS system. In this new system, f_{DDG} will also be produced using the same DDS boards that generate the electric fields. Therefore, this will not be as much of an issue.

8.3.1.1 Integer ratio between f_{rot} and f_{rf}

The timing problem becomes significantly simpler if the frequencies of the rotating electric field and the rf trapping field are chosen such that their ratio is an integer,

$$f_{\text{rot}} = r f_{\text{rf}},$$

where r is an integer. For the nominal experimental parameters, the natural choice is $r = 3$, corresponding to $f_{\text{rot}} \approx 300$ kHz and $f_{\text{rf}} \approx 100$ kHz. If the DDS tuning words are chosen such that

$$K_{\text{rot}} = 3K_{\text{rf}}.$$

In this case the ratio of these frequencies is exactly 3, even though the individual values may differ slightly from their nominal targets due to the finite frequency resolution of the DDS³.

With this choice, the common recurrence time becomes

$$T_{\text{common}} = \frac{1}{f_{\text{rf}}} = 10.000000689178753 \text{ } \mu\text{s} \approx 10 \text{ } \mu\text{s}.$$

This means that the pair of phases $(\phi_{\text{rot}}, \phi_{\text{rf}})$ repeats every (approximately) $10 \text{ } \mu\text{s}$ or every rf period. Consequently, we can make the next ion packet anytime we want, in increments of T_{common} .

In addition, synchronization with the ionization lasers becomes straightforward. If the repetition period of the DDG trigger is chosen to be an integer multiple of T_{common} ,

$$T_{\text{DDG}} = M T_{\text{common}},$$

then every ionization pulse will occur at exactly the same values of ϕ_{rot} and ϕ_{rf} . For example, choosing $M = 10^4$ gives

$$T_{\text{DDG}} \approx 0.1 \text{ s},$$

corresponding to a repetition rate very close to $9.999999 \text{ Hz} \approx 10 \text{ Hz}$.

³ Initially, we were concerned that choosing a commensurate ratio such as $f_{\text{rot}} : f_{\text{rf}} = 3 : 1$ might introduce unwanted resonances in the trap dynamics, for example by enhancing ion heating. In preliminary measurements we did not observe clear evidence of such an effect. Nevertheless, it will be important to further investigate potential systematic effects associated with commensurate rotating electric fields and rf trapping voltages.

The advantage of this configuration is that the timing becomes very simple: the packet-to-packet recurrence time is short, and the ionization laser can be synchronized by choosing the DDG period to be an integer multiple of $T_{\text{recurrence}}$. However, this choice has an important consequence: the phase evolution of the two fields is *no longer* independent.

Since $f_{\text{rot}} = 3f_{\text{rf}}$, the phase evolution of the two fields is strongly constrained and follows

$$\Delta\phi_{\text{rot}} = 3\Delta\phi_{\text{rf}} \pmod{2\pi}.$$

As a result, once the timing is adjusted so that $\mathcal{E}_{\text{rot}}$ reaches a desired phase at dissociation, the V_{rf} phase is automatically determined. In other words, the phases $\phi_{\text{rot}}^{\text{dissociation}}$ and $\phi_{\text{rf}}^{\text{dissociation}}$ cannot be chosen independently.

In practice, this may be acceptable. The REMPAD process requires a well-defined rotating-field phase, but the rf phase at dissociation may be less critical. As the ions travel toward the detector, the physical rods end before the detection region, which naturally reduces the excess micromotion that can come when the rf phase is not correct at dissociation. Therefore, choosing an integer ratio between f_{rot} and f_{rf} remains an attractive option because it greatly simplifies the timing synchronization while still allowing precise control of the rotating-field phase at dissociation.

8.3.1.2 Non-integer ratio between f_{rot} and f_{rf}

The more general strategy is to choose frequencies that are close to the nominal targets but do not have an exact integer ratio. In this case, the timing is less constrained than in the integer-ratio case, but the recurrence time can become much longer if the DDS tuning words do not share a large common divisor.

In practice, the situation can be improved by carefully selecting the DDS tuning words. Rather than simply choosing the frequencies that are closest to the nominal values of 300 kHz and 100 kHz, we can search for nearby tuning words that share a reasonably large common divisor while still avoiding an exact integer ratio between the two frequencies. We want frequencies that are close to a simple rational ratio, but not so close that the two phases become rigidly locked. Increasing

$\gcd(K_{\text{rot}}, K_{\text{rf}})$ reduces $T_{\text{recurrence}}$ and therefore improves both the packet-to-packet recurrence time and the spacing between equivalent trap times.

As an example, we can choose DDS tuning words of $K_{\text{rot}} = 12879748$ and $K_{\text{rf}} = 4337058$. These tuning words give

$$f_{\text{rot}} = 299880.0016939640 \text{ Hz}, \quad f_{\text{rf}} = 100980.0005704164 \text{ Hz},$$

which are both within 2 kHz of the nominal target frequencies. For this choice,

$$T_{\text{recurrence}} = 326.797383774900 \text{ } \mu\text{s}.$$

If the DDG repetition period is chosen to be $306 T_{\text{recurrence}}$, then the corresponding trigger rate would be

$$f_{\text{DDG}} = 10.0000000564881 \text{ Hz},$$

and both the rotating-field phase and the rf phase repeat exactly from shot to shot. This is a substantial improvement over the ~ 43 s recurrence time obtained from the nearest-to-target tuning words.

As before, we would adjust DDS starting phases such that the desired pair $(\phi_{\text{rot}}^{\text{ionization}}, \phi_{\text{rf}}^{\text{ionization}})$ is realized exactly at the moment of ionization. Even in the non-integer ratio case, we still do not have arbitrary control over $\phi_{\text{rot}}^{\text{dissociation}}$ and $\phi_{\text{rf}}^{\text{dissociation}}$. Once the frequencies and ionization phases are fixed, the dissociation phase pair is determined entirely by the trap time τ .

As previously mentioned, if there exists a trap time τ_0 for which the desired dissociation phase pair is reached, how often this point repeats in time is determined by $T_{\text{recurrence}}$ of Eq. (8.11). The code to find $T_{\text{recurrence}}$ is shown in Appendix C.

Finally, one must also choose the DDG trigger rate⁴ so that the phase slip from shot to shot is acceptable. If the DDG repetition rate is f_{DDG} , then the phase slips per shot are

$$\Delta\phi_{\text{rot}}^{\text{shot}} = 2\pi \left(\frac{f_{\text{rot}}}{f_{\text{DDG}}} - \text{round}\left(\frac{f_{\text{rot}}}{f_{\text{DDG}}}\right) \right),$$

⁴ Again, in the new DDS system, this will be less of a problem as long as the laser trigger rate is commensurate with both the f_{rot} and f_{rf} . For more detail, see Appendix C.

$$\Delta\phi_{\text{rf}}^{\text{shot}} = 2\pi \left(\frac{f_{\text{rf}}}{f_{\text{DDG}}} - \text{round}\left(\frac{f_{\text{rf}}}{f_{\text{DDG}}}\right) \right).$$

Using the above numbers for f_{rot} and f_{rf} and the fact that our DDG has 1 μHz resolution such that $f_{\text{DDG}} = 10.000000 \text{ Hz}$, the phase slips per shot comes out to be

$$\Delta\phi_{\text{rot}}^{\text{shot}} \approx 0.06098^\circ \quad \Delta\phi_{\text{rf}}^{\text{shot}} \approx 0.02054^\circ.$$

This corresponds to ϕ_{rot} slip of about 6° in 100 shots or 10 seconds. In the batch mode, this may be fine because we will only take about 2-3 seconds to prepare all the ion packets. However, this will be a problem for the throughput mode operation. Therefore, in the non-integer-ratio case, frequency optimization plays three roles. First, the frequencies should be chosen so that the packet-to-packet recurrence time is short enough for practical operation. Second, the resulting phase trajectory should allow the desired dissociation phase pair to be reached at useful trap times. Third, the frequencies should be compatible with the available DDG timing resolution so that the shot-to-shot phase slip remains acceptably small.

In summary, when f_{rot} and f_{rf} do not have a simple integer ratio, the initial ionization phases can still be chosen arbitrarily through the DDS start phases, while the dissociation phases are controlled by the trap time τ . The frequencies must then be optimized so that the recurrence time is short enough for practical use, the desired dissociation phase pair is accessible, and the phase evolution remains compatible with the laser trigger timing.

Chapter 9

Generation III ThF⁺ eEDM signal

9.1 Experimental procedure

9.1.1 Experimental block and switch states

During the edm measurement campaign, we record our data in “blocks”. A block is a self-contained set of Ramsey measurements taken under a specific experimental configurations, or *switches*. Each switch $\tilde{S}$ corresponds to an experimental parameter whose sign can be reserved between each measurement. Gen. II had four main switches per block: $\tilde{B}$, $\tilde{R}$, $\tilde{I}$, and $\tilde{D}$. During a block, we do not apply any feedback (except for the laser frequencies, which are always being monitored and corrected). Generally, we design a block with appropriate switches such that it produces an estimate of d_e and of various imperfections in experimental condition that are sufficiently accurate and precise to be useful for meaningful feedback/feedforward corrections that we apply before the next block. When we design a protocol for a block, there are three main key points to keep in mind:

- (1) **Evenly weighted configurations.** Within each block, we want each switch configuration to contribute with the same (or as nearly as possible) statistical weight. In practice, this means we want comparable numbers of shots or signal (and also comparable SNR) in each switch configuration; this is because we want the linear combinations that we form at the end the block to represent the true symmetry projections, rather than being biased by uneven sampling.

- (2) **Differences between switch states should reflect the intended physics.** The flipping of a switch state $\tilde{S}$ in a block should change the measurement *only* through physical effects that are genuinely $\tilde{S}$ -odd (or $\tilde{S}$ -even). Therefore, ideally, we design the sequence of the switch states that is i) fast compared to relevant drifts and ii) cleanly implemented, with minimal collateral changes.
- (3) **Ramsey fringes near the extrema and zero crossings.** In a Ramsey measurement, the detected asymmetry can be written in a sinusoidal function of the accumulated phase

$$A(\tau) \propto C \cos(\phi_R) = C \cos(2\pi f\tau + \phi) \quad (9.1)$$

so the sensitivity to a small frequency shift depends on where you sit on the fringe. Near a zero crossing (i.e. $\phi_R \sim \pi/2$), the slope $|dA/df|$ is maximal, so any small changes in f produce the largest measurable change in the signal—this is optimal for statistical precision on f , and therefore for d_e . Near a top or bottom of a fringe (i.e. $\phi_R \sim 0, \pi$), the slope is near zero, and thus it's less sensitive to small frequency shifts. These points serve to diagnose and decouple the amplitude/contrast offsets and any other non-phase related systematics from the phase. Therefore, we include Ramsey times that sample both: near the zero crossings for maximum EDM sensitivity and also near top/bottom to constrain other parameters (i.e. contrast, offset, etc.). Having these points make the phase extraction more robust against drifts and contrast changes because the data contain enough degrees of freedom to separate “phase-like” changes from “amplitude/offset”-like changes.

In the Gen. II implementation, a standard block was constructed from $2^3 = 8$ experimental permutations of the three switches $\tilde{B}$, $\tilde{R}$, and $\tilde{I}$. The fourth switch $\tilde{D}$ was automatically applied because the two Stark doublets are always detected simultaneously in each shot. Then, the Ramsey times for each switch state were adjusted (to correct for any slow drift in the system) independently based on the previous block data to make sure that the time points are as close as possible to both the sides of the fringes, where the frequency sensitivity is largest, and close to the beat condition¹

, where common-mode noise rejection is best and several systematics shifts are minimized.

For Gen. III, the experimental protocols will be *different*. For instance, because we have fewer ions per shot² and longer coherence time compared to Gen. II, it will take longer to determine where the late-time zero-crossings occur for a fringe. We still want to measure all the switch combinations close together in time to suppress long-term drift between each block, we may only have $\tilde{B}$ and $\tilde{D}$ in a block, with additional switches introduced only when combining several blocks into a larger “superblock”. The final protocol should therefore be chosen only after the dominant drift timescales in the apparatus have been characterized.

9.1.2 Introduction to switch-state and parity-channel notation

The different experimental switches help us understand and suppress systematics. Therefore, we need a notation to discuss all the different switch state and the experimental configurations in which the data are taken. In the next section, I introduce a compact notation that will be used by Gen. III.

To analyze block data, it is useful to distinguish between two types of quantities: 1) the raw values measured in individual experimental configurations and 2) the linear combinations formed from them. I will refer to the former as *switch-state* values and the latter as *parity-channel* values.

Here, I assume that a block consists of the eight combinations of the three switches $\tilde{B}$, $\tilde{D}$, and $\tilde{R}$ switches. The extension to N switches is immediate: every factor of $1/8$ below is replaced by 2^N , and additional switch labels are appended in the obvious way.

A particular experimental configuration is specified by the ordered switch-state label

$$\mathbf{s} = (\tilde{B}, \tilde{D}, \tilde{R}), \quad \tilde{B}, \tilde{D}, \tilde{R} \in \{\pm 1\}. \quad (9.2)$$

Then, let $V(\tilde{B}, \tilde{D}, \tilde{R})$ denote the value of some measured or inferred quantity in that switch state.

¹ Here, the beat condition refers to when the two doublets are in phase.

² or at least so far in the prototype

$$V(\tilde{B}, \tilde{D}, \tilde{R}) = V^{\mathbf{s}}, \quad \mathbf{s} \in \left\{ \begin{array}{l} (+1 +1 +1) \\ (-1 +1 +1) \\ (+1 -1 +1) \\ (-1 -1 +1) \\ (+1 +1 -1) \\ (-1 +1 -1) \\ (+1 -1 -1) \\ (-1 -1 -1) \end{array} \right\}. \quad (9.3)$$

In practice, V may represent the signed Ramsey frequency, or any of the fit parameters appearing in Eq. (9.1), or for an experimental parameter such as a stray transverse magnetic field component B_x .

It is first helpful to recall the single-switch decomposition. For a quantity V and a single switch $\tilde{S}$, we define the even and odd parts as³

$$V^{\mathcal{S}} = \frac{1}{2} [V(\tilde{S} = +1) + V(\tilde{S} = -1)], \quad V^S = \frac{1}{2} [V(\tilde{S} = +1) - V(\tilde{S} = -1)]. \quad (9.4)$$

An *unslashed* switch label denotes the part of the signal that is odd under that switch $\tilde{S} \rightarrow -\tilde{S}$, whereas the *slashed* switch label denotes the part that is even under that switch. These expressions are just parity projections of the raw switch-state values. Simple averaging isolates the part of the quantity V that is even under the switch $\tilde{S}$, while multiplying by the switch sign before averaging isolates the odd part.

³ Note that the equations here can equivalently be written as

$$V^{\mathcal{S}} = \frac{1}{2} [V^{\tilde{S}=+1} + V^{\tilde{S}=-1}], \quad V^S = \frac{1}{2} [V^{\tilde{S}=+1} - V^{\tilde{S}=-1}],$$

using the definition of $\mathbf{s}$ shown in Eq (9.3). However, here, I use the form of Eq (9.4) to be extra explicit and also because for a single switch, the switch-state label $\mathbf{s}$ is either V^{+1} or V^{-1} , which seems like an exponent.

In the case of three switches, there are a total of eight parity channels $V^{\mathbf{p}}$:

$$V^{\mathbf{p}}, \quad \mathbf{p} \in \left\{ \begin{array}{l} (\not{B}\not{D}\not{R}) \\ (B\not{D}\not{R}) \\ (\not{B}D\not{R}) \\ (BD\not{R}) \\ (\not{B}\not{D}R) \\ (B\not{D}R) \\ (\not{B}DR) \\ (BDR) \end{array} \right\}. \quad (9.5)$$

For the three-switch case, the same logic of averaging applies. For instance, simple averaging isolates the part of the quantity that is even under now all the switches:

$$V^{\not{B}\not{D}\not{R}} = \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} V(\tilde{B}, \tilde{D}, \tilde{R}). \quad (9.6)$$

On the other hand, multiplying by the switch sign before averaging isolates the odd part:

$$V^{B\not{D}\not{R}} = \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{D} V(\tilde{B}, \tilde{D}, \tilde{R}). \quad (9.7)$$

A quantity that is odd under a combination of switches is obtained by:

$$V^{BD\not{R}} = \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{B}\tilde{D} V(\tilde{B}, \tilde{D}, \tilde{R}) \quad (9.8)$$

For brevity, I generally omit the slashed (even) switches when writing parity-channel values, whenever the meaning is clear from context. In this abbreviated notation, the channel is even under all switches is denoted by a superscript zero:

$$V^{B\not{D}\not{R}} \equiv V^B, \quad V^{\not{B}\not{D}\not{R}} \equiv V^0. \quad (9.9)$$

This compact notion was first introduced by the ACME group. Although this shorthand is convenient, sometimes it is hard to see how the quantity is formed. For example, the symbol f^B may denote the Ramsey frequency that is fully projected B -odd channel, obtained by averaging over

all other switch combinations. It may also refer more loosely to a difference formed from only two measurements with $\tilde{B} = \pm 1$, while $\tilde{D}$ and $\tilde{R}$ are held fixed. In practice, the intended meaning is usually clear from context. When that distinction matters, however, we use the more explicit notation $f^{B\tilde{D}\tilde{R}}$ to indicate that the quantity is odd under $\tilde{B}$ and averaged over $\tilde{D}$ and $\tilde{R}$.

More generally, this notation can be also used to indicate the value of a switch that is held fixed while parity combinations are formed over the remaining switches. For example, suppose we take a block of data in which all eight combinations of $\tilde{B}$, $\tilde{D}$, and $\tilde{R}$ are sampled, while another switch $\tilde{P}$ is held fixed at $\tilde{P} = -1$. From that block, we can then form parity-channel values such as

$$V^{\tilde{B}\tilde{D}\tilde{R};\tilde{P}=-1}.$$

Later, suppose we take a second block under the same conditions, but with $\tilde{P} = +1$. Then, we can combine the two blocks⁴ to form a single “superblock”. This superblock then would contain 16 distinct switch states and can be projected onto 16 parity channels, including channels involving P , such as

$$f^{\tilde{B}\tilde{D}\tilde{R}P}.$$

In matrix form, the transformation between the switch-state $\mathbf{s}$ and the parity-state $\mathbf{p}$ can be written in the form of $\{V^{\mathbf{p}}\} = T\{V^{\mathbf{s}}\}$ as:

$$\begin{pmatrix} V^0 \\ V^B \\ V^D \\ V^{BD} \\ V^R \\ V^{BR} \\ V^{DR} \\ V^{BDR} \end{pmatrix} = \frac{1}{8} \begin{pmatrix} + & + & + & + & + & + & + & + \\ + & - & + & - & + & - & + & - \\ + & + & - & - & + & + & - & - \\ + & - & - & + & + & - & - & + \\ + & + & + & + & - & - & - & - \\ + & - & + & - & - & + & - & + \\ + & + & - & - & - & - & + & + \\ + & - & - & + & - & + & + & - \end{pmatrix} \begin{pmatrix} V_{\text{sgn}}^{+++} \\ V_{\text{sgn}}^{--+} \\ V_{\text{sgn}}^{+-+} \\ V_{\text{sgn}}^{---} \\ V_{\text{sgn}}^{++-} \\ V_{\text{sgn}}^{--+-} \\ V_{\text{sgn}}^{+--+} \\ V_{\text{sgn}}^{----} \end{pmatrix}. \quad (9.10)$$

⁴ given that they have comparable precision so that they carry roughly equal statistical weight

The numerical entries in this matrix depend on the ordering convention chosen for the switch states and parity channels, but the logic is always the same. If we expand to include not just three but N switches, the value of $1/8$ in the matrix is replaced by $1/2^N$.

In the matrix in Eq (9.10), I write the measured quantity as V_{sgn} to emphasize that the analysis is carried out using the *signed* value of V , rather than its magnitude alone.

This convention is one of the main differences between the earlier generations of the experiments and Gen. III. In Gen. I and Gen. II, the extracted Ramsey frequencies were treated as absolute quantities rather than signed phase-evolution frequencies. For instance, when the magnetic field direction was reversed ($\tilde{B} \rightarrow -\tilde{B}$), the spin precession direction was also reversed, but the reported frequency did not. This usage of absolute-value frequencies breaks the intuition that “averaging over $\tilde{B}$ ” results in the B -even part of the signal; the dominant Zeeman contribution is intrinsically B -odd in the underlying Hamiltonian but becomes “ f^0 ”, which appears to be B -even, because we used $|f(\tilde{B} = +1)|$ and $|f(\tilde{B} = -1)|$ as our measured frequencies. As a result, the book-keeping of channel labels was effectively shifted. What was treated as the “baseline” frequency was actually $\tilde{B}$ -odd, and the superscripts on small shifts must have been interpreted *relative* to that $\tilde{B}$ -odd baseline. In other words, because of the usage of absolute frequencies, the largest term f_0^0 (in Gen. II’s language) is $\tilde{B}$ -odd, and f^{DB} is $\tilde{D}$ -odd but $\tilde{B}$ -even. This confused me so much for so many years. For Gen. III, we plan to avoid this ambiguity and confusion by using a signed quantity for the measured frequency as in Eq (9.10); we will define the phase-evolution frequency with a consistent sign convention across $\tilde{B}$. The same rule for switch-state superscripts applies to ζ and Δ , which will be discussed in the next section.

We can also invert the transformation in Eq.(9.10)

$$\{V^{\mathbf{s}}\} = T^{-1}\{V^{\mathbf{p}}\},$$

where T^{-1} is the inverse of the transformation matrix T . This inverse relation reconstructs the value of V in a particular switch state from the corresponding parity-channel values. In fact, the inverse matrix has the same pattern of signs as T , but without the overall factor of $1/8$.

A simple way to understand the inverse expansion is to begin with a simple case of a single switch $\tilde{S} = \pm 1$. Using the definitions of the even and odd components from Eq. (9.4), we can take a sum and a difference of $V^{\$}$ and V^S as:

$$V^{\$} + V^S = V(\tilde{S} = +1) \quad V^{\$} - V^S = V(\tilde{S} = -1).$$

Then it is immediate to see that the two relations can be combined into the expression

$$V(\tilde{S}) = V^{\$} + \tilde{S}V^S. \quad (9.11)$$

In other words, the switch-state value is reconstructed by adding the even component and the odd component, weighted by the sign of the switch.

In the three-switch case, the expansion follows the same logic and becomes

$$\begin{aligned} V(\tilde{B}, \tilde{D}, \tilde{R}) = & V^0 + \tilde{B}V^B + \tilde{D}V^D + \tilde{B}\tilde{D}V^{BD} \\ & + \tilde{R}V^R + \tilde{B}\tilde{R}V^{BR} + \tilde{D}\tilde{R}V^{DR} + \tilde{B}\tilde{D}\tilde{R}V^{BDR}, \end{aligned} \quad (9.12)$$

where again $V^0 = V^{\$ \# \# \#}$.

Although this result can be obtained directly from the inverse matrix, the expansion is more transparent and obvious. For instance, the switch-state value $V(+1, -1, +1)$ corresponds to

$$V(+1, -1, +1) = V^0 + V^B - V^D - V^{BD} + V^R + V^{BR} - V^{DR} - V^{BDR}.$$

9.1.3 Switch-parity expansion

We can model the evolution within either Stark doublet as an effective two-level system. Introducing notation that differs slightly from previous generations, we parameterize the effective Hamiltonian for each doublet as

$$H_{\text{eff}} = \frac{h}{2} \begin{pmatrix} \zeta & \Delta \\ \Delta & -\zeta \end{pmatrix}. \quad (9.13)$$

The diagonal term ζ contains all contributions that shift the relative energy of the $m_F = \pm 3/2$ states, while the off-diagonal term Δ contains all contributions that mix them. In the Gen. II

notation the diagonal term was denoted by f_0 ; here I instead reserve the symbol f for the measured frequency and use ζ for the diagonal energy shift.⁵

Following Ref. [51], we decompose each matrix element into a dominant modeled baseline plus residual perturbations:

$$\zeta(\tilde{B}, \tilde{D}, \tilde{R}) = \tilde{B} \zeta_0^B + \delta\zeta(\tilde{B}, \tilde{D}, \tilde{R}), \quad (9.14)$$

$$\Delta(\tilde{B}, \tilde{D}, \tilde{R}) = \tilde{R}(\Delta_0^R + \tilde{D}\Delta_0^{DR}) + \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R}), \quad (9.15)$$

where $\tilde{B}, \tilde{D}, \tilde{R} \in \{\pm 1\}$ label the experimental switch states. Throughout this section we assume

$$|\delta\zeta| \ll |\zeta_0^B|, \quad |\Delta_0^R|, |\Delta_0^{DR}|, |\delta\Delta| \ll \zeta_0^B, \quad |\delta\Delta| \ll |\Delta_0^R|, |\Delta_0^{DR}|$$

In the Gen. II notation of Ref. [51], the quantities ζ_0^B , $\delta\zeta$, Δ_0^R , Δ_0^{DR} , and $\delta\Delta$ were denoted by f_0^0 , $\delta\mathcal{F}_0$, Δ_0 , Δ^D , and $\delta\mathcal{D}$ respectively.

In this new notation, the subscript $_0$ denotes the dominant modeled baseline amplitudes. The superscript denotes the switch-parity channel of the full Hamiltonian contribution. Therefore,

$$\tilde{B}\zeta_0^B, \quad \tilde{R}\Delta_0^R, \quad \tilde{D}\tilde{R}\Delta_0^{DR}$$

represent the B -odd, R -odd, and DR -odd baseline contributions, respectively.

We use the convention that ζ_0^B is defined to be a positive baseline amplitude,

$$\zeta_0^B > 0. \quad (9.16)$$

The sign convention used will be discussed more in the later section. With this convention, note that ζ_0^B itself does not change sign when the magnetic-field switch is reversed; the sign reversal is carried by the explicit factor of $\tilde{B}$ in Eq. (9.14).

The off-diagonal baseline amplitudes Δ_0^B and Δ_0^{DR} are defined in a similar fashion; they may be positive or negative numbers, but they do not themselves change sign under the rotation switch. The R parity of the corresponding Hamiltonian contributions is carried by the explicit factor of $\tilde{R}$ in Eq. (9.15).

⁵ The other graduate student in Gen. III Anzhou Wang and I collectively decided that we will use the symbol f to refer to only the **measured** frequency and use ζ for any diagonal energy shift in the effective Hamiltonian. This removes an ambiguity present in earlier notations and avoids introducing additional subscripts and superscripts.

For instance, using the experimental parameters of $f_{\text{rot}} = 300$ kHz and $\mathcal{E}_{\text{rot}} = 60$ V/cm, these baseline terms are approximately

$$\begin{aligned}\zeta_0^B &\approx 100 \text{ Hz}, \\ \Delta_0^R &\approx 8 \text{ Hz}, \\ \Delta_0^{DR} &\approx -3.8 \text{ Hz}.\end{aligned}\tag{9.17}$$

The off-diagonal couplings arise at fourth order in perturbation theory from the combined effects of rotation and Ω -doubling. Following Ref. [51], the corresponding baseline amplitudes are

$$\begin{aligned}h\Delta_0^R &= \frac{3\hbar\omega_{ef}}{2} \left(\frac{\hbar\omega_{\text{rot}}}{d_{\text{mf}}\mathcal{E}_{\text{rot}}} \right)^3 \left(\frac{18A_{||}^2 - 19d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2}{A_{||}^2 - d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2} \right), \\ h\Delta_0^{DR} &= \frac{3\hbar\omega_{ef}}{2} \left(\frac{\hbar^3\omega_{\text{rot}}^3}{d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2 A_{||}} \right) \left(\frac{9A_{||}^2 - 8d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2}{A_{||}^2 - d_{\text{mf}}^2\mathcal{E}_{\text{rot}}^2} \right).\end{aligned}\tag{9.18}$$

In writing these definitions, I take ω_{rot} to be the positive rotation frequency. It is important to remember that the dependence on the rotation-direction switch is *not* included in the amplitudes themselves.⁶

In Eqs. (9.14) and (9.15), the quantities

$$\delta\zeta(\tilde{B}, \tilde{D}, \tilde{R}), \quad \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R}),$$

represent any residual perturbations that arise from experimental imperfections in the system. Both quantities follow the same switch-parity notation introduced in the previous section. For instance, if there is any residual diagonal shift from some imperfection that is B odd, this will show up in the channel

$$\delta\zeta^B \quad (\text{or, equivalently } \delta\zeta^{B\mathbb{D}\mathbb{R}}).$$

If there is any residual perturbative off-diagonal shift that is odd in DR , it will show up in the channel

$$\delta\Delta^{DR} \quad (\text{or, equivalently } \delta\Delta^{\mathbb{B}DR})$$

⁶ Again, the switch dependence is carried explicitly by the factors of $\tilde{R}$ and $\tilde{D}$ in Eq. (9.15). Therefore, Δ_0^R and Δ_0^{DR} are switch-independent baseline amplitudes, i.e. do not themselves change the signs based on the switch state. The full baseline contributions

$$\tilde{R}\Delta_0^R, \quad \tilde{R}\tilde{D}\Delta_0^{DR}$$

are R -odd and DR -odd, respectively. The additional D parity of the second contribution is encoded by the explicit factor of $\tilde{D}$ and arises from the doublet-dependent structure entering the expression.

Comparing coefficients in the switch-parity expansion then gives ⁷

$$\zeta^B = \zeta_0^B + \delta\zeta^B.$$

Similarly, the full DR -odd off-diagonal component becomes

$$\Delta^{DR} = \Delta_0^{DR} + \delta\Delta^{DR}.$$

In a perfect model, one would have $\delta\zeta(\tilde{B}, \tilde{D}, \tilde{R}) = \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R}) = 0$. However, in practice $\delta\zeta^S$ and $\delta\Delta^S$ parameterize deviations from the modeled baseline due to imperfections and uncontrolled couplings. The summary of the change of notation in the definitions of Eqs (9.14) and (9.15) is given in Table 9.1.

Quantity type	Gen. II notation	Gen. III notation
Dominant diagonal baseline term	f_0^0	ζ_0^B
Residual diagonal term with old superscript S	δf_0^S	$\delta\zeta^{BS}$
Dominant off-diagonal baseline term, $\tilde{R}$ -odd, $\tilde{D}$ -even	Δ^0	Δ_0^R
Dominant off-diagonal baseline term, $\tilde{R}$ -odd, $\tilde{D}$ -odd	Δ^D	Δ_0^{DR}
Residual off-diagonal term with old superscript S	$\delta\Delta^S$	$\delta\Delta^{SR}$

Table 9.1: Dictionary between the Gen. II old convention and the Gen. III new notation used in this chapter. In the Gen. III notation, superscripts denote the total parity-state of the switches for that quantity. For diagonal terms, the dominant baseline contribution is B -odd; therefore, the old superscript S maps to the new parity-state label BS . For off-diagonal terms, the dominant baseline contribution is R -odd, so an old superscript S maps to the new parity-state label RS .

With these definitions, we introduce the signed measured frequency

$$f(\tilde{B}, \tilde{D}, \tilde{R}) \equiv \tilde{B} f_{\text{abs}}(\tilde{B}, \tilde{D}, \tilde{R}) = \tilde{B} \sqrt{\zeta(\tilde{B}, \tilde{D}, \tilde{R})^2 + \Delta(\tilde{B}, \tilde{D}, \tilde{R})^2}. \quad (9.19)$$

This convention differs from the Gen. II choice, in which the measured frequency was taken to be positive by definition. Here the explicit factor of $\tilde{B}$ removes the special role of the $\tilde{B}$ switch in

⁷ The form here is already very intuitive. However, mathematically, we can do the expansion using Eq. (9.12) very explicitly as the following:

$$\begin{aligned} & (\zeta^0 + \tilde{B}\zeta^B + \tilde{D}\zeta^D + \tilde{B}\tilde{D}\zeta^{BD} + \tilde{R}\zeta^R + \tilde{B}\tilde{R}\zeta^{BR} + \tilde{D}\tilde{R}\zeta^{DR} + \tilde{B}\tilde{D}\tilde{R}\zeta^{BDR}) \\ & = \tilde{B}\zeta_0^B + (\delta\zeta^0 + \tilde{B}\delta\zeta^B + \tilde{D}\delta\zeta^D + \tilde{B}\tilde{D}\delta\zeta^{BD} + \tilde{R}\delta\zeta^R + \tilde{B}\tilde{R}\delta\zeta^{BR} + \tilde{D}\tilde{R}\delta\zeta^{DR} + \tilde{B}\tilde{D}\tilde{R}\delta\zeta^{BDR}) \end{aligned}$$

Then, we look for terms with $\tilde{B}$ from both sides and equate then we get $\zeta^B = \zeta_0^B + \delta\zeta^B$. Note that the explicit sum over all the parity-state with that particular switch multiplied, as shown in Eq. (9.7), establishes exactly this.

the channel decomposition: the dominant Zeeman term appears in the B -odd channel f^B , and a perturbation that is odd under $\tilde{D}$ appears directly in the D -odd channel f^D .

We expand Eq. (9.19) in powers of the small perturbations to obtain (with the convention $\zeta_0^B > 0$),

$$\begin{aligned}
f(\tilde{B}, \tilde{D}, \tilde{R}) &\approx \tilde{B} \zeta_0^B + \delta\zeta + \frac{\tilde{B} \Delta^2}{2\zeta_0^B} - \frac{\delta\zeta \Delta^2}{2(\zeta_0^B)^2} - \frac{\tilde{B} \Delta^4}{8(\zeta_0^B)^3} + \dots \\
&\approx \tilde{B} \zeta_0^B + \delta\zeta \\
&\quad + \frac{\tilde{B}}{2\zeta_0^B} \left[(\Delta_0^R)^2 + (\Delta_0^{DR})^2 + 2\tilde{D} \Delta_0^R \Delta_0^{RD} + 2\tilde{R}(\Delta_0^R + \tilde{D} \Delta_0^{RD}) \delta\Delta + (\delta\Delta)^2 \right] \\
&\quad - \frac{\delta\zeta}{2|\zeta_0^B|^2} \left[(\Delta_0^R)^2 + (\Delta_0^{RD})^2 + 2\tilde{D} \Delta_0^R \Delta_0^{RD} + 2\tilde{R}(\Delta_0^R + \tilde{D} \Delta_0^{RD}) \delta\Delta + (\delta\Delta)^2 \right] \\
&\quad + \dots,
\end{aligned} \tag{9.20}$$

where $\delta\zeta = \delta\zeta(\tilde{B}, \tilde{D}, \tilde{R})$, $\Delta = \Delta(\tilde{B}, \tilde{D}, \tilde{R})$, and $\delta\Delta = \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R})$. Because ζ_0^B has been defined as a positive amplitude, there is no separate $\text{sgn}(\zeta_0^B)$ appearing in this expansion. Note with the Gen. III signed-frequency convention of Eq. (9.19), the dominant sign of the measured frequency is set by $\tilde{B}$.

Equation (9.20) may now be projected into the various parity channels. For example, the D -odd channel is

$$f^D = \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{D} f(\tilde{B}, \tilde{D}, \tilde{R}). \tag{9.21}$$

Applying this projector to Eq. (9.20), only terms whose parity-state is D survive the summation, giving

$$\begin{aligned}
f^D &= \delta\zeta^D + \frac{\Delta_0^R \delta\Delta^{BDR} + \Delta_0^{RD} \delta\Delta^{BR}}{\zeta_0^B} - \delta\zeta^D \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} - \delta\zeta^0 \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} \\
&\quad + \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{D} \left[\frac{\tilde{B} (\delta\Delta)^2}{2\zeta_0^B} - \frac{\delta\zeta}{2(\zeta_0^B)^2} \left(2\tilde{R}(\Delta_0^R + \tilde{D} \Delta_0^{RD}) \delta\Delta + (\delta\Delta)^2 \right) \right] \\
&\quad + \dots,
\end{aligned} \tag{9.22}$$

where I have again used the shorthand notation for $\delta\zeta = \delta\zeta(\tilde{B}, \tilde{D}, \tilde{R})$ and $\delta\Delta = \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R})$. In the Gen. III notation, the eEDM-sensitive channel is f^D (which used to be f^{DB} in Gen. II notation),

and the $\delta\zeta^D$ term includes the edm-induced energy shift. Since the dominant baseline term $\tilde{B}\zeta_0^B$ cancels in the summation, the leading nonzero contributions to f^D are already perturbative.

The appearance of $\delta\Delta^{BDR}$ and $\delta\Delta^{BR}$ in Eq. (9.22) is purely a consequence of the Gen. III new notation convention. In the Gen. II notation, superscripts on the Hamiltonian coefficients were “net” switch labels defined relative to the dominant switch parity of each matrix element, rather than the literal parity of the full term. For the diagonal element, the dominant term was $\tilde{B}f_0^0$, so an old diagonal superscript S corresponded to literal parity channel BS . For the off-diagonal element, the dominant term was $\tilde{R}\Delta^0$, so an old off-diagonal superscript S corresponded to literal parity channel RS . In particular, the Gen. II baseline coefficients Δ^0 and Δ^D did not themselves denote the literal parities of the full terms: $\tilde{R}\Delta^0$ is R -odd, while $\tilde{R}\tilde{D}\Delta^D$ is RD -odd. In Gen. III, we therefore write these as $\tilde{R}\Delta_0^R$ and $\tilde{R}\tilde{D}\Delta_0^{RD}$, reserving superscripts exclusively for the actual parity-state of these quantities and subscript to denote that these quantities are the modeled baseline contributions. With this convention, the Gen. II quantities δf_0^{DB} , $\delta\Delta^{DB}$, and $\delta\Delta^B$ correspond in the Gen. III literal notation to $\delta\zeta^D$, $\delta\Delta^{BDR}$, and $\delta\Delta^{BR}$, respectively.

Although the coefficients appearing in the projected channel expansion no longer carry explicit switch factors, their channel assignment can be checked by tracing each term back to its parent term in the full frequency expansion. For example, the contribution $\Delta_0^R \delta\Delta^{BDR}/\zeta_0^B$ originates from the term $(\tilde{B}\tilde{R}/\zeta_0^B)\Delta_0^R \delta\Delta$ in Eq. (9.20). Then, for that contribution only we have

$$\begin{aligned}
(\tilde{B}\tilde{R}/\zeta_0^B)\Delta_0^R \delta\Delta &= \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{D} \left(\frac{\tilde{B}\tilde{R}}{\zeta_0^B} \Delta_0^R \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R}) \right) \\
&= \frac{\Delta_0^R}{8\zeta_0^B} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{B}\tilde{D}\tilde{R} \left(\delta\Delta^0 + \tilde{B}\delta\Delta^B + \tilde{D}\delta\Delta^D + \tilde{R}\delta\Delta^R \right. \\
&\quad \left. + \tilde{B}\tilde{D}\delta\Delta^{BD} + \tilde{B}\tilde{R}\delta\Delta^{BR} + \tilde{D}\tilde{R}\delta\Delta^{DR} + \tilde{B}\tilde{D}\tilde{R}\delta\Delta^{BDR} \right). \\
&= \frac{\Delta_0^R}{8\zeta_0^B} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \left(\tilde{B}\tilde{D}\tilde{R}\delta\Delta^0 + \tilde{D}\tilde{R}\delta\Delta^B + \tilde{B}\tilde{R}\delta\Delta^D + \tilde{B}\tilde{D}\delta\Delta^R \right. \\
&\quad \left. + \tilde{R}\delta\Delta^{BD} + \tilde{D}\delta\Delta^{BR} + \tilde{B}\delta\Delta^{DR} + \delta\Delta^{BDR} \right) \\
&= \frac{\Delta_0^R}{\zeta_0^B} \delta\Delta^{BDR}.
\end{aligned}$$

where the only surviving term is $\delta\Delta^{BDR}$.

Similarly, we can project Eq.(9.20) into the B -odd channel, giving us

$$\begin{aligned}
f^B &= \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{B} f(\tilde{B}, \tilde{D}, \tilde{R}) \\
&= \zeta_0^B + \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2\zeta_0^B} \\
&\quad + \left[\delta\zeta^B + \frac{\Delta_0^R \delta\Delta^R + \Delta_0^{RD} \delta\Delta^{DR}}{\zeta_0^B} - \delta\zeta^B \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} - \delta\zeta^{BD} \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} \right. \\
&\quad \left. + \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{B} \left[\frac{\tilde{B}(\delta\Delta)^2}{2\zeta_0^B} - \frac{\delta\zeta}{2(\zeta_0^B)^2} \left(2\tilde{R}(\Delta_0^R + \tilde{D} \Delta_0^{RD})\delta\Delta + (\delta\Delta)^2 \right) \right] \right] \\
&\quad + \dots,
\end{aligned} \tag{9.23}$$

where $\delta\zeta = \delta\zeta(\tilde{B}, \tilde{D}, \tilde{R})$ and $\delta\Delta = \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R})$.

In this new notation, the B -odd channel f^B contains the dominant Zeeman transition frequency. Its leading term is the baseline Zeeman splitting ζ_0^B , followed by the dressed-state correction from the baseline off-diagonal mixing. The leading perturbative diagonal contribution is $\delta\zeta^B$, which represents any residual B -odd diagonal shift not captured by the baseline coefficient ζ_0^B .⁸

Unlike f^B , the fully even channel f^0 (or equivalently, $f^{\cancel{B}\cancel{D}\cancel{R}}$) has no zeroth-order baseline contribution, since the dominant Zeemna term $\tilde{B}\zeta_0^B$ cancels in the channel average.

$$\begin{aligned}
f^0 &= \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} f(\tilde{B}, \tilde{D}, \tilde{R}) \\
&= \left[\delta\zeta^0 + \frac{\Delta_0^R \delta\Delta^{BR} + \Delta_0^{RD} \delta\Delta^{BDR}}{\zeta_0^B} - \delta\zeta^0 \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} - \delta\zeta^D \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} \right. \\
&\quad \left. + \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \left[\frac{\tilde{B}(\delta\Delta)^2}{2\zeta_0^B} - \frac{\delta\zeta}{2(\zeta_0^B)^2} \left(2\tilde{R}(\Delta_0^R + \tilde{D} \Delta_0^{RD})\delta\Delta + (\delta\Delta)^2 \right) \right] \right] \\
&\quad + \dots,
\end{aligned} \tag{9.24}$$

where $\delta\zeta = \delta\zeta(\tilde{B}, \tilde{D}, \tilde{R})$ and $\delta\Delta = \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R})$.

⁸ To reiterate, in the new Gen. III notation, baseline quantities are denoted by the subscript ₀. Any residual perturbations are written with δ with the superscript indicating their switch parity.

In f^0 , the leading diagonal term is $\delta\zeta^0$, i.e. the part of the residual diagonal shift that is even under $\tilde{B}$, $\tilde{D}$, and $\tilde{R}$ per $\delta\zeta^0 = \delta\zeta^{\mathcal{B}\mathcal{D}\mathcal{R}}$. In practice, this term is a natural place for effects such as a non-reversing magnetic field or an imperfect B -chop, provided those effect are also even under $\tilde{D}$ and $\tilde{R}$. In that sense, the Gen. III channel f^0 now plays the role of the old Gen. II f^B channel. For instance, if the actual diagonal Zeeman term is

$$\zeta(\tilde{B}, \tilde{D}, \tilde{R}) = \tilde{B}\zeta_0^B + \delta\zeta^0,$$

then ζ_0^B represents the positive amplitude of the reversing magnetic field contribution, while $\delta\zeta^0$ represents the non-reversing magnetic field contribution. In terms of magnetic fields, this convention corresponds to⁹

$$\zeta_0^B = 3g_F\mu_B B_{\text{rev}}, \quad |\delta\zeta^0| = |3g_F\mu_B \delta B_{\text{nr}}|$$

for reversing $B_{\text{rev}} = \mathcal{B}_{\text{rot}}$ and non reversing B_{nr} magnetic fields.

For the R -odd channel, we have

$$\begin{aligned} f^R &= \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{R} f(\tilde{B}, \tilde{D}, \tilde{R}) \\ &= \left[\delta\zeta^R + \frac{\Delta_0^R \delta\Delta^B + \Delta_0^{RD} \delta\Delta^{BD}}{\zeta_0^B} - \delta\zeta^R \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} - \delta\zeta^{DR} \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} \right. \\ &\quad \left. + \frac{1}{8} \sum_{\tilde{B}, \tilde{D}, \tilde{R}=\pm 1} \tilde{R} \left[\frac{\tilde{B}(\delta\Delta)^2}{2\zeta_0^B} - \frac{\delta\zeta}{2(\zeta_0^B)^2} \left(2\tilde{R}(\Delta_0^R + \tilde{D}\Delta_0^{RD})\delta\Delta + (\delta\Delta)^2 \right) \right] \right] \\ &\quad + \dots \end{aligned} \tag{9.25}$$

with $\delta\zeta = \delta\zeta(\tilde{B}, \tilde{D}, \tilde{R})$ and $\delta\Delta = \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R})$. The R -odd channel f^R is perturbative at leading order. The first nonzero diagonal contribution is $\delta\zeta^R$, which contains residual R -odd diagonal shifts. In practice, this channel contains R -odd effects such as the charging-current systematic and Berry-phase-related shifts.

⁹ B_{nr} can be written as δB^0

9.1.4 Sign definitions for chops

Here, I define the sign conventions for the main experimental chops used in the Gen. III used in this chapter. For the magnetic field chop, $\tilde{B} = +1$ corresponds to the case in which the effective rotating magnetic bias is parallel to $\mathcal{E}_{\text{rot}}$. Equivalently, this is the sign for which the time-averaged bias $\tilde{B}B^{\text{eff}}$ points along $\hat{\mathcal{E}}_{\text{rot}}$.

For the doublet chop, I follow the Gen. II convention: $\tilde{D} = +1$ denotes the upper Stark doublet, while $\tilde{D} = -1$ denotes the lower Stark doublet.

For the rotation chop, I follow the laboratory-frame convention

$$\mathcal{E}_{\text{rot}} = \mathcal{E}_{\text{rot}} [\hat{\mathbf{x}} \cos(\omega_{\text{rot}} t) + \tilde{R} \hat{\mathbf{y}} \sin(\omega_{\text{rot}} t)].$$

With this definition, $\tilde{R} = +1$ corresponds to rotation from $+\hat{x}$ towards $+\hat{y}$, in other words, positive rotation about the $+\hat{z}$ axis.¹⁰

A summary of these sign conventions is given in Table 9.2.

Chop	Definition of the +1 state
$\tilde{B}$	effective rotating magnetic bias parallel to $\mathcal{E}_{\text{rot}}$
$\tilde{D}$	upper Stark doublet
$\tilde{R}$	$\mathcal{E}_{\text{rot}}$ rotates about the $+\hat{z}$ direction (towards South in the prototype)

Table 9.2: Sign conventions for the main experimental chops used in this chapter.

Additional chops will be defined later as needed. For example, the $\tilde{I}$ chop denotes which doublet lands on which side of the phosphor screen. The $\tilde{P}$ chop labels the polarization state set by the quarter-wave plate angle of the strobing laser which pumps the m_F states.

9.1.5 Quick note about the sign convention for the measured frequency

In Gen. III, the measured frequency is defined as a signed quantity, as in Eq. (9.19). The sign of the reported frequency is tied to the switch label $\tilde{B}$. Actually, once the measured frequency is treated as a signed quantity, the corresponding Ramsey phase accumulation in Eq. (9.1) must also carry that sign convention.

¹⁰ This points South, in the prototype apparatus.

We also need to choose a convention for the sign of ζ_0^B ; this depends on how we assign the sign of ζ_0^B to the two physical orientations of $\mathcal{B}_{\text{rot}}$. With the convention above, $\tilde{B} = +1$ corresponds to the case in which the effective magnetic field is parallel to the rotating electric field. At the time of writing, we do not have an experimental data on the sign of g_F for ThF^+ . For HfF^+ , $g_F < 0$, therefore here, I further assume that this is also true for ThF^+ . If we take

$$\Delta m_F = \left(-\frac{3}{2}\right) - \left(+\frac{3}{2}\right) = -3,$$

then the dominant Zeeman shift will be

$$\zeta_0^B = -3g_F\mu_B\mathcal{B}_{\text{rot}}/h > 0$$

for the $\tilde{B} = +1$ state.¹¹

No matter what convention we use for the measured frequency, determining the sign of d_e from a measured f^D requires a consistent set of sign conventions.

9.2 Brief notes on systematics

As we saw, with the Gen. III signed-frequency convention and new parity-state notation, several channel names differ from their Gen. II counterparts even though the underlying physical effects are the same. In Gen. II, channel labels were tied to the older “net-superscript” convention, whereas in Gen. III they follow the actual parity-state of the measured signed frequency. As a result, effects that previously appeared in channels such as f^{BR} or f^{DB} are relabeled in the new notation. For example, the charging-current systematic and the Berry’s-phase contribution, which were previously associated with the BR channel, now appear in the R channel; the eEDM signal moves from f^{DB} to f^D ; the dominant frequency f^0 becomes f^B ; and the old f^B channel becomes the switch-even channel f^0 , corresponding to a non-reversing magnetic field. Likewise, the transverse-magnetic-field effect at the second harmonic of $\mathcal{E}_{\text{rot}}$, which in the Gen. II convention contributed

¹¹ This should be regarded only as a convenient bookkeeping convention for the sign of the effective Zeeman splitting. In the actual Ramsey measurement, the relevant signal is constructed from combinations of the *in-in* and *anti-in* measurements, and these two configurations have different signs for Δm_F . Thus, the sign assignment should not be interpreted too literally as arising from a single fixed choice of Δm_F .

to f^B and f^{BR} , is relabeled as contributing to f^0 and f^R in Gen. III. Table 9.3 summarizes several common examples of this relabeling. In each case the underlying systematic or signal is unchanged; only the channel name shifts because the explicit $\tilde{B}$ factor is absorbed into the Gen. III signed frequency.

Physical quantity / effect	Gen. II channel	Gen. III channel
Dominant transition frequency	f^0	f^B
Non-reversing magnetic field	f^B	f^0
Charging-current effect	f^{BR}	f^R
Berry’s phase	f^{BR}, f^B	f^R, f^0
eEDM signal	f^{DB}	f^D
Transverse B with $2\omega_{\text{rot}}$	f^B, f^{BR}	f^0, f^R

Table 9.3: Translation of several commonly used channel labels from the Gen. II convention to Gen. III convention. The physical effects are unchanged; only the channel names differ because the signed-frequency definition moves the dominant $\tilde{B}$ factor from the channel label into the measured frequency itself.

Table 9.3 is useful as a dictionary between the Gen. II and Gen. III channel labels. However, although the Gen. II and Gen. III experiments share many common features, there are also important quantitative differences in their systematic errors due to differences in the underlying spectroscopic parameters. In particular, the off-diagonal mixing terms are larger in ThF^+ under the experimental parameters, so several familiar imperfections can produce larger systematic shifts in Gen. III than they did in earlier generations.

In this section, I highlight some systematic mechanisms that require additional care in Gen. III. The goal here is not to provide a fully exhaustive or rigorous list of all possible Gen. III systematics, but rather to motivate a somewhat different “attitude” toward both old and new systematics in light of the stronger state mixing and the signed-frequency notation introduced above.

9.2.1 Systematic consequences of enhanced off-diagonal mixing

The larger off-diagonal terms Δ_0^R and Δ_0^{RD} in ThF^+ result in the measured transition frequency being more strongly “dressed” compared to HfF^+ . More importantly, these off-diagonal mixing terms generate cross-terms between diagonal and off-diagonal effects in systematics. This

can be seen explicitly in the D -odd channel. Retaining the leading terms from Eq. (9.22), one finds

$$f^D \approx \left[\delta\zeta^D - \delta\zeta^D \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} - \delta\zeta^0 \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} + \frac{\Delta_0^R \delta\Delta^{BDR} + \Delta_0^{RD} \delta\Delta^{BR}}{\zeta_0^B} \right], \quad (9.26)$$

where the omitted terms are higher order in the residual perturbations.

Equation (9.26) shows that larger Δ_0^R and Δ_0^{RD} modify the systematics picture in three distinct ways. First, the usual diagonal D -odd magnetic response, represented by $\delta\zeta^D$, is renormalized by the factor

$$1 - \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2}.$$

Second, a residual switch-even diagonal perturbation $\delta\zeta^0$ —for example, a non-reversing magnetic field—can leak directly into the D -odd channel through the cross-term

$$-\delta\zeta^0 \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2}.$$

A non-reversing magnetic field is dangerous not only through any ordinary D -odd g -factor mismatch, but also because off-diagonal dressing converts part of the switch-even magnetic shift into a false f^D signal.

To make this more explicit, suppose the usual diagonal magnetic contribution is written as

$$\delta\zeta^D \Big|_{B_{nr}} = \left(\frac{\delta g_F}{g_F} \right) \delta\zeta^0.$$

Then Eq. (9.26) becomes

$$f^D \approx \delta\zeta^0 \left[\frac{\delta g_F}{g_F} - \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} - \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} \frac{\delta g_F}{g_F} \right] + \frac{\Delta_0^R \delta\Delta^{BDR} + \Delta_0^{RD} \delta\Delta^{BR}}{\zeta_0^B}. \quad (9.27)$$

This motivates the interpretation of the bracketed quantity as an effective magnetic sensitivity of the dressed states. In particular, even if the ordinary $\delta g_F/g_F$ contribution is tuned small, the term proportional to $\Delta_0^R \Delta_0^{RD}/(\zeta_0^B)^2$ can still produce a residual D -odd shift.

The effects mentioned above about the off-diagonal terms have been extensively studied during the Gen. II systematics studies (for instance, see Refs. [15, 51, 60]). However, these effects will become more critical in Gen. III because the terms Δ_0^R and Δ_0^{RD} are larger. This can be easily

seen in Eq (9.18): both coefficients scale linearly with the Ω -doubling frequency ω_{ef} , so a larger ω_{ef} leads directly to stronger rotation-induced mixing between the $m_F = \pm 3/2$ states.

With the larger off-diagonal terms, Gen. III has to be think carefully about systematics as these effect will play a more prominent role in determining how technical imperfections project into the measured frequency channels.

9.2.2 One subtle consequence of large off-diagonal mixing: Berry phase

In Gen. III, we have developed a method of applying localized $\pi/2$ pulses using RF drive and rings. One very useful consequence of that work is that additional RF drives may also be used to “shim” quantities such as $\delta g_F/g_F$ or the baseline mixing terms Δ_0^R and Δ_0^{RD} . This project was led primarily by another graduate student Anzhou Wang, and a more detailed discussion is given in his thesis.

During the development of these techniques, Anzhou Wang and I spent a great deal of what our advisor calls “unstructured playtime” where we discussed possible consequences of such shimming. In particular, it seems that with an additional RF drive, it might be possible to tune the effective magnetic sensitivity, represented by the bracketed quantity in Eq. (9.27) to zero.

With the magnetic sensitivity shimmed to zero, one might naively think that the Berry-phase contribution to be removed as well. The reason is not that Berry phase is itself a magnetic effect, but rather that it can produce a shift in a nonscience channel such as f^0 and f^R , which then enters the science channel through the familiar leakage pathway

$$f^D \sim \left(\frac{\delta g_F}{g_F} \right) f^0.$$

From this point of view, tuning the magnetic leakage coefficient $\delta g_F/g_F$ to zero would seem to suppress any contribution to f^D arising from a Berry-induced shift in f^0 .

This conclusion, however, is not generally correct once the eigenstates are appreciably mixed by the off-diagonal coupling Δ . In the pure $|m_F| = 3/2$ limit, the dressed-state correction is zero, but when Δ is sufficiently large, the eigenstates are no longer pure basis sates. Then, their σ_z

expectation value is reduced to

$$\langle \sigma_z \rangle = \frac{\zeta_0^B}{\sqrt{(\zeta_0^B)^2 + \Delta^2}} \approx \left(1 - \frac{\Delta^2}{2(\zeta_0^B)^2} \right).$$

It is therefore natural to define an effective Berry coefficient for the dressed state,

$$g_{\text{Berry,eff}} \approx g_{\text{Berry}} \left[1 - \frac{(\Delta_0^R)^2 + (\Delta_0^{RD})^2}{2(\zeta_0^B)^2} - \tilde{D} \frac{\Delta_0^R \Delta_0^{RD}}{(\zeta_0^B)^2} \right], \quad (9.28)$$

where g_{Berry} is the Berry coefficient seen by the pure $|m_F| = 3/2$ states, and $g_{\text{Berry,eff}}$ is the corresponding coefficient seen by the dressed eigenstates.

This expression contains both $\tilde{D}$ -even and $\tilde{D}$ -odd components. In other words, because the two doublets experience different baseline mixing, the Berry response is not the same for $\tilde{D} = +1$ and $\tilde{D} = -1$. The resulting D -odd fractional mismatch is

$$\delta g_{\text{Berry}} / g_{\text{Berry}} \approx -\Delta_0^R \Delta_0^{RD} / (\zeta_0^B)^2.$$

Thus even if the ordinary magnetic sensitivity is tuned small using an additional RF drive, a residual Berry-phase contribution can remain.

This term looks a lot like the term that shows up in the “ordinary” effective magnetic sensitivity in Eq (9.27). The similarity is not accidental: both arise because the dressed eigenstates have reduced σ_z character once the states are mixed by Δ . However, the two quantities describe different physical responses. The ordinary effective g factor in Eq (9.27) characterizes the response of the science channel to a real residual magnetic field, whereas g_{Berry} characterizes the response to a Berry-phase-induced frequency shift. Although both are dressed by the same off-diagonal mixing, they do not track one another. The magnetic response contains an intrinsic mismatch $\delta g_F / g_F$, while the Berry mismatch arises primarily from the differential dressing of the two doublets. Because of this, tuning the magnetic response of f^D to zero does not generally eliminate the Berry-phase response.

I emphasize this point here because it illustrates a broader lesson about the relationship between Gen. II and Gen. III. The two generations share many experimental techniques, but subtle

quantitative differences can alter the relative importance of particular systematics. In Gen. II, the Berry phase was treated as affecting both doublets identically to a very good approximation. That approximation was reasonable because the corresponding mixing terms were much smaller. Using the Gen. II values, $\Delta_0^R = 0.6$ Hz and $\Delta_0^{RD} = -0.9$ Hz, and taking $\zeta_0^B \sim 100$ Hz, one finds

$$\frac{\delta g_{\text{Berry}}}{g_{\text{Berry}}} \approx -\frac{(0.6)(-0.9)}{100^2} = 5.4 \times 10^{-5}.$$

Compared to the ordinary $\delta g_F/g_F \sim 2 \times 10^{-3}$, this is a small effect.

However, for Gen. III, the corresponding parameters are much larger. Taking $\Delta_0^R = 8$ Hz and $\Delta_0^{RD} = -3.8$ Hz, and again $\zeta_0^B \sim 100$ Hz gives

$$\frac{\delta g_{\text{Berry}}}{g_{\text{Berry}}} \approx -\frac{(8)(-3.8)}{100^2} = 3.04 \times 10^{-3}.$$

This is now comparable to $\delta g_F/g_F \sim 2 \times 10^{-3}$. Therefore, if Gen. III operates in a regime with $\zeta_0^B \sim 100$ Hz, the Berry phase should no longer be treated purely in the $|m_F = 3/2|$ picture, but it should be analyzed in the dressed-state basis.¹²

Considering Berry phase in the dressed-state picture is only one consequence of having larger baseline mixing terms Δ_0^R and Δ_0^{RD} . Larger mixing terms are also expected to increase the importance of magnetic systematics, as Trevor Wright explored for Gen. III in Ref. [60]. More generally, the enhanced off-diagonal mixing in Gen. III will have (many more) additional consequences that remain to be explored.

9.2.3 Brot, Erot?

So far, in Gen. I and II and also in the Gen. III prototype experiments, we have used the laboratory-frame magnetic field gradient as a way to apply a Zeeman shift; the micromotion driven

¹² Note that this estimate quantifies only the differential dressing of a Berry-phase response that is diagonal in the pure-state basis (i.e. a contribution entering through ζ^0 term. It does not include Berry-related imperfections that enter through residual off-diagonal perturbations like the $\delta\Delta^{BR}$ and $\delta\Delta^{BDR}$ terms in Eq (9.22). These terms provide an additional pathway by which rotating-field imperfections can enter the D -odd channel, and their sizes will depend on the detailed structure of the field inhomogeneities, rather than only on Δ_0^R , Δ_0^{RD} , and ζ_0^B . These effects will, in general, become more important in Gen. III as the baseline mixing terms are larger.

by the rotating electric field samples this magnetic field gradient to result in the effective rotating-frame magnetic bias:

$$B^{\text{eff}}(\mathbf{r}) \equiv \langle \mathbf{B}(\mathbf{r} + \mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \rangle_t. \quad (9.29)$$

In this way of generating the magnetic field, the spatial inhomogeneity of the rotating electric field enters mainly through the micromotion orbit $\mathbf{r}_{\text{rot}}$ and its position dependence. This is the same mechanism that was discussed in Chapter 7, where a spatially varying $r_{\text{rot}}(\mathbf{r})$ could change the effective Zeeman shift sampled by the ions.

For the Gen. III apparatus, we are considering the possibility of applying an actual rotating magnetic field $\mathbf{B}_{\text{rot}}$, using a set of wires carrying currents driven at ω_{rot} . In this case, the rotating electric and rotating magnetic fields are two distinct fields, each of which can have its own spatial structure. Then, instead of Eq. (9.29), it is more natural to define the longitudinal rotating-frame magnetic field as

$$B_{\parallel}^{\text{eff}}(\mathbf{r}) \equiv \left\langle \mathbf{B}_{\text{rot}}(\mathbf{r}, t) \cdot \hat{\mathcal{E}}_{\text{rot}}(\mathbf{r}, t) \right\rangle_t, \quad (9.30)$$

where $\hat{\mathcal{E}}_{\text{rot}}$ is the instantaneous direction of the rotating electric field

$$\hat{\mathcal{E}}_{\text{rot}} \equiv \frac{\mathcal{E}_{\text{rot}}(\mathbf{r}, t)}{|\mathcal{E}_{\text{rot}}(\mathbf{r}, t)|}, \quad (9.31)$$

which sets the molecular quantization axis. In the direct-drive of the rotating magnetic field, the longitudinal Zeeman shift is controlled by the overlap between the rotating magnetic and electric field, rather than by the magnetic field gradient sampled by the micromotion orbit.

With the new Gen. III notation, the longitudinal overlap $B_{\parallel}^{\text{eff}}$ contributes to the diagonal shift ζ , whereas the imperfections of the overlap generate additional mixing and contribute to the off-diagonal term Δ . The imperfections in the rotating fields appear in the perturbative quantities

$$\delta\zeta(\tilde{B}, \tilde{D}, \tilde{R}), \quad \delta\Delta(\tilde{B}, \tilde{D}, \tilde{R}).$$

Suppose that we have the fields of the form:

$$\mathcal{E}_{\text{rot}}(\mathbf{r}, t) = \mathcal{E}_0(\mathbf{r}) \left[\cos(\omega_{\text{rot}}t + \phi_E(\mathbf{r}))\hat{\mathbf{x}} + \sin(\omega_{\text{rot}}t + \phi_E(\mathbf{r}))\hat{\mathbf{y}} \right], \quad (9.32)$$

and similarly

$$\mathcal{B}_{\text{rot}}(\mathbf{r}, t) = \mathcal{B}_0(\mathbf{r}) \left[\cos(\omega_{\text{rot}} t + \phi_B(\mathbf{r})) \hat{\mathbf{x}} + \sin(\omega_{\text{rot}} t + \phi_B(\mathbf{r})) \hat{\mathbf{y}} \right], \quad (9.33)$$

where ϕ_E and ϕ_B are phases of the rotating electric and magnetic fields, respectively. With this, Eq. (9.30) becomes¹³

$$B_{||}^{\text{eff}}(\mathbf{r}) = \mathcal{B}_0(\mathbf{r}) \cos(\theta_{BE}(\mathbf{r})), \quad (9.34)$$

where I define a relative phase or misalignment angle by

$$\theta_{BE}(\mathbf{r}) \equiv \phi_B(\mathbf{r}) - \phi_E(\mathbf{r}). \quad (9.35)$$

Equation (9.34) shows that the effective longitudinal Zeeman bias is controlled by both the local magnitude of the magnetic field B_0 and the local alignment between the rotating magnetic and electric fields, which is very intuitive.

The ideal situation corresponds to no spatial structure in the rotating magnetic field $\mathcal{B}_0$ such that

$$\mathcal{B}_0(\mathbf{r}) = \bar{\mathcal{B}}_0, \quad \theta_{BE}(\mathbf{r}) = 0 \text{ or } \pi,$$

giving us

$$B_{||}^{\text{eff}}(\mathbf{r}) = \pm \bar{\mathcal{B}}_0.$$

In other words, the rotating magnetic field is either parallel or antiparallel to the rotating electric field direction, and both fields are spatially uniform.

To see how deviations of the fields contribute to this ideal case, we write

$$\mathcal{B}_0(\mathbf{r}) = \bar{\mathcal{B}}_0 + \delta\mathcal{B}_0(\mathbf{r}), \quad \theta_{BE}(\mathbf{r}) = \bar{\theta}_{BE} + \delta\theta_{BE}(\mathbf{r}), \quad (9.36)$$

with $|\delta\mathcal{B}_0| \ll \bar{\mathcal{B}}_0$ and $|\delta\theta_{BE}| \ll 1$. We can then expand Eq. (9.34)

$$B_{||}^{\text{eff}}(\mathbf{r}) \approx \bar{\mathcal{B}}_0 \cos(\bar{\theta}_{BE}) + \delta\mathcal{B}_0(\mathbf{r}) \cos(\bar{\theta}_{BE}) - \bar{\mathcal{B}}_0 \sin(\bar{\theta}_{BE}) \delta\theta_{BE}(\mathbf{r}), \quad (9.37)$$

to first order. We see that the first term is the intended common-mode bias. The second term shows how the spatial variation of the rotating magnetic field amplitude contribute to the shift.

¹³ I use the trig identity $\cos a \cos b + \sin a \sin b = \cos(a - b)$

The third term is the shift caused by the spatial variation of the alignment between the rotating magnetic and electric fields.

The rotating field imperfections can be described by using the same parity-state notation introduced earlier in this chapter.

$$\begin{aligned}\delta\zeta^{\mathbf{P}} &\propto (\delta B_{\parallel}^{\text{eff}})^{\mathbf{P}}, \\ \delta\Delta^{\mathbf{P}} &\propto (\delta B_{\perp}^{\text{eff}})^{\mathbf{P}}\end{aligned}\tag{9.38}$$

where $\mathbf{p}$ are all the possible parity state¹⁴ shown in Eq. (9.5) and $\delta B_{\perp}^{\text{eff}}$ is any transverse component of the effective magnetic field.

Expressing this way, once the residual field imperfections are written as contributions to $\delta\zeta(\tilde{B}, \tilde{D}, \tilde{R})$ and $\delta\Delta(\tilde{B}, \tilde{D}, \tilde{R})$, their effect on the measured Ramsey can be organized by their switch channel. For the EDM measurement, the parity channels that matter most are those in the D -odd channel as in Eq. (9.22). At leading order, these terms are

$$\delta\zeta^D, \quad \delta\zeta^0, \quad \delta\Delta^{BDR}, \quad \delta\Delta^{BR}.$$

The diagonal terms $\delta\zeta^D$ and $\delta\zeta^0$ arise from imperfections that modify the effective longitudinal bias field $B_{\parallel}^{\text{eff}}$. For example, a D -odd change in the overlap between the rotating magnetic and electric field contributes to $\delta\zeta^D$. A nonreversing longitudinal component contributes to $\delta\zeta^0$. By contrast, the off-diagonal terms $\delta\Delta^{BDR}$ and $\delta\Delta^{BR}$ describe imperfections that do not simply shift the longitudinal Zeeman splitting, but instead generate mixing between the dressed basis state. Physically, these could arise from transverse components of the rotating magnetic field relative to $\hat{\mathcal{E}}_{\text{rot}}$ due to imperfect synchronization between the rotating electric and magnetic fields, ellipticity, or higher-harmonic structure in either field.

One subtle point is that the magnitude of $\mathcal{E}_0(\mathbf{r})$ of the rotating electric field does not appear explicitly in Eq. (9.37). This is because the longitudinal Zeeman shift depends only on the direction of the rotating electric field $\hat{\mathcal{E}}_{\text{rot}}$, rather than on the field magnitude itself.¹⁵

¹⁴ Note, with the rotating magnetic and electric field, in theory, there would be two rotation-related chops; one for the electric field, $\tilde{R}_E$ and another for the magnetic field $\tilde{R}_B$.

¹⁵ If I pause and think about it, this is pretty obvious, but it is still amusing because I've spent a lot more time

However, for systematics, this does not mean that the inhomogeneity of $\mathcal{E}_{\text{rot}}$ is irrelevant. A spatially or temporally varying electric field magnitude still can alter the dressing of the states and therefore change quantities such as $\zeta_0^B, \delta\zeta, \Delta_0^R, \Delta_0^{DR}$, and $\delta\Delta$. In this sense, we still have to care about the homogeneity of the rotating electric field magnitude.

The main point of this discussion is that once the rotating magnetic field is generated independently, imperfections in the two fields must be analyzed as separate sources of systematic error. As shown in Eq. (9.37), the imperfections in the rotating magnetic field relative to the rotating electric field can shift the longitudinal magnetic field, and therefore shift the Zeeman splitting. Therefore, spatial inhomogeneity, transverse components, ellipticity, and higher harmonics of the applied rotating magnetic field provide new pathways to $\delta\zeta$ and $\delta\Delta$.

9.3 Sensitivity in Gen. III prototype

In this section, I estimate experimental sensitivity of the Gen. III prototype.

9.3.1 Long coherence time and scatter

By using the octupole trap configuration discussed in Sec. 6.2, we were able to observe a Ramsey fringe out to late times of approximately 18 s, as shown in Fig. 9.1.

The initial Ramsey contrast is not ideal, and the ion number is also limited by several signal-loss mechanisms discussed in Chapters 4 and 5. We do observe a decay in the Ramsey contrast at long times. At the moment, however, it is not clear whether this reduction is due to a genuine dephasing mechanism, imperfect field stability/homogeneity, or simply the decreasing signal-to-noise ratio at long interrogation times.

Nonetheless, with this data, we can try to estimate the quantum projection noise (QPN) limit and compare it to several experimentally observed scatters, following the general procedure

thinking about the inhomogeneity of the rotating electric field being mapped onto the Ramsey frequency with the applied B_{axgrad} . In a way, by using the rotating magnetic field, we have essentially transferred the issue of the inhomogeneity to the magnetic field.




Figure 9.1: Ramsey asymmetry and noise analysis in the Gen. III prototype. **Top:** measured Ramsey asymmetry for the upper and lower doublets as a function of a Ramsey interrogation time. Ramsey fringes are observed out to approximately 18 seconds. **Bottom:** corresponding *raw* differential-asymmetry scatter estimated in several ways. the uncorrelated upper and lower scatter (green), ion-number normalized uncorrelated scatter, the correlated differential scatter from simultaneous doublet detection, and the quantum projection noise (QPN) limit. The total averaged number of ions (including both upper and lower doublets) for each time segments are approximately: 180, 180, 170, 160, 130, 80, 65, 40, 30 for two adjacent shots of *in-in* and *in-out*.

of Ref. [41]. Returning to Eq. 9.1, the Ramsey measurement gives

$$A(t_{\text{Ramsey}}) \sim C \cos(\phi_R) = C \cos(2\pi f t_{\text{Ramsey}} + \phi),$$

where C is the Ramsey contrast and ϕ_R is the accumulated Ramsey phase.

Near a zero crossing of the Ramsey fringe, the measured asymmetry is approximately linear in phase. I define the upper- and lower-doublet asymmetries as

$$\varphi_u = \frac{N_{\downarrow u} - N_{\uparrow u}}{N_{\downarrow u} + N_{\uparrow u}} \sim C \phi_R^u, \quad \varphi_l = \frac{N_{\downarrow l} - N_{\uparrow l}}{N_{\downarrow l} + N_{\uparrow l}} \sim C \phi_R^l, \quad (9.39)$$

where ϕ_R^u and ϕ_R^l are the Ramsey phases for the upper and lower doublets, respectively. Here, $N_{\uparrow u}$ refers to the number of ions detected to be in the state $\uparrow^u$ of the upper doublet, and similarly for the other populations. When the upper and lower doublets are in phase, the phase difference is proportional to

$$\Delta\varphi = \varphi_u - \varphi_l \sim C(\phi_R^u - \phi_R^l).$$

According to Ref. [41], the QPN limit for this differential asymmetry can be expressed as

$$\sigma_{\Delta\varphi}^{\text{QPN}} = \sqrt{\frac{1 - \varphi_u^2}{N_u} + \frac{1 - \varphi_l^2}{N_l}} \quad (9.40)$$

where N_u and N_l are the total detected ion numbers for the upper and lower doublets, respectively, formed from the two adjacent cycles used to construct the asymmetries. Explicitly,

$$N_u = N_{\downarrow u} + N_{\uparrow u}, \quad N_l = N_{\downarrow l} + N_{\uparrow l}.$$

These numbers are the total ion numbers that enters the calculation of each doublet asymmetry. Therefore, when we are forming the scatter on the asymmetry, these numbers are not divided by 2, even at the zero crossing.¹⁶

It is also useful to compare this QPN floor with several experimentally measured scatter metrics. For instance, if we assume no correlation between the upper and lower doublet asymmetries,

¹⁶ Although each output port, $N_{\downarrow}$ and $N_{\uparrow}$, contains approximately half of the ions at the zero crossing, the variance of the asymmetry goes as

$$\sigma_{\varphi}^2 \propto 1/N,$$

with N the total number of ions in both outputs.

then the total scatter is just

$$\sigma_{\text{uncorr}} = (\sigma_{\varphi_u}^2 + \sigma_{\varphi_l}^2)^{1/2}, \quad (9.41)$$

where σ_{φ_u} and σ_{φ_l} are the standard deviations of φ_u and φ_l , respectively. We can also define this total uncorrelated scatter with ion number normalization to reduce the contribution from shot-to-shot ion production noise.

With simultaneous detection of both doublets, we can form correlated noise rejection. In a given experimental shot, $N_{\downarrow u}$ and $N_{\downarrow l}$ are measured simultaneously, and therefore they share the same fluctuations in ion production, state preparation, and state detection. The same is true for between $N_{\uparrow u}$ and $N_{\uparrow l}$. In addition, when the upper and lower doublet Ramsey fringes are in phase, they also share any common-mode phase fluctuations. We can use these correlations to calculate the correlated differential scatter as

$$\sigma_{\Delta\varphi}^{\text{corr}} = \sigma(\varphi_u - \varphi_l) \quad (9.42)$$

9.3.2 eEDM sensitivity

We can now use these scatter estimates to estimate the corresponding sensitivity to d_e . In this calculation below, I use the QPN limit, assuming the apparatus can be operated near that limit.

For each doublet, the measured Ramsey asymmetry is

$$A_u \equiv \varphi_u \sim C\phi_R^u, \quad A_l \equiv \varphi_l \sim C\phi_R^l, \quad (9.43)$$

near a zero crossing of the fringe. Here, C is the Ramsey contrast, and ϕ_R^u and ϕ_R^l are the Ramsey phases of upper and lower doublets, respectively.

There are two equivalent ways to define an EDM-sensitive differential channel. The first is the raw differential asymmetry:

$$\Delta A \equiv A_u - A_l, \quad (9.44)$$

for which

$$\Delta A \sim C(\phi_R^u - \phi_R^l) \equiv C\Delta\phi_R. \quad (9.45)$$

The second is the half of the raw differential channel

$$A_D \equiv \frac{A_u - A_l}{2}, \quad (9.46)$$

for which

$$A_D \sim C \frac{\phi_R^u - \phi_R^l}{2} \equiv C \phi_R^D. \quad (9.47)$$

The eEDM is inferred through one of these measured asymmetry channels. Using a generic asymmetry A ,

$$\delta A = \frac{\partial A}{\partial d_e} \delta d_e.$$

Then, using the noise on the measured signal A to be σ_A , the uncertainty in d_e is

$$\delta d_e = \frac{\sigma_A}{|\partial A / \partial d_e|} = \frac{\sigma_A}{|\partial A / \partial \phi_R| |\partial \phi_R / \partial d_e|}, \quad (9.48)$$

where in the second equality, I used chain rule to directly relate d_e to the measured A . The side of the fringe, where the slope is maximal, is the operating point for the eEDM measurement, i.e. $|\partial A / \partial \phi_R| = C$.

If we use the raw differential asymmetry $\Delta A = A_u - A_l$, and therefore the raw differential phase measurement $\Delta \phi_R$, then

$$\frac{\partial(\Delta \phi_R)}{\partial d_e} = \frac{4E_{\text{eff}}T}{\hbar},$$

Combining everything together, we get

$$\delta d_e^{\text{pair}} = \frac{\hbar}{4CE_{\text{eff}}T_{\text{coherence}}} \sigma_{\Delta A}. \quad (9.49)$$

Here, δd_e^{pair} denote the sensitivity of d_e from a pair of measurement *in-in* and *anti-in* that forms an asymmetry.¹⁷

From Fig. 9.1, the QPN-limited scatter of the raw differential asymmetry is approximately

$$\sigma_{\Delta \phi}^{\text{QPN}} \sim 0.25$$

¹⁷ Note if I instead used the half-difference channel $A_D = \frac{A_u - A_l}{2}$, then $\delta d_e^{\text{pair}} = \frac{\hbar}{2CE_{\text{eff}}T_{\text{coherence}}} \sigma_{A_D}$. Since $\sigma_{A_D} = \sigma_{\Delta A}/2$, both expressions give the same final sensitivity.

around 12 s. Using $C \sim 0.3$, this yields an estimation of

$$\delta d_e^{\text{pair}} \approx 2.5 \times 10^{-28} \text{ e cm.} \quad (9.50)$$

The sensitivity per $\sqrt{\text{Hz}}$ must then be corrected for the full acquisition time of one *in-in/anti-in* pair. Since one pair requires two shots total,

$$\eta_{d_e} = \delta d_e^{\text{pair}} \sqrt{2(T_{\text{coherence}} + T_{\text{dead}})} \approx \sqrt{26} \times 2.5 \times 10^{-28} \approx 1.3 \times 10^{-27} \text{ e cm}/\sqrt{\text{Hz}}, \quad (9.51)$$

where for the last one, I have used 1-second dead time.

This crude estimate is also assuming that we can determine upper and lower doublets unambiguously. In the prototype apparatus, that assumption is not strictly valid, since the doublet contrast is less than 1. In the future generation apparatus, we expect the statistical sensitivity to improve—we will have higher ion number through an improved trap geometry, operation at larger f_{rot} , have a higher contrast¹⁸, and potentially reduced residual hydrogen partial pressure. We also expect to improve the experimental duty cycle, which will further enhance the statistical sensitivity. Most importantly, we will be able to operate either with batch or continuous-mode, which should increase the sensitivity further.

¹⁸ This is an ongoing effort in our group to increase initial contrast

Chapter 10

Conclusion

The work present in this thesis focused on understanding and extending the coherence time of the Gen. III ThF^+ eEDM prototype experiment. The limitations on the coherence (interrogation) time were broadly divided into two categories: loss of signal and loss of coherence. I also presented several proof-of-principle steps towards continuous-mode operation, as well as the demonstration of long-coherence Ramsey spectroscopy in the prototype apparatus.

Looking back, reaching long Ramsey coherence in the prototype took considerably longer than we initially expected. In the earlier JILA eEDM experiments, many of the mechanisms discussed in this thesis were either absent or small enough not to be limiting. Extending the interrogation time from the few-second scale to the tens-of-seconds scale changed this; effects that were previously secondary became important and often the experiment improved only by solving one limitation and then discovering the next one.¹ Even the “simple” task of holding on to the ions for long enough became a nontrivial experimental challenge.

The Gen. III apparatus is now being designed and constructed with the lessons from the prototype in mind. The prototype has clarified several essential design requirements: improved vacuum and material choices to reduce residual-gas collisions, better magnetic “cleanliness” and field control, more robust rotating-field generation, improved trap geometries with larger effective depth, and timing protocols compatible with high-throughput operation. It is likely that new limitations will appear as the apparatus improves, but the work in this thesis provides a framework

¹ It never ends...

for identifying them and separating their effects on signal, trap lifetime, and Ramsey contrast.

The JILA eEDM experiment, especially the Gen. III apparatus, is a very complex experiment. I think it is very remarkable that such a large and technically demanding apparatus can be operated by a relatively small team of graduate students and postdocs. As the group moves from the prototype toward the full Gen. III apparatus, we plan to decommission the prototype soon.² The prototype has served its purposes: it demonstrated long coherence Ramsey spectroscopy with ThF^+ and identified the main experimental requirements for turning that capability into a high-sensitivity eEDM measurement. I hope that the studies presented here will serve as a useful foundation for the next stage of the experiment.

² Big sad face

Bibliography

- [1] V. Andreev, D. G. Ang, D. DeMille, J. M. Doyle, G. Gabrielse, J. Haefner, N. R. Hutzler, Z. Lasner, C. Meisenhelder, B. R. O’Leary, C. D. Panda, A. D. West, E. P. West, X. Wu, and ACME Collaboration. Improved limit on the electric dipole moment of the electron. Nature, 562:355–360, 2018.
- [2] Jonathan J. Hudson, Dhiren M. Kara, I. J. Smallman, Ben E. Sauer, Michael R. Tarbutt, and E. A. Hinds. Improved measurement of the shape of the electron. Nature, 473(7348):493–496, 2011.
- [3] William B. Cairncross, Daniel N. Gresh, Matt Grau, Kevin C. Cossel, Tanya S. Roussy, Yiqi Ni, Yan Zhou, Jun Ye, and Eric A. Cornell. Precision measurement of the electron’s electric dipole moment using trapped molecular ions. Physical Review Letters, 119:153001, 2017.
- [4] Brent Graner, Y. Chen, E. G. Lindahl, and B. R. Heckel. Reduced limit on the permanent electric dipole moment of ^{199}Hg . Physical Review Letters, 116(16):161601, 2016.
- [5] Michael Bishof, Richard H. Parker, Kevin G. Bailey, John P. Greene, Roy J. Holt, Mukut R. Kalita, Wolfgang Korsch, Nathan D. Lemke, Zheng-Tian Lu, Peter Mueller, Thomas P. O’Connor, Jaideep T. Singh, and Matthew R. Dietrich. Improved limit on the ^{225}Ra electric dipole moment. Physical Review C, 94(2):025501, 2016.
- [6] F. Kuchler, E. Babcock, M. Burghoff, T. Chupp, S. Degenkolb, I. Fan, P. Fierlinger, F. Gong, E. Kraegeloh, W. Kilian, et al. A new search for the atomic edm of ^{129}Xe at frm-ii. Hyperfine Interactions, 237(1):95, 2016.
- [7] E. R. Tardiff, E. T. Rand, G. C. Ball, T. E. Chupp, A. B. Garnsworthy, P. Garrett, M. E. Hayden, C. A. Kierans, W. Lorenzon, M. R. Pearson, et al. The radon edm apparatus. Hyperfine Interactions, 225(1–3):197–206, 2014.
- [8] Tanya S. Roussy, Luke Caldwell, Trevor Wright, William B. Cairncross, Yuval Shagam, Kia Boon Ng, Noah Schlossberger, Sun Yool Park, Anzhou Wang, Jun Ye, and Eric A. Cornell. An improved bound on the electron’s electric dipole moment. Science, 381(6653):46–50, 2023.
- [9] William B. Cairncross and Jun Ye. Atoms and molecules in the search for time-reversal symmetry violation. Nature Reviews Physics, 1:510–521, 2019.
- [10] A. D. Sakharov. Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe. JETP Letters, 5:24–27, 1967. Originally published in Pisma Zh. Eksp. Teor. Fiz. 5, 32–35 (1967).

- [11] Michael Dine and Alexander Kusenko. The origin of the matter-antimatter asymmetry. Reviews of Modern Physics, 76(1):1–30, 2003.
- [12] T. E. Chupp, P. Fierlinger, M. J. Ramsey-Musolf, and J. T. Singh. Electric dipole moments of atoms, molecules, nuclei, and particles. Reviews of Modern Physics, 91(1):015001, 2019.
- [13] Maxim Pospelov and Adam Ritz. Ckm benchmarks for electron electric dipole moment experiments. Physical Review D, 89(5):056006, 2014.
- [14] M. S. Safronova, D. Budker, D. DeMille, Derek F. Jackson Kimball, A. Derevianko, and C. W. Clark. Search for new physics with atoms and molecules. Reviews of Modern Physics, 90(2):025008, 2018.
- [15] Tanya S. Roussy. A new limit on the electron’s electric dipole moment. PhD thesis, University of Colorado, Boulder, CO, August 2022. Ph.D. thesis; JILA PI advisor: Eric A. Cornell.
- [16] Timothy Chupp and Michael Ramsey-Musolf. Electric dipole moments: A global analysis. Physical Review C, 91(3):035502, 2015.
- [17] J. J. Sakurai and Jim Napolitano. Modern Quantum Mechanics. Addison-Wesley, Boston, 2 edition, 2011.
- [18] Norman F. Ramsey. A new molecular beam resonance method. Physical Review, 76:996, 1949.
- [19] Norman F. Ramsey. A molecular beam resonance method with separated oscillating fields. Physical Review, 78:695–699, 1950.
- [20] Norman F. Ramsey. Molecular Beams. Oxford University Press, Oxford, 1956.
- [21] Norman F. Ramsey. Experiments with separated oscillatory fields and hydrogen masers. Reviews of Modern Physics, 62(3):541–552, 1990.
- [22] Marlan O. Scully and M. Suhail Zubairy. Quantum Optics. Cambridge University Press, Cambridge, 1997.
- [23] W. M. Itano, J. C. Bergquist, J. J. Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland. Quantum projection noise: Population fluctuations in two-level systems. Physical Review A, 47(5):3554–3570, May 1993.
- [24] D. J. Wineland, J. J. Bollinger, W. M. Itano, and D. J. Heinzen. Squeezed atomic states and projection noise in spectroscopy. Physical Review A, 50(1):67–88, July 1994.
- [25] David E. Morrissey and Michael J. Ramsey-Musolf. Electroweak baryogenesis. New Journal of Physics, 14:125003, 2012.
- [26] Jorinde van de Vis, Jordy de Vries, and Marieke Postma. Bubble trouble: a review on electroweak baryogenesis. arXiv preprint arXiv:2508.09989, 2025. Review article, preprint.
- [27] Kaori Fuyuto, Junji Hisano, and Eibun Senaha. Toward verification of electroweak baryogenesis by electric dipole moments. Physics Letters B, 755:491–497, 2016.

- [28] Yuan-Zhen Li, Michael J. Ramsey-Musolf, and Jiang-Hao Yu. Does the electron EDM preclude electroweak baryogenesis? arXiv preprint arXiv:2404.19197, 2024. Preprint.
- [29] A. E. Leanhardt, J. L. Bohn, H. Loh, P. Maletinsky, E. R. Meyer, L. C. Sinclair, R. P. Stutz, and E. A. Cornell. High-resolution spectroscopy on trapped molecular ions in rotating electric fields: A new approach for measuring the electron electric dipole moment. Journal of Molecular Spectroscopy, 270(1):1–25, 2011.
- [30] Huanqian Loh, Kevin C. Cossel, Matt Grau, Kang-Kuen Ni, Edmund R. Meyer, John L. Bohn, Jun Ye, and Eric A. Cornell. Precision spectroscopy of polarized molecules in an ion trap. Science, 342(6163):1220–1222, 2013.
- [31] Edmund R. Meyer, John L. Bohn, and Michael P. Deskevich. Candidate molecular ions for an electron electric dipole moment experiment. Physical Review A, 73:062108, 2006.
- [32] Edmund R. Meyer and John L. Bohn. Prospects for an electron electric-dipole moment search in metastable tho and thf^+ . Physical Review A, 78:010502, 2008.
- [33] Kia Boon Ng, Yan Zhou, Lan Cheng, Noah Schlossberger, Sun Yool Park, Tanya Roussy, Luke Caldwell, Yuval Shagam, Antonio Vigil, Eric A. Cornell, and Jun Ye. Spectroscopy on the electron-electric-dipole-moment-sensitive states of thf^+ . Physical Review A, 105:022823, 2022.
- [34] P. G. H. Sandars. Measurability of the proton electric dipole moment. Physical Review Letters, 19:1396–1398, 1967.
- [35] John M. Brown and Alan Carrington. Rotational Spectroscopy of Diatomic Molecules. Cambridge University Press, Cambridge, 2003.
- [36] Daniel N. Gresh, Kevin C. Cossel, Yan Zhou, Jun Ye, and Eric A. Cornell. Broadband velocity modulation spectroscopy of ThF^+ for use in a measurement of the electron electric dipole moment. Journal of Molecular Spectroscopy, 319:1–9, January 2016.
- [37] William Cairncross. Searching for Time-Reversal Symmetry Violation with Molecular Ions: Quantum State Control and Photofragment Imaging. Ph.d. thesis, University of Colorado Boulder, Boulder, Colorado, July 2019.
- [38] Kia Boon Ng. The ThF^+ eEDM experiment: concept, design, and characterization. PhD thesis, University of Colorado Boulder, Boulder, CO, 2023. PDF hosted by JILA; record lists online posting month as 2024-05.
- [39] J. M. Brown, A. S.-C. Cheung, and A. J. Merer. λ -type doubling parameters for molecules in δ electronic states. Journal of Molecular Spectroscopy, 124:464–475, 1987.
- [40] Stephen Eckel, Paul Hamilton, Emil Kirilov, H. W. Smith, and David DeMille. Search for the electron electric dipole moment using ω -doublet levels in pbo . Physical Review A, 87:052130, 2013.
- [41] Yan Zhou, Yuval Shagam, William B. Cairncross, Kia Boon Ng, Tanya S. Roussy, Tanner Grogan, Kevin Boyce, Antonio Vigil, Madeline Pettine, Tanya Zelevinsky, Jun Ye, and Eric A. Cornell. Second-scale coherence measured at the quantum projection noise limit with hundreds of molecular ions. Physical Review Letters, 124:053201, 2020.

- [42] Kang-Kuen Ni, Huanqian Loh, Matt Grau, Kevin C. Cossel, Jun Ye, and Eric A. Cornell. State-specific detection of trapped hff^+ by photodissociation. Journal of Molecular Spectroscopy, 300:12–15, 2014.
- [43] Kia Boon Ng, Sun Yool Park, Anzhou Wang, Addison Hartman, Rohan Kompella, Lan Cheng, Xing Fan, Stephan Malbrunot-Ettenauer, Jun Ye, and Eric Cornell. High-efficiency quantum-state detection of thf^+ with resonance-enhanced multiphoton asymmetric dissociation. Physical Review A, 112:053104, 2025.
- [44] J. Alberto Beswick and Richard N. Zare. On the quantum and quasiclassical angular distributions of photofragments. The Journal of Chemical Physics, 129:164315, 2008.
- [45] M. McDonald, B. H. McGuyer, F. Apfelbeck, C.-H. Lee, I. Majewska, R. Moszynski, and T. Zelevinsky. Photodissociation of ultracold diatomic strontium molecules with quantum state control. Nature, 535:122–126, 2016.
- [46] I. Majewska, S. S. Kondov, C.-H. Lee, M. McDonald, B. McGuyer, R. Moszynski, and T. Zelevinsky. Experimental and theoretical investigation of the crossover from the ultracold to the quasiclassical regime of photodissociation. Physical Review A, 98:043404, 2018.
- [47] Kevin C. Cossel, Daniel N. Gresh, Laura C. Sinclair, Tyler Coffey, Leonid V. Skripnikov, Alexander N. Petrov, Nikolai S. Mosyagin, Anatoly V. Titov, Robert W. Field, Edmund R. Meyer, Eric A. Cornell, and Jun Ye. Broadband velocity modulation spectroscopy of HfF^+ : Towards a measurement of the electron electric dipole moment. Chemical Physics Letters, 546:1–11, 2012.
- [48] A. N. Petrov, L. V. Skripnikov, and A. V. Titov. Zeeman interaction in the $^3\delta_1$ state of HfF^+ to search for the electron electric-dipole-moment. Physical Review A, 96:022508, 2017.
- [49] L. V. Skripnikov. Theoretical study of HfF^+ cation to search for the t, p -odd interactions. The Journal of Chemical Physics, 147:021101, 2017.
- [50] N. J. Stone. Table of nuclear magnetic dipole and electric quadrupole moments. Atomic Data and Nuclear Data Tables, 90(1):75–176, 2005.
- [51] Luke Caldwell, Tanya S. Roussy, Trevor Wright, William B. Cairncross, Yuval Shagam, Kia Boon Ng, Noah Schlossberger, Sun Yool Park, Anzhou Wang, Jun Ye, and Eric A. Cornell. Systematic and statistical uncertainty evaluation of the hff^+ electron electric dipole moment experiment. Phys. Rev. A, 108:012804, Jul 2023.
- [52] L. V. Skripnikov and A. V. Titov. Theoretical study of ThF^+ in the search for T, P -violation effects: Effective state of a Th atom in ThF^+ and ThO compounds. Physical Review A, 91:042504, 2015.
- [53] Malika Denis, Morten S. Nørby, Hans Jørgen Aagaard Jensen, André Severo Pereira Gomes, Malaya K. Nayak, Stefan Knecht, and Timo Fleig. Theoretical study on ThF^+ , a prospective system in search of time-reversal violation. New Journal of Physics, 17:043005, 2015.
- [54] Wolfgang Paul. Electromagnetic traps for charged and neutral particles. Reviews of Modern Physics, 62(3):531–540, 1990.

- [55] Fouad G. Major, Viorica N. Gheorghe, and Günther Werth. Charged Particle Traps: Physics and Techniques of Charged Particle Field Confinement. Springer, Berlin, Heidelberg, 2005.
- [56] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. Reviews of Modern Physics, 75(1):281–324, 2003.
- [57] R. P. Stutz. Towards Measuring the Electron Electric Dipole Moment Using Trapped Molecular Ions. Ph.D. thesis, University of Colorado Boulder, Boulder, Colorado, 2010.
- [58] Huanqian Loh. Search for an Electron Electric Dipole Moment with Trapped Molecular Ions. Ph.D. thesis, University of Colorado Boulder, Boulder, Colorado, 2013.
- [59] Noah Schlossberger. An apparatus for measuring the electron’s electric dipole moment in trapped ThF^+ . PhD thesis, University of Colorado Boulder, Boulder, CO, 2023.
- [60] Trevor H. Wright. Stray Fields and The Electron’s Electric Dipole Moment. PhD thesis, University of Colorado Boulder, Boulder, CO, May 2024. Ph.D. thesis; JILA PI advisors: Eric Cornell and Jun Ye.
- [61] Kuang Chen, Scott T. Sullivan, Wade G. Rellergert, and Eric R. Hudson. Measurement of the coulomb logarithm in a radio-frequency paul trap. Physical Review Letters, 110(17):173003, April 2013.
- [62] J. A. Krommes. An introduction to the physics of the coulomb logarithm, with emphasis on quantum-mechanical effects. Journal of Plasma Physics, 85(1):925850101, 2019.
- [63] Martin Ligare. Classical thermodynamics of particles in harmonic traps. American Journal of Physics, 78(8):815–819, 2010.
- [64] Peter F. Bernath. Spectra of Atoms and Molecules. Oxford University Press, 4 edition, 2020.
- [65] Nicolas Vanhaecke and Olivier Dulieu. Precision measurements with polar molecules: the role of the black body radiation. Molecular Physics, 105(11–12):1723–1731, 2007.
- [66] David J. Griffiths. Introduction to Quantum Mechanics. Pearson Prentice Hall, 2 edition, 2005.
- [67] George Gioumousis and D. P. Stevenson. Reactions of gaseous molecule ions with gaseous molecules. v. theory. The Journal of Chemical Physics, 29(2):294–299, 1958.
- [68] National Institute of Standards and Technology. Nist computational chemistry comparison and benchmark database. NIST Standard Reference Database Number 101, Release 22, May 2022, Editor: Russell D. Johnson III, 2022.
- [69] Julia E. Rice, Peter R. Taylor, Timothy J. Lee, and Jan Almlof. The determination of accurate dipole polarizabilities α and γ for the noble gases. The Journal of Chemical Physics, 94(7):4972–4979, 1991.
- [70] Jim Mitroy, Marianna S. Safronova, and Charles W. Clark. Theory and applications of atomic and ionic polarizabilities. Journal of Physics B: Atomic, Molecular and Optical Physics, 43(20):202001, 2010.

- [71] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a paul trap. Journal of Applied Physics, 83(10):5025–5033, 1998.
- [72] T. Baba and I. Waki. Sympathetic cooling rate of gas-phase ions in a radio-frequency-quadrupole ion trap. Applied Physics B, 74(4-5):375–382, 2002.
- [73] Yoshiki Moriwaki, Maki Tachikawa, Yoshiharu Maeno, and Tadao Shimizu. Collision cooling of ions stored in quadrupole radio-frequency trap. Japanese Journal of Applied Physics, 31(11B):L1640–L1643, 1992.
- [74] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. Journal of Research of the National Institute of Standards and Technology, 103(3):259–328, 1998.
- [75] David R. Leibbrandt, Bernard Yurke, and Richard E. Slusher. Modeling ion trap thermal noise decoherence. Quantum Information and Computation, 7(1–2):52–72, 2007.
- [76] Wolfgang Paul and Helmut Steinwedel. Notizen: Ein neues massenspektrometer ohne magnetfeld. Zeitschrift für Naturforschung A, 8(7):448–450, 1953. German; often cited in English as “A new mass spectrometer without a magnetic field”.
- [77] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. Physical Review Letters, 74(20):4091–4094, 1995.
- [78] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. Physical Review Letters, 75(25):4714–4717, 1995.
- [79] H. Häffner, C. F. Roos, and R. Blatt. Quantum computing with trapped ions. Physics Reports, 469(4):155–203, 2008.
- [80] R. Blatt and D. Wineland. Entangled states of trapped atomic ions. Nature, 453(7198):1008–1015, 2008.
- [81] D. Main, P. Drmota, D. P. Nadlinger, E. M. Ainley, A. Agrawal, B. C. Nichol, R. Srinivas, G. Araneda, and D. M. Lucas. Distributed quantum computing across an optical network link. Nature, 638(8050):383–388, 2025.
- [82] Minzhao Liu, Ruslan Shaydulin, Pradeep Niroula, Matthew DeCross, Shih-Han Hung, Wen Yu Kon, Enrique Cervero-Martín, Kaushik Chakraborty, Omar Amer, Scott Aaronson, Atithi Acharya, Yuri Alexeev, K. Jordan Berg, Shouvanik Chakrabarti, Florian J. Curchod, Joan M. Dreiling, Neal Erickson, Cameron Foltz, Michael Foss-Feig, David Hayes, Travis S. Humble, Niraj Kumar, Jeffrey Larson, Danylo Lykov, Michael Mills, Steven A. Moses, Brian Neyenhuis, Shaltiel Eloul, Peter Siegfried, James Walker, Charles Lim, and Marco Pistoia. Certified randomness using a trapped-ion quantum processor. Nature, 640(8058):343–348, 2025.
- [83] S. A. Diddams, Th. Udem, J. C. Bergquist, et al. An optical clock based on a single trapped $^{199}\text{Hg}^+$ ion. Science, 293(5531):825–828, 2001.
- [84] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. Reviews of Modern Physics, 87:637–701, 2015.

- [85] Mason C. Marshall, Daniel A. Rodriguez Castillo, Willa J. Arthur-Dworschack, Alexander Aeppli, Kyungtae Kim, Dahyeon Lee, William Warfield, Joost Hinrichs, Nicholas V. Nardelli, Tara M. Fortier, Jun Ye, David R. Leibrandt, and David B. Hume. High-stability single-ion clock with 5.5×10^{-19} systematic uncertainty. Physical Review Letters, 135(3):033201, July 2025.
- [86] T. Lindvall, T. Fordell, K. J. Hanhijärvi, M. Doležal, J. M. Rahm, S. L. Weyers, and A. E. Wallin. $^{88}\text{Sr}^+$ optical clock with 7.9×10^{-19} systematic uncertainty and measurement of its absolute frequency with 9.8×10^{-17} uncertainty. Physical Review Applied, 24(4):044082, October 2025.
- [87] Charles Cheung, Sergey G. Porsev, Dmytro Filin, Marianna S. Safronova, Malte Wehrheim, Lukas J. Spieß, Shuying Chen, Alexander Wilzewski, José R. Crespo López-Urrutia, and Piet O. Schmidt. Finding the ultranarrow $^3p_2 \rightarrow ^3p_0$ electric quadrupole transition in Ni^{12+} ion for an optical clock. Physical Review Letters, 135(9):093002, August 2025.
- [88] S. Willitsch. Coulomb-crystallised molecular ions in traps: methods, applications, prospects. International Reviews in Physical Chemistry, 31(2):175–199, 2012.
- [89] B. R. Heazlewood. Cold ion chemistry within coulomb crystals. Molecular Physics, 117(14):1934–1941, 2019.
- [90] O. A. Krohn and H. J. Lewandowski. Cold ion–molecule reactions in the extreme environment of a Coulomb crystal. The Journal of Physical Chemistry A, 128(10):1737–1752, 2024.
- [91] Harvard Natural Sciences Lecture Demonstrations. Rotating saddle. Video. Accessed 2026-02-11.
- [92] T. Hasegawa and J. J. Bollinger. Rotating-radio-frequency ion traps. Phys. Rev. A, 72(4):043403, Oct 2005.
- [93] A. V. Gaponov and M. A. Miller. Potential wells for charged particles in a high-frequency electromagnetic field. Soviet Physics JETP, 7(1):168–169, 1958. Russian original: Zh. Eksp. Teor. Fiz. 34, 242–243 (1958).
- [94] Richard C. Fernow. Principles of Magnetostatics. Cambridge University Press, 2017.
- [95] Andrew Zangwill. Modern Electrodynamics. Cambridge University Press, 2013.
- [96] Daniel V. Schroeder. An Introduction to Thermal Physics. Oxford University Press, 2021.
- [97] Leo Radzihovsky. Physics 7230: Statistical mechanics, lecture set 3: Canonical ensemble, 2021. Lecture notes, dated 20 January 2021.
- [98] J. Goldwin, S. Inouye, M. L. Olsen, and Deborah S. Jin. Cross-dimensional relaxation in bose-fermi mixtures. Physical Review A, 71(4):043408, 2005.
- [99] Fang Fang, Shun Wu, Aaron Smull, Joshua A. Isaacs, Yimeng Wang, Chris H. Greene, and Dan M. Stamper-Kurn. Cross-dimensional relaxation of atomic gas mixtures in a spherical-quadrupole magnetic trap. Physical Review A, 101(1):012703, 2020.

- [100] A. Patscheider, L. Chomaz, G. Natale, D. Petter, M. J. Mark, S. Baier, B. Yang, R. R. W. Wang, J. L. Bohn, and F. Ferlaino. Determination of the scattering length of erbium atoms. Physical Review A, 105(6):063307, 2022.
- [101] Scott D. Baalrud and Jérôme Daligault. Temperature anisotropy relaxation of the one-component plasma. Contributions to Plasma Physics, 57(6-7):238–251, 2017.
- [102] R. Kubo. Note on the stochastic theory of resonance absorption. Journal of the Physical Society of Japan, 9(6):935–944, 1954.
- [103] Jonathan Coslovsky, Gadi Afek, Alexander Mil, Ido Almog, and Nir Davidson. Spectroscopic measurement of the softness of ultra-cold atomic collisions. Physical Review A, 96(3):032713, 2017.
- [104] Yoav Sagi, Ido Almog, and Nir Davidson. Universal scaling of collisional spectral narrowing in an ensemble of cold atoms. Physical Review Letters, 105(9):093001, 2010.
- [105] Markus Zahn. Chapter 4: Electric field boundary value problems. https://ocw.mit.edu/courses/res-6-002-electromagnetic-field-theory-a-problem-solving-approach-spring-2008/659bc3f2465095e4d9308538846c5a5a/MITRES_6_002S08_chapter4.pdf, 2008. In Electromagnetic Field Theory: A Problem Solving Approach, pp. 257–312; accessed 2026-03-19.
- [106] James A. Fedchak, Patrick J. Abbott, Jay H. Hendricks, Paul C. Arnold, and Neil T. Peacock. Recommended practice for calibrating vacuum gauges of the ionization type. Journal of Vacuum Science & Technology A, 36(3):030802, 2018.
- [107] Robert L. Summers. Empirical observations on the sensitivity of hot cathode ionization type vacuum gages. Technical Report NASA TN D-5285, NASA, 1969.
- [108] Stanford Research Systems. Models RGA100, RGA200, and RGA300 Residual Gas Analyzer Manual. Stanford Research Systems.
- [109] MKS Instruments. Baratron Capacitance Manometers Family Brochure. MKS Instruments, 2023.
- [110] Stanford Research Systems. Gas Correction Factors for Bayard-Alpert Ionization Gauges. Stanford Research Systems. Application note, accessed 2026-04-06.
- [111] T. N. Olney, N. M. Cann, G. Cooper, and C. E. Brion. Absolute scale determination for photoabsorption spectra and the calculation of molecular properties using dipole sum rules. Chemical Physics, 223(1-2):59–98, 1997.
- [112] Frederick Reif. Fundamentals of Statistical and Thermal Physics. Waveland Press, 2009.
- [113] Richard Fitzpatrick. Thermodynamics and statistical mechanics: The equipartition theorem, 2006. Lecture notes.
- [114] Harvey Gould and Gene F. Mazenko. Microscopic theory of self-diffusion in a classical one-component plasma. Physical Review A, 15(3):1274–1287, 1977.
- [115] David J. Griffiths. Introduction to Electrodynamics. Pearson, 4 edition, 2013.

- [116] J. M. Brown, J. T. Hougen, K.-P. Huber, J. W. C. Johns, I. Kopp, H. Lefebvre-Brion, A. J. Merer, D. A. Ramsay, J. Rostas, and R. N. Zare. The labeling of parity doublet levels in linear molecules. Journal of Molecular Spectroscopy, 55:500–503, 1975.
- [117] Leonard Mandel and Emil Wolf. Optical Coherence and Quantum Optics. Cambridge University Press, Cambridge, 1995.
- [118] T. Hasegawa, M. J. Jensen, and J. J. Bollinger. Stability of a penning trap with a quadrupole rotating electric field. Phys. Rev. A, 71(2):023406, February 2005.
- [119] D. Bauer, P. Mulser, and W. H. Steeb. Relativistic ponderomotive force, uphill acceleration, and transition to chaos. Phys. Rev. Lett., 75(25):4622–4625, December 1995.
- [120] Yuval Shagam, William B. Cairncross, Tanya S. Roussy, Yan Zhou, Kia Boon Ng, Daniel N. Gresh, Tanner Grogan, Jun Ye, and Eric A. Cornell. Continuous temporal ion detection combined with time-gated imaging: Normalization over a large dynamic range. Journal of Molecular Spectroscopy, 368:111257, 2020.

Appendix A

Calibration of IG and RGA

For the neutral gas measurements discussed in Chapter 4, we needed an estimate of the absolute partial pressure of gas in the apparatus. Both ion gauges and residual gas analyzers (RGAs) have species-dependent responses, so the raw pressure reading is generally not equal to the true partial pressure of the gas of interest. In this appendix, I describe the calibration procedure we used to determine empirical response factors for the ion gauge and the RGA.

The ion gauge used here is a Bayard-Alpert ionization gauge (Agilent UHV-24p). Standard Bayard-Alpert gauges conventionally report nitrogen-equivalent pressures, so gas-dependent correction factors must be applied for other species [110]. Such correction factors for Bayard-Alpert gauges are reported in several references, such as the ion gauge manuals [110]. However, our RGA (SRS RGA200) did not come with an equivalent calibration for our operating conditions. We therefore calibrated both instruments against a common reference.

A.1 Calibration of ion-gauge and RGA readings

To obtain quantitative partial pressures of the ion gauge and the RGA, we calibrated both against a Baratron capacitance manometer. The calibration setup is shown in Fig. A.1. This follows the standard approach for vacuum-gauge calibration, where the instrument under test is compared against a reference gauge under well-controlled conditions. Because a capacitance manometer is essentially insensitive to gas species, the Baratron served as the reference gauge for this calibration [106, 109].




Figure A.1: Calibration vacuum chamber used to calibrate the ion gauge and the residual gas analyzer (RGA) against a Baratron capacitance manometer. From a small upstream reservoir, gas is admitted through a leak valve into the calibration chamber. The pressure decay in the reservoir with the known volume (outlined with dotted line) determines the gas throughput. An orifice placed just upstream from the turbomolecular pump limits the effective pumping speed at the chamber to about ~ 10 L/s. The resulting steady-state chamber pressure provides the Baratron-based reference pressure, and this is compared to the readings from the ion gauge and the RGA.

Upstream of the calibration vacuum chamber, we prepared a small gas reservoir consisting of the leak-valve assembly and associated fittings, with volume of $V_{\text{res}} \approx 12.5 \text{ cc}$.¹ For each gas species, the reservoir was filled with the leak valve closed. Then the gas was admitted into the calibration chamber as we opened up the leak valve, and the pressure of the reservoir was monitored with the Baratron.

Rather than attempting to model the conductance of the leak valve directly which is hard to do, we instead inferred the gas throughput from the pressure decay of the upstream reservoir. Assuming isothermal expansion, we have

$$Q = -\frac{d(PV)}{dt} = -V_{\text{res}} \frac{dP_{\text{res}}}{dt}, \quad (\text{A.1})$$

where Q is the gas throughput from the reservoir into the calibration chamber in Torr L/s and P_{res} is the reservoir pressure.

The calibration chamber was pumped by a turbomolecular pump through an orifice placed immediately upstream of the pump. The size of the orifice was designed such that for N_2 , the effective pumping speed was $\sim 10 \text{ L/s}$.² The steady-state gas pressure in the calibration chamber is then

$$P_{\text{ref,g}} = \frac{Q}{S_{\text{eff}}}. \quad (\text{A.2})$$

This pressure provides the reference value against which the ion-gauge and RGA readouts are compared.

As an example, Fig. A.2(a) shows a representative Baratron measurement with N_2 . After the reservoir was filled with the leak valve closed, the leak valve was opened and the Baratron pressure in the reservoir began to decrease approximately linearly. From the measure slope, we have

$$\left| \frac{dP_{\text{res}}}{dt} \right| \approx 8.7 \times 10^{-4} \text{ Torr/s},$$

¹ The reservoir volume was measured by filling the assembly with methanol. We then weighed the contained methanol and converted the measured mass to volume using the density of methanol.

² For a thin aperture in the molecular-flow regime, the conductance scales approximately as $\propto M^{-1/2}$ for a gas species with mass M . Accordingly, the effective pumping speed was calculated separately for each gas species.

together with $V_{\text{res}} = 12.5 \times 10^{-3} \text{ L}$ and $S_{\text{eff},\text{N}_2} \approx 10 \text{ L/s}$, we infer a chamber pressure

$$P_{\text{ref},\text{N}_2} \approx \frac{(8.7 \times 10^{-4} \text{ Torr/s})(12.5 \times 10^{-3} \text{ L})}{(10 \text{ L/s})} \approx 1.1 \times 10^{-6} \text{ Torr}.$$

At the same time, the RGA measured a nitrogen partial pressure of approximate $9.9 \times 10^{-7} \text{ Torr}$, while the ion gauge read $1.05 \times 10^{-6} \text{ Torr}$. Therefore, for N_2 , both instruments agree with the Baratron-based reference at the $\sim 10\%$ level.




Figure A.2: Example of a Baratron measurement for N_2 . After the upstream reservoir is filled, the leak valve was opened slightly around 11:47, which resulted in an approximately linear decrease in the reservoir pressure. The measured slope, indicated with the green line, together with the known reservoir volume of $\sim 12.5 \text{ cc}$ and effective pumping speed of $\sim 10 \text{ L/s}$, is used to infer the steady-state N_2 pressure in the calibration chamber, with Eqs.(A.1) and (A.2).

We repeated this procedure for several gas species. The resulting ratios $P_{\text{meas}}/P_{\text{ref}}$ are summarized in Table A.1 as

$$C_{\text{RGA},g} = \frac{P_{\text{meas}}}{P_{\text{ref}}}, \quad C_{\text{IG},g} = \frac{P_{\text{meas}}}{P_{\text{ref}}}.$$

For the ion gauge, these ratios broadly follow the gas-correction factors found in the manual [110].

The RGA, by contrast, exhibits a stronger species dependence.

To determine the pressure range over which the calibration remained linear, we repeated the measurements over approximately 10^{-9} to $5 \times 10^{-6} \text{ Torr}$. Figure A.3 shows the result for N_2 . The ion-gauge response remained approximately proportional to the reference pressure throughout this

Gas	Mass (amu)	RGA response factor $C_{\text{RGA},g} = P_{\text{meas}}/P_{\text{ref}}$	IG response factor $C_{\text{IG},g} = P_{\text{meas}}/P_{\text{ref}}$	IG correction factor from Ref. [110]
N ₂	28	1.1	1.0	1.0
Ne	20	0.3	0.3	0.3
CO	28	1.3	1.0	1.05
CH ₄	16	1.4	1.5	1.5
H ₂	2	2.4	0.5	0.46
He	4	0.6	0.2	0.18
Kr	84	0.5	1.7	1.7

Table A.1: Measured response of the RGA and ion gauge relative to the Baratron-based reference for several calibration gases. The quantities $C_{\text{RGA},g}$ and $C_{\text{IG},g}$ are defined as the ratio of the instrument reading to the reference pressure $P_{\text{meas}}/P_{\text{ref}}$, for gas species g . The final column lists the published ion-gauge correction factors for comparison.

range. The RGA, on the other hand, began to deviate systematically from linearity above about 2×10^{-6} Torr. Generally, this behavior is consistent with the known nonlinearity response in RGA measurements at higher pressure due to ionizer space-charge effects and ion-neutral scattering [108].

A.2 Inferring the pressure at the octagon

In the main apparatus, the leak valve and the RGA were mounted on the same octagon port as shown in Fig. A.4(a), with the leak valve located between the octagon and the RGA. In this geometry, gas introduced through the leak valve produces a local pressure “enhancement” near the RGA. As a result, the pressure at RGA is not a good representation of the pressure the ions experience at the octagon.

To keep the discussion clear, I distinguish between the raw readings from the instruments and the actual pressures. For a gas species g , we define

$$P_{\text{RGA}}^{(g,\text{corr})} = \frac{P_{\text{RGA}}^{\text{raw}}}{C_{\text{RGA},g}}, \quad P_{\text{IG}}^{(g,\text{corr})} = \frac{P_{\text{IG}}^{\text{raw}}}{C_{\text{IG},g}}, \quad (\text{A.3})$$

where $C_{\text{RGA},g}$ and $C_{\text{IG},g}$ are the species-dependent response factors obtained from the calibration described above. With this convention, $P^{(g,\text{corr})}$ is the best estimate of the true partial pressure of the gas species g inferred from the corresponding instrument.




Figure A.3: Measured N_2 response of the ion gauge and the RGA as a function of the Baratron-based reference pressure. The ion-gauge response remains approximately linear to the reference pressure of the full measured range of 10^{-9} to 5×10^{-6} Torr. The RGA begins to deviate systematically from linearity above about 2×10^{-6} Torr.




Figure A.4: Geometry of the experiment used to infer the gas pressure at octagon. (a) CAD drawing of the relevant portion of the main apparatus, showing the relative locations of the RGA, leak valve, octagon, differential-pumping chamber, and ion gauge. (b) Simplified schematics of the two gas-injection configurations used to determine the geometric corrections G_g and L_g . Left shows the direct injection configuration through the leak valve near the RGA, used to measure L_g . Right shows the remote injection through the differential-pumping turbo backing line to measure G_g .

The geometric correction was determined empirically using two separate configurations. First, gas was introduced through the backing line of the differential-pumping chamber turbo pump, as the right panel in Fig. A.4(b). In this configuration, the gas enters the chamber away from the RGA, so the RGA reading is representative of the pressure at the octagon, rather than being dominated by a local source, i.e. the leak valve.³ We therefore define

$$P_{\text{OCT}}^{(g)} = G_g P_{\text{IG}}^{(g,\text{corr})}, \quad (\text{A.4})$$

where G_g is the ratio between the measured pressure at the octagon (in this case with the RGA) and the pressure measured by the ion gauge downstream in the main chamber. For N_2 , we measured

$$G_{\text{N}_2} \approx 8,$$

as shown in Fig. A.5(b). Therefore,

$$P_{\text{OCT}}^{(\text{N}_2)} \approx 8 P_{\text{IG}}^{(\text{N}_2,\text{corr})}. \quad (\text{A.5})$$

Second, gas was introduced directly through the leak valve, as shown in the left column of Fig. A.4(b). In this case, the RGA sees a locally enhanced pressure near the injection point. We parameterize this by

$$P_{\text{RGA}}^{(g,\text{corr})} = L_g P_{\text{IG}}^{(g,\text{corr})}, \quad (\text{A.6})$$

where here, L_g describes the ratio between the corrected RGA reading and the corrected ion-gauge reading with the leak-valve injection. For N_2 , we found


$$L_{\text{N}_2} \approx 80,$$

as shown in Fig. A.5(a).


Combining the two relations of Eqs. (A.4) and (A.6), we can obtain a direct relation between the RGA reading during leak-valve injection and the pressure at the octagon:

$$P_{\text{OCT}}^{(g)} = \frac{G_g}{L_g} P_{\text{RGA}}^{(g,\text{corr})} = \frac{G_g}{L_g} \frac{P_{\text{RGA}}^{\text{raw}}}{C_{\text{RGA},g}} \quad (\text{A.7})$$

³ Here, I implicitly assume that there is no additional conductance between the center of the octagon to the RGA, which may not be completely true, given that the RGA ionizer is on for the measurement.




(a)




(b)

Figure A.5: Empirical geometric calibration for N_2 . (a) Corrected RGA pressure as a function of corrected ion-gauge pressure when N_2 is introduced through the leak valve. The slope defines the local enhancement factor $L_{N_2} \approx 80$. (b) Pressure at the octagon (inferred from the RGA under remote gas injection through differential-pumping chamber) as a function of the corrected ion-gauge pressure. The slope defines the octagon-to-ion-gauge factor $G_{N_2} \approx 8$. These two geometric factors together determine the relation between the RGA reading during leak-valve injection and the N_2 pressure at the octagon.



(a)



(b)

Figure A.6: Same analysis as Fig. A.5, but for H_2 after applying the corresponding species-dependent response factors for both RGA and IG from Table A.1.

A.2.1 Summary of geometric factors

For N_2 , the obtain values of L_g and G_g are:

$$L_{N_2} \approx 80, \quad G_{N_2} \approx 8.$$

For H_2 , the obtained values of L_g and G_g are

$$L_{H_2} \approx 34, \quad G_{H_2} \approx 6,$$

as shown in Fig. A.6. To calculate H_2 pressures in Fig. 4.8, the factors $C_{RGA,g}$ and $C_{IG,g}$ from Table A.1 and L_{H_2} and G_{H_2} from Fig. A.6 are used.

The uncertainty on L_g and G_g is largest for the lowest-pressure points, where both signals are only slightly above baseline.

A.3 Procedure for obtaining the H_2 pressures in Fig. 4.8

Here I explicitly describe how the H_2 pressures plotted in Fig. 4.8 were obtained.

Without any additional injection of H_2 , the RGA measures

$$P_{bg,RGA}^{raw} = 6 \times 10^{-10} \text{ Torr},$$

at mass 2 amu. Assuming this is from hydrogen, the true baseline pressure of H_2 is

$$P_{bg}(H_2) = \frac{P_{bg,RGA}^{raw}}{C_{RGA,H_2}} \approx 2.4 \times 10^{-10} \text{ Torr}.$$

This corresponds to the the first point, i.e. the “baseline” pressure, of H_2 in Fig. 4.8.

When additional H_2 is introduced through the leak valve, only the increase from this baseline is propagated through the geometric correction. In other words, the H_2 pressure at the octagon is inferred from the raw RGA reading as

$$P_{OCT}(H_2) = \frac{P_{bg,RGA}^{raw}}{C_{RGA,H_2}} + \frac{G_{H_2}}{L_{H_2}} \frac{P_{RGA}^{raw} - P_{bk,RGA}^{raw}}{C_{RGA,H_2}},$$

with $C_{RGA,H_2} \approx 2.5$, $G_{H_2} \approx 6$, and $L_{H_2} \approx 34$. For the measurement, the raw RGA readings are

$$2 \times 10^{-9}, \quad 5 \times 10^{-9}, \quad 1.1 \times 10^{-8}, \quad \text{and} \quad 5.3 \times 10^{-8} \text{ Torr}.$$

The corresponding inferred H_2 pressures at the octagon are approximately

$$3.4 \times 10^{-10}, 5.5 \times 10^{-10}, 9.7 \times 10^{-10}, \text{ and } 3.9 \times 10^{-9} \text{ Torr},$$

respectively.

A key caveat is that this procedure assumes the baseline mass-2 signal measured by the RGA is representative of the background H_2 already present in the vacuum system. However, this may not be strictly true, because the active RGA ionizer itself can contribute a hydrogen-rich background through outgassing or desorption. If that is the case, then the first H_2 pressure point in Fig. 4.8 is overestimated, and the true baseline H_2 pressure in the chamber would be smaller than the value plotted.

The injected-gas points are less sensitive to this issue, since the pressure increases until this issue does not have a big contribution to the measurement. This caveat is important when interpreting whether the observed baseline $J = 1$ decay rate is already hydrogen-limited.

Appendix B

Anti-Helmholtz configuration

In this appendix, I derive the relation between the on-axis field derivatives of an anti-Helmholtz coil pair and the frequency shift parameters used in the main text.

In the current-free region near the trap center, the magnetic field satisfies

$$\nabla \times \mathbf{B} = 0, \quad \nabla \cdot \mathbf{B} = 0. \quad (\text{B.1})$$

The magnetic scalar potential is then introduced such that

$$\mathbf{B}(\mathbf{r}) = -\nabla \Phi_B(\mathbf{r}), \quad \nabla^2 \Phi_B(\mathbf{r}) = 0. \quad (\text{B.2})$$

Consider a circular loop of radius R centered at $z = z_0$ with current I and N turns. Along the symmetry axis z , the magnetic field is given by [94, 115]

$$B_z^{\text{loop}}(z, z_0) = \frac{\mu_0 N I}{2} \frac{R^2}{[(z - z_0)^2 + R^2]^{3/2}}. \quad (\text{B.3})$$

Depending on the direction of the current I , the generated magnetic field changes its sign.

For the anti-Helmholtz configuration, we place two identical loops at $z = \pm d$ with opposite currents. The on-axis field from this anti-Helmholtz pair is

$$B_z(z) = \frac{\mu_0 N I R^2}{2} \left[\frac{1}{[R^2 + (z - d)^2]^{3/2}} - \frac{1}{[R^2 + (z + d)^2]^{3/2}} \right]. \quad (\text{B.4})$$

We expand Eq. (B.4) about $z = 0$,

$$B_z(z) = \partial_z B_z|_0 z + \frac{1}{3!} \partial_z^3 B_z|_0 z^3 + O(z^5) \quad (\text{B.5})$$

Derivative at the origin	Value for the anti-Helmholtz pair
$B_z(0)$	0
$\partial_z B_z _0$	$\frac{3\mu_0 N I R^2 d}{(R^2 + d^2)^{5/2}}$
$\partial_z^2 B_z _0$	0
$\partial_z^3 B_z _0$	$\frac{15\mu_0 N I R^2 d(4d^2 - 3R^2)}{(R^2 + d^2)^{9/2}}$

Table B.1: On-axis magnetic-field derivatives at the center of an anti-Helmholtz pair consisting of two identical circular coils of radius R located at $z = \pm d$ with opposite currents.

The derivatives of Eq. (B.4) are listed in Table B.1. As expected from symmetry, $B_z(z)$ is odd in z , and all even derivatives vanish at the origin.

Plugging in these derivatives into the Taylor expansion, we obtain

$$B_z(z) \approx \left[\frac{3\mu_0 N I R^2 d}{(R^2 + d^2)^{5/2}} \right] z + \frac{1}{3!} \left[\frac{15\mu_0 N I R^2 d(4d^2 - 3R^2)}{(R^2 + d^2)^{9/2}} \right] z^3 + O(z^5). \quad (\text{B.6})$$

We now relate the two leading on-axis derivatives to the rotating-frame effective magnetic field. Let

$$L \equiv \partial_z B_z|_0, \quad C \equiv \partial_z^3 B_z|_0. \quad (\text{B.7})$$

The linear term in the anti-Helmholtz field is

$$\mathbf{B}^{(1)}(\mathbf{r}) = -\frac{L}{2}x\hat{\mathbf{x}} - \frac{L}{2}y\hat{\mathbf{y}} + Lz\hat{\mathbf{z}}. \quad (\text{B.8})$$

from $\nabla \cdot \mathbf{B} = 0$ and axial symmetry.

Projecting this field to the rotating quantization axis and averaging over one rotation cycle, we get

$$B_{\text{eff}}^{(1)} = \left\langle \mathbf{B}^{(1)}(\mathbf{r} + \mathbf{r}_{\text{rot}}(t)) \cdot \hat{\mathbf{r}}_{\text{rot}}(t) \right\rangle_t = -\frac{L}{2}r_{\text{rot}}. \quad (\text{B.9})$$

As expected, the leading anti-Helmholtz gradient produces a common-mode rotating frame magnetic field proportional to $\frac{1}{2}\partial_z B_z|_0$, which we call B_{axgrad} .

The next contribution comes from the cubic term. The axially symmetric field for the cubic

field can be written near the origin as

$$\begin{aligned} B_x^{(3)}(\mathbf{r}) &= \frac{C}{16}x(x^2 + y^2 - 4z^2). \\ B_y^{(3)}(\mathbf{r}) &= \frac{C}{16}y(x^2 + y^2 - 4z^2). \\ B_z^{(3)}(\mathbf{r}) &= \frac{C}{6}z^3 - \frac{C}{4}z(x^2 + y^2). \end{aligned} \tag{B.10}$$

which satisfy Eq. (B.1). Along the axis, we get

$$B_z^{(3)}(0, 0, z) = \frac{C}{6}z^3, \tag{B.11}$$

as required by the Taylor expansion in Eq. (B.6).¹

Performing the rotation averaging to Eq. (B.10), we get

$$B_{\text{eff}}^{(3)}(\mathbf{r}) = \frac{Cr_{\text{rot}}}{16}[r_{\text{rot}}^2 + 2(x^2 + y^2 - 2z^2)]. \tag{B.12}$$

The second term is quadratic in the secular coordinates and therefore produces the spatially varying frequency shift of interest.

In the main text, I parameterize the induced frequency shift as

$$f(x, y, z) = f_0 + f_2 \left(-\frac{x^2}{2} - \frac{y^2}{2} + z^2 \right). \tag{B.13}$$

Therefore, comparing Eqs. (B.12) and (B.13), we obtain

$$\begin{aligned} \frac{f_0^{(20)}}{I} &= \frac{1}{2} \frac{1}{I} \partial_z B_z|_0 \times \left(\frac{50 \text{ Hz}}{3 \text{ mG/cm}} \right), \\ \frac{f_2^{(40)}}{I} &= \frac{1}{4} \frac{1}{I} \partial_z^3 B_z|_0 \times \left(\frac{50 \text{ Hz}}{3 \text{ mG/cm}} \right), \end{aligned} \tag{B.14}$$

using the conversion factor for $f_{\text{rot}} = 147.5 \text{ kHz}$ and $\mathcal{E}_{\text{rot}} = 60 \text{ V/cm}$. Here, $f_0^{(20)}$ denotes the common-mode shift induced by the leading anti-Helmholtz gradient, and $f_2^{(40)}$ denotes the quadratic curvature induced by the next axially symmetric term.

¹ The linear and cubic terms L and C are the fields associated with the axially symmetric $l = 2, m = 0$ and $l = 4, m = 0$ magnetic multipoles, respectively in the notation used in the main text.

Appendix C

Timing code

In this appendix, I describe the code used to find short recurrence times near specified values of f_{rot} and f_{rf} and the corresponding f_{laser} , which is the triggering rate for the pulsed lasers.

As discussed in the main text, for a DDS with clock frequency f_{clk} and n number of bits, the frequency spacing Δf_{DDS} and the output frequency f_i are given by

$$\Delta f_{\text{DDS}} = \frac{f_{\text{clk}}}{2^n}, \quad f_i = K_i \Delta f_{\text{DDS}},$$

where K_i is an integer tuning word.

For the rotating-field and rf frequencies, we define

$$K_{\text{rot}} = g \, p, \quad K_{\text{rf}} = g \, q, \quad (\text{C.1})$$

with

$$\text{gcd}(p, q) = 1, \quad g = \text{gcd}(K_{\text{rot}}, K_{\text{rf}}).$$

Therefore, p/q is the reduced ratio

$$\frac{f_{\text{rot}}}{f_{\text{rf}}} = \frac{K_{\text{rot}}}{K_{\text{rf}}} = \frac{p}{q},$$

as mentioned in the main text. Defined this way, the integer g is the greatest common divisor of the two DDS tuning words.

For specified tolerances around the target frequencies of $f_{\text{rot}}^{(0)}$ and $f_{\text{rf}}^{(0)}$, the allowed integer turning word ranges from $K_{\text{rot},\text{min}}$ to $K_{\text{rot},\text{max}}$ defined as

$$K_{\text{rot},\text{min}} = \left\lceil \frac{f_{\text{rot}}^{(0)} - \delta f_{\text{rot}}}{\Delta f_{\text{DDS}}} \right\rceil, \quad K_{\text{rot},\text{max}} = \left\lfloor \frac{f_{\text{rot}}^{(0)} + \delta f_{\text{rot}}}{\Delta f_{\text{DDS}}} \right\rfloor,$$

and similarly for $K_{\text{rf},\min}$ and $K_{\text{rf},\max}$.

Then, the code loops through candidate reduced ratios $p : q$. For each reduced ratio considered by the code, the allowed values of g must satisfy

$$K_{\text{rot},\min} \leq gp \leq K_{\text{rot},\max}, \quad K_{\text{rf},\min} \leq gq \leq K_{\text{rf},\max},$$

then, the range of g considered by the code corresponds to

$$g_{\min} = \left[\min\left(\frac{K_{\text{rot},\min}}{p}, \frac{K_{\text{rf},\min}}{q}\right) \right], \quad g_{\max} = \left[\max\left(\frac{K_{\text{rot},\max}}{p}, \frac{K_{\text{rf},\max}}{q}\right) \right].$$

Since we are interested in the shortest recurrence time, g is set to the maximum possibility, i.e. $g = g_{\max}$.

With this g , the recurrence time is simply

$$T_{\text{recurrence}} = \frac{1}{g\Delta f_{\text{DDS}}} = \frac{2^n}{f_{\text{clk}}g}.$$

As discussed in the main text, this suggests that during one recurrence time, the rotating field completes p cycles and the rf field completes q cycles and therefore the phase pair $(\phi_{\text{rot}}, \phi_{\text{rf}})$ returns to the original value.

The code also calculates a laser triggering frequency close to the desired value $f_{\text{laser}}^{(0)}$:

$$f_{\text{laser}} = K_{\text{laser}}\Delta f_{\text{DDS}}.$$

In order to make the laser trigger commensurate with the recurrence time, we require that g/K_{laser} to be an integer. Then,

$$f_{\text{laser}} = K_{\text{laser}}\Delta f_{\text{DDS}} = \frac{g}{M}\Delta f_{\text{DDS}} = \frac{1}{MT_{\text{recurrence}}}.$$

Therefore, the laser trigger period is an integer multiple of $T_{\text{recurrence}}$, and the laser pulses do not accumulate phase slip relative to either f_{rot} or f_{rf} . In the code, I search over possible integer values of K_{laser} that is closest to the desired laser repetition rate.¹

¹ Here, I implicitly assumed that the laser trigger comes from the same instrument with the same f_{clk} and n , which will be the case in the Gen. III experiment, with the updated DDS system.

```

import math
import pandas as pd

""" User inputs """

f_clk = 125e6          # DDS clock frequency, Hz
nbits = 32            # DDS bits

set_f_rot = 300e3      # set f_rot, Hz, nominally 300 kHz
set_f_rf = 100e3       # set f_rf, Hz, nominally 100 kHz
set_f_laser = 10.0     # set f_laser, Hz, nominally 10 Hz

f_rot_tol = 5000.0     # allowed tolerance from set_f_rot, Hz
f_rf_tol = 5000.0      # allowed tolerance from set_f_rf, Hz

n_result = 30          # number of rows to print

max_term_integer = 100 # for p/q ratio, max integer allowed for p or q
exclude_integer_ratio = True # True if want to avoid p/q=integer

""" Calculate DDS settings """
df_Hz = f_clk/2**nbits # DDS frequency resolution

### calculate minimum and maximum tuning words for f_rot, given f_rot_tol
K_rot_min = math.ceil((set_f_rot-f_rot_tol)/ df_Hz)
K_rot_max = math.floor((set_f_rot+f_rot_tol)/ df_Hz)

K_rf_min = math.ceil((set_f_rf-f_rf_tol)/ df_Hz)
K_rf_max = math.floor((set_f_rf+f_rf_tol)/ df_Hz)

""" Calculate candidates """
results = []

for p in range(1, max_term_integer+1):
    for q in range(1, max_term_integer+1):

        # if a and b are not reduced ratio, stop and skip
        if math.gcd(p,q) != 1:

```

```

        continue

# skip integer ratios p/q
if exclude_integer_ratio and q==1:
    continue

# calculate minimum and maximum g, given this ratio p:q
g_min = math.ceil(max(K_roat_min/p, K_rf_min/q))
g_max = math.floor(min(K_rot_max/p, K_rf_max/q))

# if no allowed g, stop and skip
if g_min>g_max or g_max <=0:
    continue

# finding the maximum g, which is shortest recurrence time
g = g_max

# given g, calculate f_rot, f_rf and T_rec according to definitions
K_rot = g*p
K_rf = g*q

f_rot = K_rot*df_Hz
f_rf = K_rf*df_Hz

T_rec = 2**nbits/(f_clk*g)

# given T_rec, find f_laser such that K_laser divides g and f_laser is closest to set_f_laser
K_laser_target = round(set_f_laser/df_Hz)
K_laser_min = 1
K_laser_max = min(g, 2*K_laser_target) # factor of 2 makes sure it also looks at f_laser >10 Hz

best_K_laser=None
best_f_laser=None
best_error=float("inf") # set to inf initially

for K_laser_try in range(K_laser_min, K_laser_max+1):
    if g % K_laser_try != 0: # if K_laser_try does not divide g, it is wrong, keep looking
        continue

```

```

f_laser_try = K_laser_try*df_Hz
error = abs(f_laser_try-set_f_laser)

if error<best_error:                # loop through M_try until the error is minimized
    best_error=error
    best_K_laser = K_laser_try
    best_f_laser = f_laser_try

K_laser = best_K_laser
f_laser = best_f_laser
M = g // K_laser

# now append everything
results.append({
    "T_rec_us": T_rec*1e6,           # T_reccurent in us
    "M": M,                          # integer for f_laser
    "f_laser_Hz": f_laser,           # set f_laser to be commensurate with T_rec
    "error_f_laser_Hz": abs(set_f_laser-f_laser),

    "K_rot": K_rot,                 # Tuning word for f_rot
    "K_rf": K_rf,                   # Turning word for f_rf
    "K_laser": K_laser,              # Tuning word for f_laser
    "f_rot_Hz": f_rot,               # f_rot in Hz
    "f_rf_Hz": f_rf,                # f_rf in Hz

    "gcd": g,                       # gcd of f_rot and f_rf
    "reduced_ratio": f"{p}:{q}",     # reduced ratio between f_rot:f_rf
})

""" Sort by shortest recurrence time """
df = pd.DataFrame(results)

df = df.sort_values(
    by = ["T_rec_us", "error_f_laser_Hz"], ascending=[True, True],).reset_index(drop=True)

cols = ["T_rec_us", "M", "f_laser_Hz", "error_f_laser_Hz", "K_rot", "K_rf",
        "K_laser", "f_rot_Hz", "f_rf_Hz", "gcd", "reduced_ratio"]

```

```
print(df[cols].head(n_result).to_string(index=False))
```

Listing C.1: Python code used to search for commensurate DDS frequencies.